A Family of Turn-Based Strategy Games with Moose

Daniel Ashlock, Joseph Alexander Brown, Connor Gregor, and Munir Makhmutov

***Abstract*—The game that is the focus for this study is a simple turn-based strategy game used for competition in an AI class at Innopolis University for two years. In this study, we continue to investigate the game with evolved game-playing agents. The space of the moose competition games is expanded and the hypothesis that changing the resource profiles of the game will lead to the evolution of different types of agents is tested. The fields that the moose forage in may have enhanced rate of plant growth or enhanced richness of plant biomass. Increasing these parameters is found to decrease conflict between moose, after a minimum level is passed, and enhancing richness is found to modify moose behavior in intuitive ways. In some parts of the resource model space, the agents exhibit negative density-dependent selection that increases the rate of confrontation. This study demonstrates that the family of moose games provides a rich palette of different games for testing students and agent AIs.**

I. Introduction

Moose can be quite aggressive, though the bulk of this aggression is targeted towards their own species; While living in the wild, moose make use of their horns in order to fight over potential mates and also ward away other moose from what the moose considers to be their territory [10]. The game that is used as the focus of this study places two moose in competition for the resources available in three foraging areas referred to as fields; each moose attempts to graze as much as they can while behaving without regard for the needs of the other moose. The game occurs over discrete time steps; at a time step, each moose is given the option to choose one of the three forage area fields and then attempt to gather the resources that have accrued there. Each forage area uses a logistic growth model in order to determine the rate at which they replenish resources over many time steps. Once each moose has chosen a field, the fields that have no more than a single moose upon them increase the amount of resource that they contain. If the moose finds themselves to be alone, they take all of the resource in the forage area for themselves, reducing it down to a baseline level. In the case where the two moose both chose the same foraging area, they threaten one another, tear up the area, waste some of the resource, and gain none of it for themselves. Since untouched fields accrue extra resources, this game creates a type of competition that contains the potential for cooperation. Should the moose forage in separate areas, then they gain access to forage resources without wasting any of them. The most amicable version of this competition would have each moose remain within one forage area – this totally avoids competition, but also ensures that each moose gains the minimum single time step value from the logistic growth model; the moose can add these minimal values to their fitness while leaving the third forage area alone. This third area forms a strong strategic temptation over time as it continues to compile more and more vegetation upon its field. Since the moose that attain the most forage are the ones with the best reproductive success in our model, the strategy of picking a field and sticking with it is a poor one. That third forage area calls out to be exploited. In this aspect, the moose game is similar to the classic prisoner's dilemma [11]; the similarity is drawn as mutual exploitation of the third field ends up being detrimental to both moose as they waste a time step fighting and ruin the resource rich field. Meanwhile, exploitation committed only by one moose grants them a considerable fitness advantage over the other moose.

The moose game simulates interactions between two moose. The environment in which the moose live contains three fields in which they can forage. In these fields, various vegetation grows in a sigmoid fashion, using a logistic curve [16] given by Equation 1.

$$f(x) = \frac{Ce^x}{1 + e^x} \tag{1}$$

Where C is a parameter that creates an upper bound on the amount of resource that can be contained within a field should it remain undisturbed indefinitely. If the field has no accessible vegetation at all, then the value of the vegetation is given by $f(0)$. It is assumed that the initial value of x for all fields is $x = 1$ as each field has some vegetation prior to the arrival of the moose. If no more than a single moose visits a field, its vegetation continues to grow; if one moose is in the field, then they harvest the available vegetation, immediately resetting x to zero. When a moose harvests the resource of field k, the moose eats the vegetation based on the amount available and adds $f(x_k) - f(0)$ to its fitness values; the vertical shift of $-f(0)$ is used so that grazing upon an empty field does nothing to benefit a moose's fitness. If two moose visit a field, then they will fight. Fighting is exhausting, prevents either moose from eating, and damages the local area. Neither moose adds any value to their fitness and x is decremented by one for that field, down to a minimum of zero. Since a field replenishes at the start of a time step, a moose never has to graze upon an empty field.

Fig. 1 demonstrates how time steps of the moose foraging game play out based on the decisions made by each moose; Each row in the figure shows the state of each field at the end of a single time step. The top row of fields shows the fifth time step of the game; the two moose have been continually

Daniel Ashlock and Connor Gregor are with the Department of Mathematics and Statistics at the University of Guelph, in Ontario, Canada, {dashlock|cgregor}@uoguelph.ca

Joseph Alexander Brown and Munir Makhmutov are with the *Artificial Intelligence in Games Development Lab*, Innopolis University, Innopolis, Russia, {j.brown|m.makhmutov}@innopolis.ru

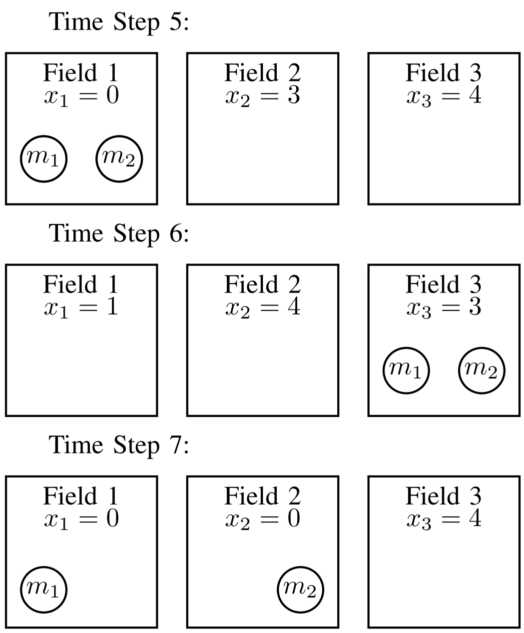

Fig. 1. A sequence of time steps which show two moose being in competition upon the same field of resources and then splitting apart so that they can both successfully forage.

fighting over the first field and, as a result, the other two fields have become rich with resources. At the sixth time step, both moose attempt to migrate to field three which is most rife with resources; by doing this, the moose enter conflict at field three. Neither moose gains any score to add to their fitness and x_3 is decremented. The other two fields that were left undisturbed then increment their x values. At the seventh time step, the moose cooperate as moose one and two migrate to fields one and two respectively. Because they each have their foraging field, the fitness of moose one and two are increased by $f(2) - f(0)$ and $f(5) - f(0)$ respectively. The lone field three increments x_3 and once again becomes the most desirable field for the moose to travel to during the next time step.

The original work [1] contains the results of the contest between student agents conducted for the Game Theory course at Innopolis University. The students of 2020 and 2021 have participated in the contest, the average results of the latter outperformed the first ones. It was claimed that the communication limitations caused by COVID-19 could be one of the reasons for this. Each pair of student agents have fought between each other for 100 rounds. The payouts of each moose for each game were summarized and compared to others. It was found out that for a given set of agents the most successful strategy assumes the application of farming and resting fields to maximize the payoffs. The farming field is required for eating vegetation; the resting field is required for waiting when the farming field will become sufficiently abundant in fodder. The students have demonstrated that the most optimal field value for eating is 5 or 6. In case of strife between moose, the winning agents were applying different forms of greedy algorithms.

For the experiments within our study, we have not put any limitations on the memory of the moose. The only limitation was that a moose can only see the value of the current field at the current move and know which field his opponent attempted to forage at last time. The moose agents used in this study play an iterated game, conditioning their choice of which field to attempt to forage in on the field chosen by their opponent previously chose and their own internal state.

This study investigates the result of making variations to the base scenario of the moose game. These variations include:

- altering the richness capacity C to C_k where $k \in \{1, 2, 3\}$ so that it is different for each of the three fields.
- changing the growth rate so that x increments by more than one when a field is left untouched.

With these variations, it has been investigated whether the moose cooperate and conflict differently under these variations or if their behavior is indifferent to the settings of these parameters.

The goal of the study is to finish characterizing the moose game (MG) sufficiently so that it can be employed as a family of games for use in evaluating the adaptability of different sorts of AI. It is also anticipated that the MG can be used as a suitable problem for the teaching of AI techniques.

The remainder of this study is structured as follows. Section II gives background information. Section III specifies the agent representation used to simulate the moose. Section IV gives the experimental design while Section V gives and discusses results. Finally, Section VI draws conclusions and discusses next steps.

II. Background

The MG is a turn-based strategy game, albeit a very simple one, that involves resource management, in the form of the logistic growth model for the plants. The first study on the MG [1] demonstrated that there are a variety of strategies that arise under the influence of evolution. Since the MG is competitive, the fitness function used – total forage acquired – does not have optima *per se*. Rather, for a given strategy there are effective and ineffective opponents. The fact that the MG has embedded conflict arising from the fact that the field not chosen by either player in the last time step is the most valuable property means that there is likely to be a rich strategy space, something confirmed in the earlier study.

While an obvious generalization of the MG is to change the number of fields or moose, in this study we will examine the effect on the evolution of agents of changing the feature of the simulation that has three identical fields as well as varying the growth rate of the plants. The MG can be thought of as simulating the emergence of instincts that can support survival in a species that is territorial. In Moose, territoriality arises from the animals need for large amounts of fodder to support its substantial size. An adult, male moose is shown in Fig. 2.

Fig. 2. An adult, male moose: credit *Wikimedia Commons*

III. Representation

We employ a finite state representation for the moose foraging behavior. An example of a 4-state machine is shown in Fig. 3. The agents are represented as a linear list of states with each state containing a response and transition destination for all three possible actions, forage in fields 1-3. The initial action is stored with the first state of the finite state machine. The variation operators are two point crossover of the list of states – with the initial action attached to the first state – together with mutation [21]. The point mutation generates a new, random, transition for one opponent action on a randomly selected state. The number of mutations is selected uniformly at random in the range 1-MNM. The *maximum number of mutations* (MNM) is set to 1 and the number of states is set to 8 for this study, based on a parameter study done in [1].

For games with a discrete set of moves, finite stare agents are a common choice, [2] and, in this case, permit direct comparison with the earlier study [1] where they were also used. Other natural choices, deferred to the future, are artificial neural nets [23], genetic programming systems [7], or lookup tables [6].

IV. Experimental Design

This study varies the growth rate of the plants in the fields visited by the moose and C: the maximum amount of forage that the model approaches, its richness capacity. Faster growing plants may make staying in the same field a more acceptable choice by decreasing the difference between available forage. The logistic growth model means that a field slows its growth the more forage there is in it, approaching the richness capacity, but never reaching it; $f(x)$ approaches C as $x \to \infty$, so $\frac{C}{2}$ is the largest amount of fitness that can be gained from single grazing. Varying the capacity of

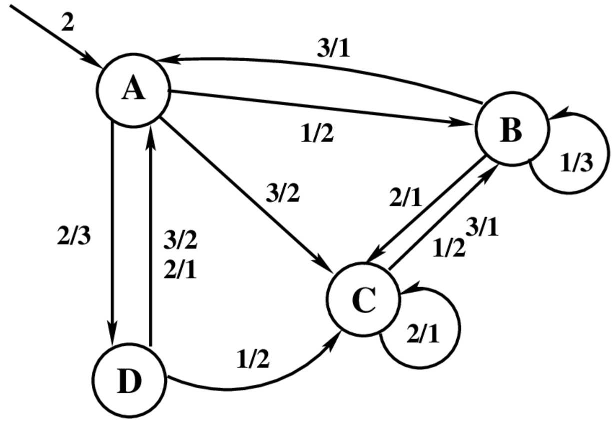

Fig. 3. An example of an evolved 4-state finite state machine. The agent's initial action appears on the sourceless arrow. Subsequent actions are generated by transitions of the form A/R where A denotes the opponent's most recent action and R the agent's response.

the fields simply means that there are better fields and worse fields as the minimal value obtained by a better field will be better than the minimal value of a worse field based on $f(x) - f(0)$. Making some of the fields more desirable changes the strategic equation. Table I specifies the choices made for each parameter in each experiment of this study.

TABLE I
Plant growth rates and field capacities used in the experiments in this study

Experiment Number	Growth Rate	Field Capacities		
1	1.0	10	10	10
2	2.0	10	10	10
3	3.0	10	10	10
4	1.0	10	10	20
5	1.0	10	10	30
6	1.0	10	20	20
7	2.0	10	20	20

A. The Evolutionary Algorithm

The experiments with the finite state moose agents used a population of 36 agents. Fitness evaluation consisted of a round-robin tournament with each pair of moose playing fifty rounds of the MG. The fitness of a single agent is equal to the average score they achieved against all opponents and games so that the units of fitness are the same as they would be for one play of the game.

Using the representation and variation operators described in Section III, agents were evolved for 250 generations with 30 replicates of the evolutionary algorithm in each experiment. The evolutionary algorithm updates the population of 36 agents by preserving the 24 most fit, the *elite*. The remaining 12 agents are then replaced via fitness proportional selection of pairs of parents from the elite. The parents are copied, and the copies undergo crossover and mutation to generate children. This evolutionary algorithm has been used in numerous past experiments with game-playing agents, e.g. [2], [3].

B. Reporting Statistics

The reporting statistics used are the *total fitness* of the evolving population and the fraction of plays of each possible type $x : y$ with the x and y being the fields chosen by the two moose playing. The total fitness is the area under the curve traced out by average fitness over the course of evolution, shown in Fig. 4. Total fitness provides a single-statistic estimation of the efficiency of the moose agents at extracting bio-energy from their environment, though it will also reflect the result of increasing plant growth rate and field capacity.

V. Results and Discussion

Fig. 5 shows the distribution of total fitness for each of the experiments. The first three experiments increase the growth rate of the plants. As this parameter is increased, the average fitness (acquired forage) for the moose increases. Experiments four and five increase the capacity for a single rich field; this leaves the moose in Experiment four in slightly worse shape than the somewhat faster-growing plants in Experiment two, yet still well above base example in Experiment one. Increasing the maximum vegetation of the rich field from 20 to 30 between Experiments 4 and 5 has a substantial effect, yet it also proved to be the wild and varying in total fitness results. Experiment six has two richer fields, as does Experiment seven which also ties in a higher growth rate for the plants. These experiments show an increase in the rate at which the moose acquire forage, a substantial and significant one for Experiment 7.

The distributions of fitnesses shown in Fig. 5 demonstrate that the moose are responding to the availability of forage

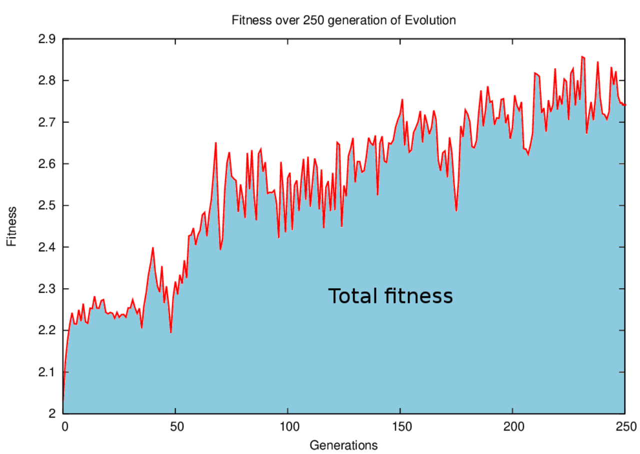

Fig. 4. Shown is a graph of the population average fitness of evolving moose agents in one run of the evolutionary algorithm. The shaded area under the curve is the *total fitness* for the experiment.

in a plausible manner. The serves as a check that the agent training algorithm is functioning nominally. With this check on the system accomplished it is time to turn to the conflict statistics.

Fig. 6 shows the fraction of encounters between moose that occurred in the same field. Experiments 1-3 test the hypothesis that faster-growing plants will decrease conflict. Between Experiment 1 and Experiment 2 the average level of conflict increased, but became much less variable. The level of conflict in Experiment 3, with the fastest-growing plants, was significantly lower. At the higher end with Experiment 7, faster-growing plants do decrease conflict. With Experiment 4-6 where the fields were given different capacity values, there was no clear trend regarding whether incidents of conflict increased or decreased. It ought to be noted that Experiment 5 had the highest number of conflict incidents; this is presumably due to the one field being so much more valuable and the moose repeating competing over harvesting from it.

Fig. 7 shows the fraction of all moves made that are in fields one, two, and three for each of the experiments. The first three experiments have all the fields of equal richness and so, nominally, should have roughly one-third of all plays in each field. This is, in fact what happens, certifying that the system is behaving well. This lets us turn to the four experiments with fields that have enhanced richness.

Experiments four and five increase field three to have an upper limit on plant growth of 20 and 30, up from 10. The statistics from these two runs are very close. This means that fraction of moves on the enriched field changes by about the same amount in spite of the greater richness of the field in Experiment 5. Looking back at Fig. 6, however, we see that the variability of conflict between the runs increases substantially between Experiments four and five. This means that the fraction of coordinated moves in which both moose try for field three has increased substantially. This may also explain why the increase in total fitness in Fig. 5 between the fourth and fifth experiments did not increase in proportion to the increase in the richness of the enhanced field.

Experiments six and seven increase the richness of fields two and three to 20. Experiment six left plant growth rates at the minimum level, while Experiment seven doubled the growth rate of the plants. First of all, the use of the field without enhanced richness was low in Experiment six and very low in experiment seven. Recall that the source of the conflict among the moose is the third, unexploited field. Considering field one as the least desirable and hence “third” field, the sharp decrease in conflict (Fig. 6) in Experiment means that the strategy of each moose sticking to a field is made practical by the lowered desirability of the third field.

Experiments one through three tested the hypothesis that increasing the growth rate of plants will decrease conflict and found that the hypothesis has support after a threshold is passed. Experiments six and seven provide additional support for this hypothesis – conflict dropped sharply between growth rates of 1.0 and 2.0 – and the threshold appears to be in

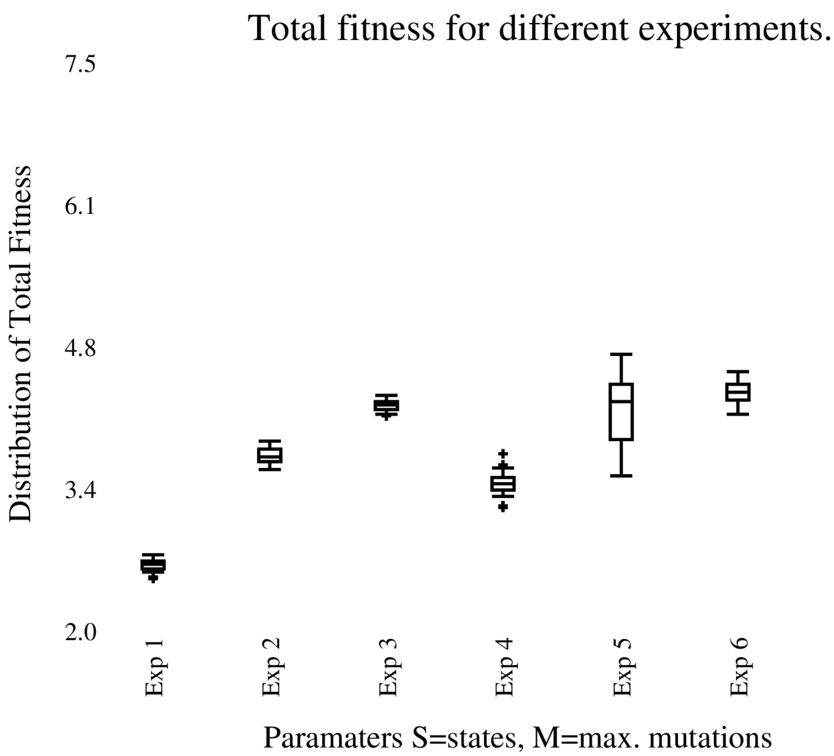

Fig. 5. Shown are the distribution of total fitness values for each of the experiments.

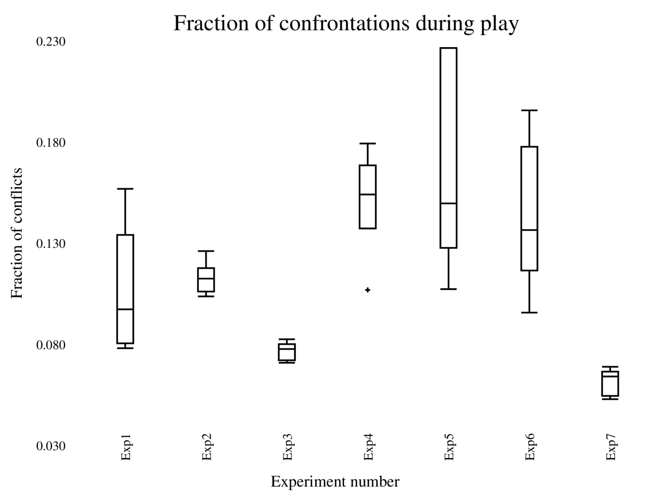

Fig. 6. Shown are the distribution, across experiments, of the fraction of plays that involved conflict. The nominal rate is 1/3, so the moose always managed to avoid one another somewhat.

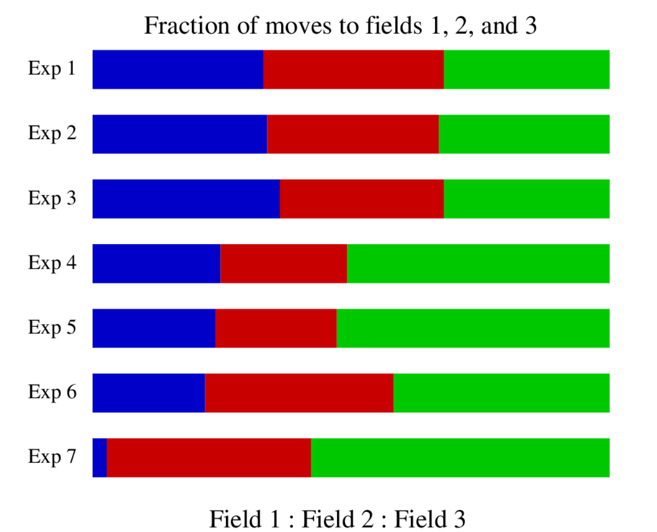

Fig. 7. Shown are the distribution, across experiments, of the fraction of plays of each type. The blue bar is the fraction of moves made to field 1, red bar for field 2, and green bar for field 3.

a different place when there are two fields with enhanced richness. This is a matter that could be profitable to studies done at greater resolution.

A. Pattern of Conflict During Evolution

The experiment with the highest rate of conflict was Experiment five, performed with a single field with richness enhanced to 30. We examined the trajectory of both fitness and percent of encounters involving conflict over the course of evolution. Three examples of such traces appear in Fig. 8. The high conflict runs were those like that in the middle of the figure – a pattern of sawing back and forth – which turns out to have a simple cause.

Negative density-dependent selection [15] is a situation in which the rarer of two species is more fit, as a result of its rarity. This causes the rarer species to have a higher

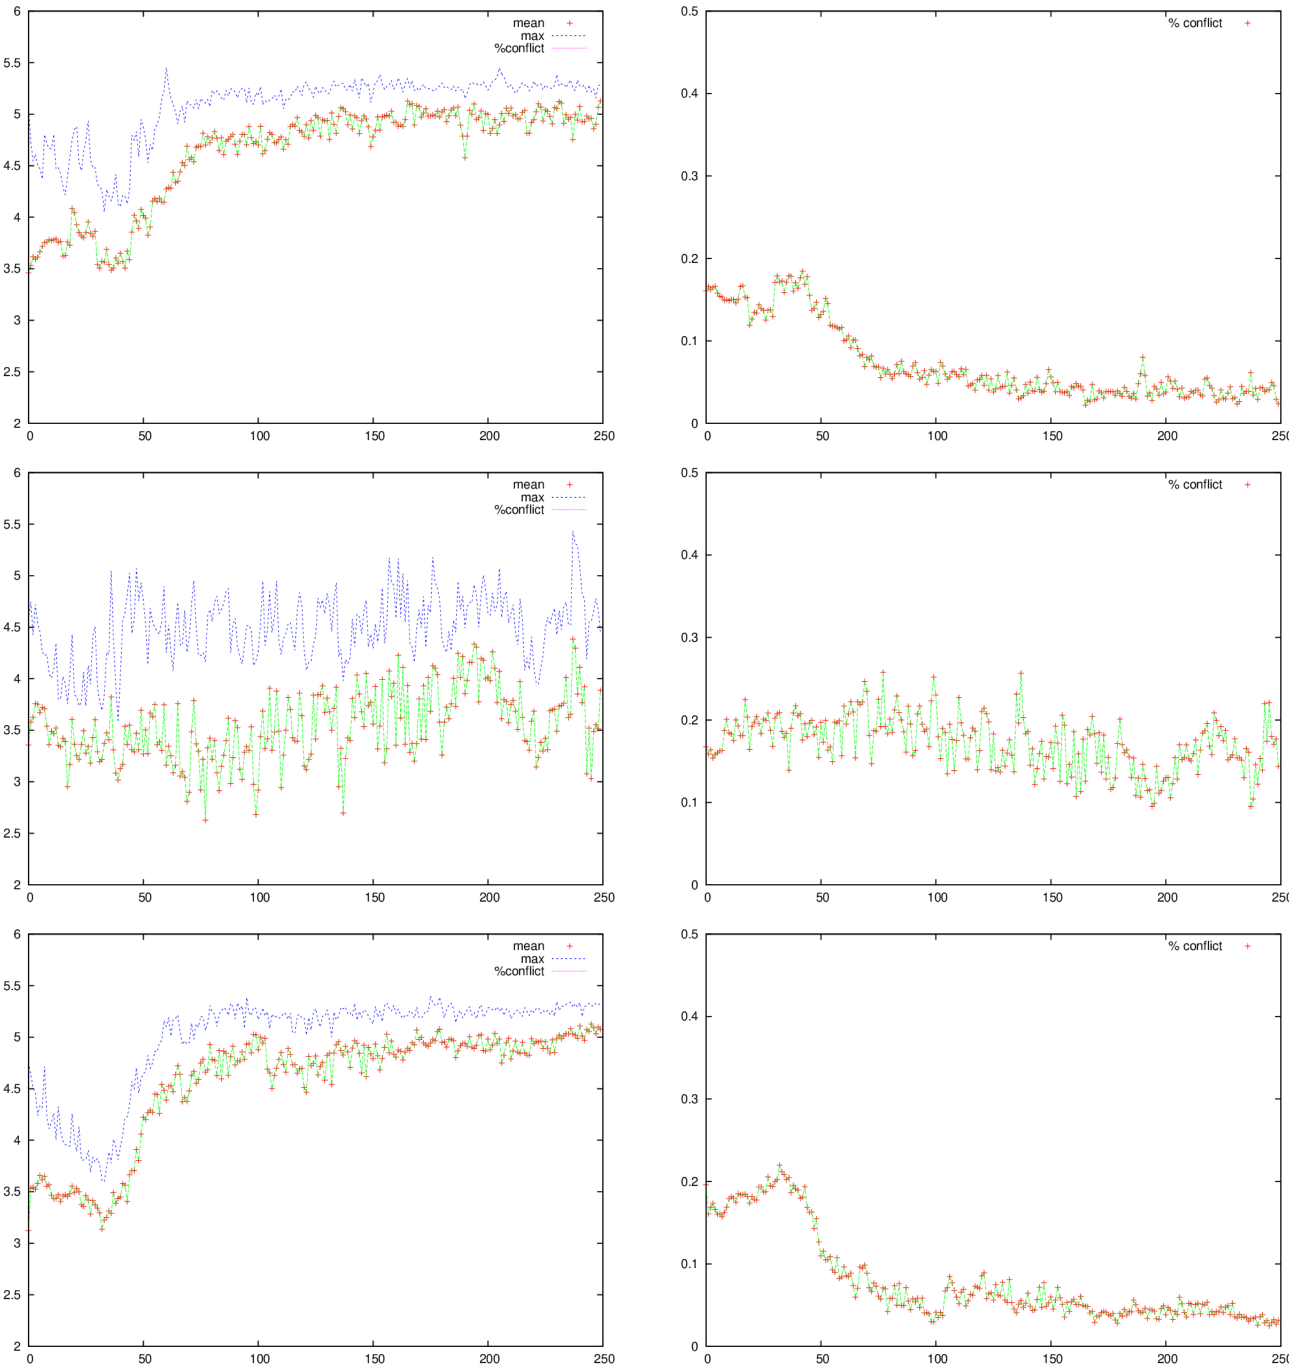

Fig. 8. Shown are the evolution of average and maximum fitness (left) and percent of encounters where moose were in conflict (right) for three of the evolutionary runs (21, 12, and 27) in Experiment 5, which had the highest levels of conflict.

reproduction rate, setting up the kind of irregular oscillations seen in the second population shown in Fig. 8. At least six of the thirty populations showed this behavior.

In this case, the mechanism of fitness as the result of rarity are two types of agents that access the enriched field at different times during the iterated play. Possibly splitting odd and even rounds or some more complex pattern. That means that, of these two types, the rarer one is more often granted untrammelled access to the enriched field as it is paired with agents of its own type less often.

VI. Conclusions and Next Steps

This study shows that for the most part, that the MG game responds in an intuitive way to the modification of its resource parameters. The hypothesis that increasing the growth rate of plants would reduce conflict turned out to be correct for portions of the plant growth rate parameter range, but was more complex. The increase above the nominal rate from Experiment 1 first increased conflict and substantially increased the variability of conflict. The additional increase did reduce conflict, this being seen in the first three experiments. Experiments six and seven also increased growth rates, but with two enriched fields. Here, the first increase in the plant growth rate yielded a large decrease in conflict.

Enhancing the richness of fields caused those fields to be more valuable. When both field richness capacity and plant growth rate were enhanced, this caused a *de facto* change in the strategic equation with the one field without enhanced richness becoming almost worthless in Experiment 7 (Fig. 7). These two parameters – field richness given by C and the growth rate represented by the per-time-step increment of the parameter x create a rich collection of games with distinct strategic landscapes.

A. Classroom Use of the MG

The MG emerged as an exercise for an AI in games class. The basic game, using the parameters in Experiment one, was used in contests for two years. This was part of a greater experiment in the encouraging of active learning techniques [18], [20], [24] due to their ability to increase student engagement and outcomes [12], see also GAs applied to computer art [8]. These techniques are difficult to employ in a mathematics or science framework [9]. The broadening of the game in this study suggests that the game can provide a good unit in an AI class and may be usable as a challenge for generalization in AI.

In a classroom setting, after explaining the game to the students, the instructor can ask the students to devise and defend strategies before having them compete in a tournament. The generalizations of the game given in this study can be used as thought experiments where the students predict the outcome of the modifications; the impact of increasing growth rate is not linear. The students can then be asked to test their hypothesis experimentally.

As seen in [1], the role of discussion of MG cannot be overestimated. The students of the second year had more opportunities for offline communication in comparison with the first-year students obliged to study remotely. For classes with many participants division into mini-groups for discussion of agents may be useful, because students can begin cooperating against other groups supporting their group members by choosing some secret "handshaking". This idea of division into teaching groups for cooperation was used by some students for getting a higher ranking in the contest.

B. As an AI Challenge

General game playing [13] seeks to find AI that can learn games while playing them, allowing a single AI to be a general game player. This is a very ambitious goal. Restricting the scope of the games to video games [19] or mathematical games [5] can yield a more accessible goal. The MG forms a family of target games with a small number of moves (the number of fields) but also with distinct strategies in different versions of the game. This is a relatively easy test problem in which to test AI generalization.

The generalization can be carried well beyond those in this study by increasing the numbers of moose and fields. A game with three fields and five moose would, for example, force conflict. In this study, different fields could have different richness, but all plants in the simulation had a single growth rate. The richness of the soil, access to water, and other natural factors make different growth rates in different fields plausible and this is not a feature that would be difficult to add to the simulation.

Finally, the moose could be placed in a natural area with an entire network of fields. A moose could see the occupancy and richness of adjacent fields, creating a strategic game with a complex resource model that might include seasonal variation in the plant growth. The game could be given additional variation by introducing natural predators to the moose that would roam between the fields. These predators would penalize the fitness of the moose if they stop in the same field as them.

C. The Empire Skin

The MG could be re-skinned into a more dramatic setting by replacing moose with army units and fields with small states or free cities on the border between two empires. Border realms left alone increase their wealth. Armies both trying to pillage the same border realm get no plunder and substantially reduce the wealth by fighting a battle. This skin may be of more interest to games researchers as it is closer to extant games.

A historical precedent occurs in the fights between the Guelphs and Ghibellines, especially *the War of the Oaken Bucket*, between the Italian towns of Bologna and Modena. This conflict, which, with the exception the Battle of Zappolino, was primarily small skirmishes and raids on the borders between the city states [14]. The attacks on actual cities would require large sieges in this time period as the cities had massive defensive walls, hence attacks upon the outlaying farms and villages would weaken the defending forces without much of your own expenditures in terms of resources while also providing resources and plunder to the attacking forces. We see also similar themes in the comments of Sun Tzu [22] on utilizing the provisions of the enemy[1], and on the avoiding of sieges[2], hence another theme for the MG mechanics could be the warlord periods of the Eastern Zhou period of China.

This themeing to be made into a producible board game would require a few changes to the mechanics, such a simplification of the field/farm production rules over time to whole values in order to allow for a token representation, and a method of ensuring a hidden deployment of forces; hidden moves are a common mechanic in a number of existing games and this should not pose a hard implementation issue.

D. Representation of Moose Agents

In [4] it was found that the collection of strategies that arise under evolution in the iterated prisoner's dilemma is substantially controlled by the representation, finite state machine, neural net, etc., chosen by the researcher. In [17] it was found that computational or informational resources, like the number of neurons or depth of reporting of past plays, could also substantially influence which strategies arose.

Prisoner's dilemma is a very simple game. A hypothesis that the authors do not think is true is that in a more complex game, representation and resources would have less of an impact on the outcome. The family of moose games is a more complex environment in which this hypothesis could be tested.

References

[1] D. Ashlock, J. A. Brown, S. Houghten, and M. Makhmutov. One moose, two moose, three fields, more. In *Proceedings of the IEEE 2021 Conference on Games*, pages 1–7, 2021.

[2] D. Ashlock and E.-Y. Kim. Fingerprinting: Automatic analysis and visualization of prisoner's dilemma strategies. *IEEE Transaction on Evolutionary Computation*, 12:647–659, 2008.

[3] D. Ashlock, E.Y. Kim, and W. Ashlock. Fingerprint analysis of the noisy prisoner's dilemma using a finite state representation. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(2):157–167, 2009.

[4] D. Ashlock, E.Y. Kim, and N. Leahy. Understanding representational sensitivity in the iterated prisoner's dilemma with fingerprints. *IEEE Transactions on Systems, Man, and Cybernetics–Part C: Applications and Reviews*, 36(4):464–475, 2006.

[5] D. Ashlock, E.Y. Kim, and D. Perez-Lebana. Toward general mathematical game playing agents. In *Proceedings of the 2018 IEEE Conference on Computational Intelligence in Games*, pages 110–116, 2018.

[6] Daniel Ashlock. *Evolutionary Computation for Opimization and Modeling*. Springer, New York, 2006.

[7] W. Banzhaf, P. Nordin, R. E. Keller, and F. D. Francone. *Genetic Programming : An Introduction*. Morgan Kaufmann, San Francisco, 1998.

[8] Joseph Alexander Brown, Hamna Aslam, Daria Miklashevskaya, and Nikita Lozhnikov. Pedagogical evolved art: An examination and results of the innopolis al art contest. In *2020 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8, 2020.

[9] Jesse Eickholt. Barriers to active learning for computer science faculty. *ArXiv*, abs/1808.02426, 2018.

[10] Rose Eveleth. Why do moose fight so much, December 2013. [Online; posted 6-December-2013].

[11] D.B. Fogel. Evolving behaviors in the iterated prisoners dilemma. *Evolutionary Computation*, 1(1), 1993.

[12] Scott Freeman, Sarah L. Eddy, Miles McDonough, Michelle Knovic Smith, Nnadozie Okoroafor, Hannah Jordt, and Mary Pat Wenderoth. Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences of the United States of America*, 111 23:8410–5, 2014.

[13] Michael Genesereth, Nathaniel Love, and Barney Pell. General game playing: Overview of the aaai competition. *AI Magazine*, pages 62–72, 2005.

[14] Matteo Griffoni. *Memoriale historicum de rebus bononiensium*, chapter Conflictus Zapolini. Città di Castello, Tipi dell'editore S. Lapii (1902 ed.), 1325.

[15] M. Hixon. Density dependence and independence. In *Encyclopedia of Life Sciences*. John Wiley & Sons, Chichester,UK, 2009.

[16] Tomonori Kawano, Nigel Wallbridge, and Carrol Plummer. Logistic models for simulating the growth of plants by defining the maximum plant size as the limit of information flow. *Plant Signaling & Behavior*, 15(2):1709718, 2020. PMID: 31984861.

[17] E. Y. Kim and D. Ashlock. Changing resources available to game playing agents: Another relevant design factor in agent experiments. *IEEE Transactions on Computational Intelligence and AI in Games*, 9(4):312–322, 2016.

[18] Manuel Mazzara, Petr Zhdanov, Mohammad Reza Bahrami, Hamna Aslam, Iouri Kotorov, Muwaffaq Imam, Hamza Salem, Joseph Alexander Brown, and Ruslan Pletnev. Education after covid-19. In *Smart and Sustainable Technology for Resilient Cities and Communities*. Springer, 2021.

[19] Diego Perez Liebana, Spyridon Samothrakis, Julian Togelius, Tom Schaul, Simon M Lucas, Adrien Couëtoux, Jerry Lee, Chong-U Lim, and Tommy Thompson. The 2014 General Video Game Playing Competition. *IEEE Transactions on Computational Intelligence and AI in Games*, 8(3):229–243, 2016.

[20] A. Sadovykh, M. Beketova, and M. Khazeev. Hackathons as a part of software engineering education: Case in tools example. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12271 LNCS:232–245, 2020.

[21] William M Spears. Adapting crossover in evolutionary algorithms. In *Evolutionary programming*, pages 367–384, 1995.

[22] Sun Tzu and Lionel Giles (translation). *The Art Of War: The Oldest Military Treatise In The World*. Allandale Online Publishing, 1910 trans.

[23] X. Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447, Sep 1999.

[24] Oksana Zhirosh, Joseph Alexander Brown, and David Tickner. Democratizing faculty development - establishing a training program at a new computer science university in russia. In *ASEE Annual Conference and Exposition, Conference Proceedings*, page Paper ID #25473, Tampa, Florida, 2019.

[1] "II.9 Bring war material with you from home, but forage on the enemy. Thus the army will have food enough for its needs." and "II.15 - Hence a wise general makes a point of foraging on the enemy. One cartload of the enemy's provisions is equivalent to twenty of one's own, and likewise a single picul of his provender is equivalent to twenty from one's own store." Other passages in this section comment on the effects of morale, paying troopers who take from the enemy their chariots, and how captured good should have their symbols (flags) changed when looted.

[2] "II.2 - When you engage in actual fighting, if victory is long in coming, then men's weapons will grow dull and their ardor will be damped. If you lay siege to a town, you will exhaust your strength."