

TEACHING GAME THEORY AND RATIONALITY TO ARTIFICIAL INTELLIGENCE MASTER'S STUDENTS

J. Jordán, V. Julián

Universitat Politècnica de València (SPAIN)

Abstract

Game theory is a branch of applied mathematics used in economics and artificial intelligence to study the behaviour of self-interested agents. In game theory, agents are rational, which means that they will always analyse the situation intelligently and seek their own benefit. The concept of rationality in game theory and artificial intelligence is based on the rigorous analysis of a conflict between self-interested agents in which they exclusively seek to maximise their utility in the defined terms of the game. Are we humans totally rational in making our decisions? Do we consider all the possible options to get the maximum benefit or do we let ourselves be influenced by other factors such as feelings?

In this paper, we present the methods used with students of the Master's Degree in Artificial Intelligence at the Universitat Politècnica de València to teach game theory and rationality. In order to ensure that the students were able to assimilate the concepts of game theory, the classic master class method has been combined with new technologies to make the lesson dynamic. In this way, the explanation of theoretical concepts has been mixed with small games carried out with the students, always guided by the professor, using tools such as Kahoot and Socrative.

The experience acquired during the lessons shows that even in the case of students coming from engineering backgrounds with a solid mathematical base, in general, they still lack this abstract understanding of rationality and act according to their feelings in many cases. This invites reflection and attempts to provide students with guidance for the correct understanding of the concept of rationality along with other complex concepts typical of game theory. On the other hand, we also show the results obtained from a series of games played in the classroom that demonstrate this lack of rationality in the students, what obviously makes them human. This can be interpreted as the fact that the utility functions of the agents that represent humans in a game do not capture all the variables, such as feelings, that a person takes into account when making certain types of decisions.

The evaluation of the knowledge acquired by the students was done through an open project that consisted in designing a game including the rules of the game, applicable concepts of equilibrium, and analysis of the rationality of the players. This allowed the assessment of the knowledge acquired by the students during the lessons carried out with the dynamic methodology obtaining very satisfactory results. In this sense, the academic results of the students regarding the average grades obtained were higher than in previous courses in which this methodology was not applied. On the other hand, student satisfaction was also high, as shown by the surveys carried out.

Keywords: Artificial intelligence, game theory, rationality.

1 INTRODUCTION

Learning the basic concepts of the branch of applied mathematics known as game theory [1] can certainly be complex, so teaching these concepts is a pedagogical challenge. This also depends on the context and the background of the students. In the case we are studying here, the students belong to the Master's Degree in Artificial Intelligence, so they have a significant mathematical background, although they have rarely worked with game theory concepts, particularly rationality.

A game represents a situation, specified based on rules, in which its participants have a series of strategies to apply, where each of these strategies provides the agent with a utility (generally benefit minus cost) depending on the strategies of the other agents. Thus, an agent acts in its own interest trying to maximise its utility, but it must consider the strategies of the others to do so. In general, the solution to a game will be an equilibrium, i.e., a situation where no agent can benefit by changing its strategy unilaterally, in which case all agents are said to be in best response (considering the strategies applied by the other agents) [2].

However, for all this to be accomplished, it is assumed that the agents are rational, that is, that they analyse the situation to seek their own benefit. The concept of rationality in game theory and artificial intelligence is based on the rigorous analysis of a conflict between self-interested agents in which they exclusively seek to maximize their utility in the defined terms of the game. Nevertheless, we humans are not usually totally rational as other factors can intervene, usually feelings, sensations, or intuitions, which prevent us from acting coldly just by assessing the direct benefit or prejudice of the situation represented by a game in the context of game theory. Regarding this issue, we could say that when this happens it is because the utility functions of the game do not capture all the variables involved in the decisions. In fact, the rationality of human is a problem studied by the psychology of reasoning [3], which is out of the scope of artificial intelligence and the scope of this work.

Therefore, the methods used to explain for the first time the concepts of game theory and the analysis of how students react to the concept of rationality throughout their learning is studied in this paper. Firstly, Section 2 explains the didactic methodology used during the lessons to explain the theoretical content using traditional lecture techniques combined with dynamic elements to make small example games typical of game theory. In addition, it also explains the works proposed to evaluate the students. Then, in Section 3 the results obtained in the games made during the classes are analysed together with the results of the work made by the students and their evaluation. Finally, Section 4 summarises the main conclusions drawn from this work.

2 METHODOLOGY

This section explains the methodological approach followed for the Automated Negotiation course of the Master's Degree in Artificial Intelligence at the Universitat Politècnica de València. Firstly, it is commented how the typical methodology of a theoretical class has been used. Then, the techniques used to support the theoretical concepts through dynamic games in the classroom are explained. Finally, the use of open projects as individual student work to evaluate the knowledge acquired is discussed.

2.1 Theoretical Classes

The teaching of certainly complex theoretical concepts such as game theory is difficult to separate from the support of a professor. While the use of flip teaching techniques for a first contact with such theoretical concepts may be debatable, it seems advisable if these are explained by a professor who can exemplify such contexts and get direct feedback from students as to whether they are understanding the subject matter. For these reasons, the classic theory class is chosen.

It should be noted that these theory lessons are well supported by slides and well-prepared explanations with small examples so that the whole set is very didactic. However, to complement the explanations of these complex concepts and at the same time stimulate the students' minds, it is necessary to include dynamic tools for this purpose, as we will see below.

2.2 Dynamic Games in the Classroom

In recent years, the arrival of new technologies in the classroom has become much more evident, and even more so in a university context, where more technological resources are usually available per classroom. In the context that concerns us, classrooms are even equipped with a computer per student, although most of them also have their own laptop or smartphone. All of this means that the teaching task can be easily supported by all this technology. Specifically, for classes in which game theory is discussed, it has been decided to use web technologies with extensive possibilities.

Firstly, the use of the web tool Kahoot! [4] is proposed, as it allows questions to be asked that are displayed both on the classroom projector and on the students' devices (either a computer or a smartphone). This provides a dynamic environment in which students can easily and quickly answer questions. In addition, the results of each question can be viewed on the fly for analysis, and the results can be saved for future purposes. Thus, this kind of tool is useful for the kind of dynamism that can be introduced in a game theory class, i.e., to provide an example game that students can answer on the fly, as will be seen in Section 3.1.1.

Secondly, the use of the web tool Socrative [5] is proposed. In this case, Socrative offers the possibility to create questions in which students can enter more complex written answers. In addition, the questions can also be launched online, and the professor can consult the students' results on the

fly and even have the results saved for future analysis. For these reasons it is proposed to use this tool to make more complex instances of games in which the information generated by the students can be collected, as will be seen in Section 3.1.2.

Finally, in some cases it is not necessary to have new technologies to obtain satisfactory and dynamic results. Thus, for certain types of games the most useful thing that can be done to help students get into the role and to know exactly what role their peers are playing may be a post it note. These techniques can also work well if used in the right context and also to break the sometimes overwhelming use of technology, especially in degrees related to computer engineering. An example of this is discussed in Section 3.2.

2.3 Open Project

The evaluation of the knowledge acquired by the students can be carried out in different ways depending on the context. In our case, an evaluation is proposed by means of an open topic project. Since the aim of the game theory lessons is that students understand the main concepts so that they can use them in Automated Negotiation, it is considered that being able to choose a topic that appeals to the student is more positive for their learning. It should be noted that we are dealing with master's students, so they already have a lot of previous experience and acquired responsibility.

This type of project offers the advantages of giving students freedom, which has repercussions on their interest in the subject and, therefore, on their understanding of the concepts. In addition, this type of open topic work makes students more motivated and the quality of the work along with the grades they can obtain is much higher. However, these assignments often require more advice from the professor, as well as a slightly more complex evaluation. However, it all depends on the number of students and the time available for completion. Thus, the advantages are considered to outweigh the disadvantages. The outcome of the proposal for this type of work is discussed in Section 3.3.

3 RESULTS

This section presents the results obtained during the classes of the Automated Negotiation course of the Master's Degree in Artificial Intelligence at the Universitat Politècnica de València. Firstly, the experience and results of the ultimatum game carried out in the classroom are discussed. Next, the development of the pirate game during the classes is also summarised to observe the evolution of the students with respect to rationality in game theory and to be able to explain the concept of equilibrium. Afterwards, the project assigned to the students is explained with which the evaluation of their knowledge is carried out, commenting also on the results obtained.

3.1 The Ultimatum Game

The ultimatum game [6] is usually used as a prototypical example to explain concepts of rationality focused on both game theory and negotiation. This game consists of two players, the proposer, and the responder. There is an amount of money to be distributed between the two players. The peculiarity of the game is that the proposer decides the amount he assigns to himself and the amount he assigns to the responder. This amount can vary between 0 and the total amount of money available. Once the proposal has been made by the proposer, the responder can either accept, so that both players keep the amounts proposed, or reject the proposal, so that neither player receives anything.

From a mathematical and rational point of view, any amount proposed by the first player should be accepted by the responder, except for the amount 0 (for the responder), which could be accepted or not since saying yes or no is the same for the responder, because his payoff is 0 in both cases. Therefore, any amount higher than 0 means a gain for the responder. This, although it may seem obvious when we explain it from the pragmatic point of view of the matter, does not occur naturally in most human beings. Therefore, before explaining rationality using this game, we propose that students play it with the professor to analyse their initial thinking and their knowledge or intuition about rationality.

3.1.1 *Playing the Ultimatum Game versus the Professor*

The first activity consists of playing the ultimatum game with the professor being the proposer, and each of the students plays the role of an individual responder using the Kahoot! tool. The advantage of using Kahoot! in the classroom (or in online classes) lies in the ease of asking questions electronically and getting all the answers from the students on the fly and being able to analyse the results on site or

later. In addition, the interactivity and dynamism of this tool encourages students to participate. In this experiment, students are asked 6 questions consisting of various divisions of the money available which are decreased for each question. So, it is like playing the ultimatum game 6 times with different proposals for the division of money, emphasizing that each game is independent from the others (although in the human mind this separation is unlikely to exist completely). In this case, the amount to be divided is 100, and no proposal is necessary in which the responder receives more than the proposer, since it would always be accepted and it is evident, although it is not superfluous to corroborate it with the students verbally.

In Fig. 1 we can see the percentage of students who accept each of the money divisions proposed by the professor. The first division consists of 51 euros for the professor and 49 euros for the student and is accepted by all the students. The second division, consisting of 40 euros for the student, is accepted by almost 80% of the students. In this case we already see how a part of the students do not consider the distribution fair and prefer to remain without benefits if they do the same to the proposer. The third distribution of 75€ for the proposer and 25€ for the responder still has an acceptance rate close to 60%, which shows that a majority is still able to suppress other feelings and focus on the benefit they get. The next distribution of 90 and 10 euros respectively for proposer and responder is accepted by about 40% of the students. In this case, it already becomes clear that for most students the incentive is certainly low, and they are led to "punish" the proposer by making both lose out. The fifth distribution is already at the limit as it gives 99 euros to the proposer, and 1 to the responder, so that only 20% of the students accept. This percentage of students might be said to have either natural notions of rationality, or a mentality cold enough to know that it still compensates them to make a profit. Finally, the last division consists of 100 for the proposer and 0 for the responder, and obviously no student accepted, however, there were several students who marked the option "perhaps", so it could be corroborated that these students already had previous training in game theory concepts.




Figure 1. Ultimatum game results professor vs. student (individual game through Kahoot!). Percentage of acceptance of the students to different divisions of the total payoff.

The realization of this game serves as an introduction to the concept of rationality of game theory and as a principle of negotiation. From the data, students can be invited to reason about the convenience of accepting something that gives us some benefit rather than keeping nothing. This, from a game theory perspective, is the basis for understanding rationality, always looking for one's own benefit within the current possibilities even if other players get more.

Furthermore, it is interesting to consider what amounts could be accepted when a person plays the ultimatum game with a stranger. On the one hand, with a small amount of about 100 euros, most people would not accept less than 30 euros because they find it abusive from the other player [7][8]. However, the division that anyone will accept varies significantly if we drastically increase the amount of money to be distributed [9]. Consider that if the amount to be distributed was 100 million euros most people (who previously had no more than several million) would accept even if they were offered only 1 million.

3.1.2 *Playing the Ultimatum Game versus a Classmate*

The next activity before continuing with more theoretical content is for students to play the ultimatum game in pairs. In this case, they choose at random who will be the proposer and who will be the responder, and they play ten rounds of the game in a row while maintaining the role. The amount of money to be distributed is again 100€. In addition, students are asked to write down the result, i.e., the utility obtained in each round to calculate the average obtained from each of them. It should be remembered that in cases where the responder rejects the proposal, the utility for both is 0. In addition, the results of the average obtained by each pair are noted down in an activity launched by the professor on the Socrative platform to be able to analyse them globally with the students at the end of the activity.

The intention of this second activity with the same ultimatum game is that students have more time to think about what kind of offers they can offer (proposer) and which ones they are willing to accept (responder). In this situation students deliberate a bit more on how they react to different offers. All this gives rise to a better explanation of rationality afterwards and makes them see that in many cases they have let themselves be carried away by feelings that have led to a decrease in their average utility.

Table 1. Average utility obtained by both roles in 10 rounds of the ultimatum game for different student profiles.

	<i>Proposer</i>	<i>Responder</i>
<i>Profile 1</i>	43,8	26,2
<i>Profile 2</i>	39,5	30,5
<i>Profile 3</i>	35,3	14,7

The results obtained by the students can be seen in Table 1 which shows the average utility of the ten rounds obtained by each role. From all the pairs formed by the students, three profiles have been chosen which have been generally repeated.

Firstly, the first profile corresponds to a responder who usually accepts most of the offers since the average utility of both players (especially the first one) is relatively high. However, even for the proposer, this does not reach 50, which shows that there have been several times in which the responder has rejected offers, thus lowering the average of both. In addition, the responder has an average that indicates that the majority of offers he will have accepted would be around 30-40. The second profile has a more balanced average utility between both players, indicating that the offers that were accepted were probably fairer. However, as the proposer's average is around 40, the responder also rejected some of the offers. Finally, the third profile could be said to have been the most aggressive; it is not possible to know whether the proposer or the responder, since the average utility of both is relatively lower than that of the other profiles. This may have been caused either by too unfair offers from the proposer, or by a nonconformist responder who has rejected the offers on many occasions. In any case, what has caused these behaviours is much less utility. In view of all the results in Table 1, it should simply be asked whether the greed of the proposers, and the non-conformity of the responders, has not also caused this decrease in average utility. Simply making offers 50 | 50 or close to this number as 60 | 40 would have made the responder always say yes, with both players getting a much higher average utility. On the other hand, if we analyse it from a strictly rational point of view, the responder should have accepted any offer always above 0, which could have caused the proposer to have a very high average utility, and the responder only what the proposer's goodwill would have left him. However, in an iterative game like this, it is not advisable either to accept all offers, especially those that are very unfair, because they only cause the proposer to be more and more greedy.

All this proves that the behaviour of the players changes when a set of rounds are played, since the previous rounds influence the following ones. In this sense, a game that is played only once, where we should accept anything above 0, is not the same as a game that is repeated, where other factors come into play and we are interested in the average utility (or the total sum) that we obtain. So, we can say that the responder can have some control over the proposer by rejecting the less fair offers, which influences the proposer to make fairer offers so as not to obtain 0 utility.

Therefore, in iterated games the idea of strict rationality is certainly diluted since there are other factors when repeating the game several times, such as strategic behaviour. This whole discussion is

a great starting point to explain the concept of rationality and see that it also depends on the context and is not just a strictly mathematical concept. Furthermore, this also helps students to become interested in the subject and to ask themselves other questions about game theory.

3.2 The Pirate Game

A very interesting game that serves to illustrate the contrariness of some equilibria is the pirate game [10], which can be considered a multi-player version of the ultimatum game. Let us remember that an equilibrium is a situation in a game in which all the players are at their best response when considering the strategies of the other players, which means that a unilateral deviation of any player can only mean utility losses. Once the concept of equilibrium has been explained and reasoned with the students, even using small examples of more typical and didactic games for this purpose, it is interesting to use the pirate game to see the reactions of the students and make them see that it is difficult to have a global vision of a game when they only control their own decisions and do not think much about the possible decisions of others. This also happens because, as we have already mentioned, people's rationality can be relative.

In the pirate game there are five rational pirates (in strict order of seniority A, B, C, D, and E) who must decide how to distribute 100 gold coins. The most senior pirate proposes a distribution, and all pirates including the proposer vote on whether to accept this distribution. If the majority accepts, they get the coins and the game ends. In case of a tie vote, the proposer has the casting vote. If the majority rejects the plan, the proposer is thrown overboard from the ship and dies, hence the next most senior pirate makes a new proposal to start the system again. The process is repeated until a distribution is accepted or if there is one pirate left. Obviously, each pirate wants to survive and maximize the gold coins he receives. The pirates do not trust each other so they will never make any honour promises.

In the classroom, it is possible to make groups of 5 to play different rounds with each other and then comment on the results globally, or we can directly select 4 volunteers, reserving the main role to the professor, and play several rounds. In any case, the aim is to get the possible results in different rounds, showing that none of them is the equilibrium, or that it is not the best for most players. The particularity of this game is that the equilibrium is somewhat strange since the distribution would be 98 coins for player A, 0 for B, 1 for C, 0 for D, and 1 for E. This would be the equilibrium since if there were only D and E left, the vote of more value would be for D and he would win, so E does not want to end up in this situation. If there are C, D, and E left, player C knows that D will offer E nothing in the next move, so by offering E 1 coin he should already have his vote. In the same way, if in addition there were still B, just by offering 1 coin to D he could already count on his vote, since if he stays only with C and E he would receive 0. So, knowing all this, A knows that he can count on the support of C and E to win the game, so he would offer the minimum to each of them, a gold coin, and he would stay for 98.

So, in class this game was played, guided by the professor, and on none of the occasions the described equilibrium was reached. This happened because the students did not have such an extreme rationality from a mathematical point of view, and what happened most of the time was that all the players (except one, or sometimes two) ended up with a result of 0. However, it is a positive experience because when trying to provoke the students to reach the equilibrium (influencing and explaining their decisions), they ended up realising the importance of rationality from the point of view of game theory. Furthermore, they did not lose the focus of another important lesson, that the rationality presented in game theory does not always fully respond to human rationality as it is influenced by other factors not captured by the utility functions of the games themselves.

Finally, it should be noted that playing a game of this type is a very fun and positive experience for the students, highly recommended as a closure of some lesson.

3.3 The Final Project

The students chose different topics that interested them within the context of problem solving and analysis with game theory, which resulted in an interesting variety of work. The analyses carried out by each of the students demonstrated the skills acquired during the theoretical and practical lessons together with the autonomous work.

In general, the evaluation that was made to most of the students was higher than 8.5 out of 10 since the level of the works presented was excellent. In addition, some of these works were a little completer and more exhaustive, which gave them the highest mark.

It should also be noted that this experience gave such good results because they had the support of the professor through tutorials that were used by most students. This meant direct advice from the professor both in the choice of work to be done and in its development. With this, the students felt more confident about the decisions they made about the direction of their work and about the results and considerations they obtained.

Finally, the conclusion to be drawn from this type of open work is that the results at the overall qualification level are superior and the learning and mastery of the subject by the students is also superior. On the other hand, the only drawback is that it requires much more time from the professor when offering this individual assessment, which may make it unfeasible depending on the time available or the number of students.

4 CONCLUSIONS

This paper has presented the methodology used in the Automated Negotiation classes of the Master's Degree in Artificial Intelligence at the Universitat Politècnica de València to teach game theory and its main concepts. The traditional master class methods have been combined with dynamic methods using online tools to make small example games that served to settle the desired concepts in the students. This meant helping students to analyse these concepts for a better understanding, as well as obtaining interesting results on their idea of rationality, which in most students was influenced by other feelings. On the other hand, a final project was proposed with which the knowledge acquired by the students was evaluated and they obtained higher final grades than in previous years.

Finally, the results of student satisfaction obtained in the student opinion surveys were very good, also improving on those obtained in previous years. This shows that this type of teaching, although it requires more direct attention from the professor in the case of the open work project and tutoring, compensates since student satisfaction is notably better. Furthermore, this is one of the objectives of teaching the Master's in Artificial Intelligence and the subject of Automated Negotiation to continue attracting students in subsequent years.

ACKNOWLEDGEMENTS

This work was partially supported by MINECO/FEDER RTI2018-095390-B-C31 project of the Spanish government.

REFERENCES

- [1] J. Von Neumann, and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton University Press, 2007.
- [2] Y. Shoham, and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.
- [3] J. P. Leighton, "Defining and describing reason" in *The Nature of Reasoning* (J. P. Leighton, and R. J. Sternberg, eds.). Cambridge University Press, 2004.
- [4] Kahoot!, Accessed 12 January, 2021. Retrieved from <https://kahoot.com/>
- [5] Socrative, Accessed 12 January, 2021. Retrieved from <https://www.socrative.com/>
- [6] J. C Harsanyi, "On the Rationality Postulates underlying the Theory of Cooperative Games", *The Journal of Conflict Resolution*, 5 (2), 179–196, 1961.
- [7] J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, and H. Gintis, *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford University Press, 2004.
- [8] H. Oosterbeek, R. Sloof, and G. van de Kuilen, "Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis". *Experimental Economics*. 7 (2), 171–188, 2004.
- [9] S. Andersen, S. Ertaç, U. Gneezy, M. Hoffman, J. A. List, "Stakes Matter in Ultimatum Games", *American Economic Review*, 101 (7), 3427–3439, 2011.
- [10] I. Stewart, "A Puzzle for Pirates", *Scientific American*, pp. 98–99, 1999.