Multi-Objective Evolutionary Game Theory: A case study in cancer therapy

Lukas Bostelmann-Arp[1], Sanaz Mostaghim[1], Andreas Braun[2] and Thomas Tüting[2]

[1] Faculty of Computer Science, Otto-von-Guericke University Magdeburg
[2] Department of Dermatology, Otto-von-Guericke University Magdeburg
Lukas.Bostelmann-Arp@ovgu.de

Abstract

In this paper, we introduce an early concept of using multi-objective optimization to study various emerging strategies in evolutionary game theory and show its application in a case study. We aim to analyze the emergent behavior when changing the game's environment through optimization. The multi-objective approach allows looking at the results of each model evaluation from different points of view. For the realization, we suggest the use of a multi-agent model to compute the outcome of a game. Such a model allows modeling even complex interrelationships and can be used as input to multi-objective optimization algorithms. Finally, we demonstrate a use case by optimizing therapy plans for melanoma through the incorporation of medications into a multi-agent model of concurring cell populations in the tumor micro environment.

Introduction

In evolutionary game theory (EGT), regular game theory (GT) is used in an evolutionary context that represents a sequence of interactions among the participants. The focus lies on populations instead of individual players. Fitness plays a fundamental role in population survival through reproduction (Sandholm, 2020). The individuals who form the populations do not actively reason about their decisions, but inherit strategies through genetic operators. Therefore, the strategies or behaviors that emerge are of great interest. Analogous to the Nash equilibrium in GT, evolutionarily stable strategies (ESS) are especially important. Once adopted by a population in a specific environment, a set of ESS cannot be substituted by a novel set of strategies.

However, the emerging set of strategies depend on the game's rules. Therefore, altering those rules might allow exploring additional emerging sets of strategies. This raises the question of how the different populations adapt to the changed environment. Note that the environment is not changed during the interactions, but only for different runs of the model. Further, the properties of the rule changes and their effects on the emerging set of strategies can be evaluated based on various, sometimes conflicting, criteria.

Therefore, the goal of this paper is to propose the incorporation of a multi-objective optimization scheme into the EGT to allow analyzing several sets of strategies which are the results of an optimization problem.

For this purpose, we study agent-based models (ABMs) as a mode to simulate the evolutionary game. In such models, autonomous agents with basic rules interact with each other. The result of these interactions is usually an emerging behavior. An overview regarding agent-based modelling and its tools is provided by Abar et al. (2017). Agent-based models, or multi-agent systems in general, require the use of a robust optimization algorithm, especially in case of stochastically determined decisions of agents. There are several works in the literature that use Evolutionary Algorithms for this purpose, e.g., Moya et al. (2021) studied the use of different evolutionary multi-objective algorithms (EMOAs) for the automatic calibration of the model itself.

Moreover, there is another reason why agent-based models can be the connecting piece between EGT and evolutionary optimization algorithms: Adami et al. (2016) compared standard EGT methods to the results of ABMs and concluded that the latter can be beneficial in the prediction of more realistic scenarios, where current mathematical tools reach their limits. Nonetheless, they and others (Hilbe and Traulsen, 2016) make clear, that both have their purpose and that mathematical methods in particular can help keeping agent-based systems from becoming too arbitrary. Regarding the computational side, a framework has been proposed (Izquierdo et al., 2019) that allows the simulation of evolutionary game dynamics through agent-based systems. Based on this, there are already a few works applied to real scenarios, e.g., Coelho and Ralha (2022) studied the land use, respectively coverage, by simulating interacting human entities.

Multi-Objective EGT (MO-EGT)

In MO-EGT, we introduce two counterparts, namely ABMs based on EGT, and the multi-objective optimization algorithm (MOA). The ABM contains several populations which interact according to the rules defined by a payoff matrix defined in an EGT framework. Our goal, in using MOA, is to obtain several optimal values for the pay-off matrix by optimizing two conflicting functions. The first function is meant to regulate the population size and the second function describes the cost in the change of environment or payoff matrix, respectively. These can be best described using our case study.

Here we consider the interaction of multiple cell populations in the tumor micro environment. This is modeled using an ABM in an EGT framework. Within this environment, many physical restrictions, but also inter cell conflicts, exert pressure on cancer cells. These influences can be seen as selective forces, and therefore justify the interpretation of the tumor progression as an evolutionary game (Wölfl et al., 2021). We define the interactions between the normal cells and tumor cells using a pay-off matrix. In the optimization process, the focus lies on therapeutic intervention, which influence the values in the pay-off matrix. For example, one value of the payoff matrix indicates the damage caused by an immune cell to a tumor cell, considering the current drug concentration at the site of interaction. A current research question deals with finding adequate sequences and intervals for various medications for therapeutic interventions. This is necessary to prevent the formation of resistances due to the cellular plasticity given by the evolutionary characteristics of cancer cells. The goal is then to keep a stable population of tumor cells, or even a declining one. Therefore, the number of surviving tumor cells is counted for the first objective function $f1$. The second objective function $f2$ for MOA concerns the cost of the therapy.

In our experiments, we take the NSGA-II (Deb et al., 2002) as the multi-objective optimization algorithm. Further, a population size of 52 with a termination criterion of 50 generations is used. Note, that the term population is used here in the context of the EMOA and is different from the population in an ABM. Regarding the variation operators, a simple one-point crossover is used together with a Gaussian mutation. Both operators are extended to handle a variable number of drug administrations, as this, together with the dosage, is part of the optimization problem.

The underlying ABM consists of five different cell populations. The most important one, the tumor cells, feature the ability to change their state. This behavior, called differentiation in a biological sense, represents the evolutionary aspect of this game.

Results

Regarding the experiment, the optimization algorithm was run five times to mitigate stochastic effects contained in the algorithm itself. Since a simulation also contains probabilistic decisions, each evaluation was repeated 16 times before computing the fitness values from the average.

The combined Pareto front is shown in Figure 1 for the two objectives $f1$ and $f2$ introduced earlier. Since the tumor started with 100 tumor cells, the results that end up with around 100 tumor cells can be considered stable regarding the size of the tumor. This is marked by the dashed vertical line in the plot. Nonetheless, the algorithm was able to produce a well distributed front with different compromises between the two objectives. Each of those solutions can be analyzed regarding the composition of the tumor and the cell's strategy or behavior, respectively. This information can help the physician better understand the tumor's response to the corresponding therapy. The next step, deals with the decision-making process. As an evolutionary multi-objective algorithm already presents a wide range of possible solutions, it is easier for the decision-maker to choose a final one, than designing one from scratch. This results in a higher confidence towards the decision. However, selecting among the solutions requires additional information about the patient which exceeds the tumor composition data required for the simulation initialization. That includes, for example, information regarding the lifestyle or psychological condition.

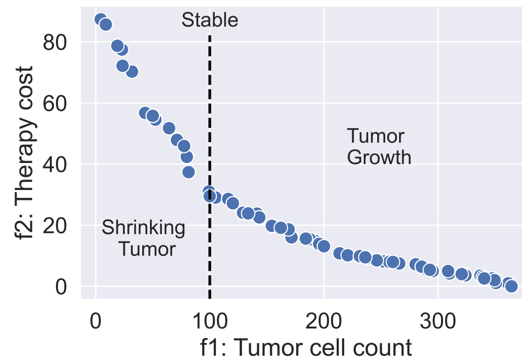

Figure 1: Combined Pareto front

Conclusion and Future Work

In this paper, we proposed a workflow and illustrated, how multi-objective optimization can be implemented into EGT. So far, EGT is being used to compute a certain optimal strategy, in most cases, the so called evolutionarily stable strategy. Here, we study several sets of strategies which are the results of a multi-objective optimization problem that altered the game's rules.

While the presented case study showed that our proposed method works, there are still some gaps to fill. This concerns above all the explainability. As the original EGT is based on mathematical tools, its correctness can be proofed. While it may be possible for simple agent-based models that apply the same rules used in EGT, it is hard for more complex ones and especially for the optimization algorithm itself. In addition, the possible outcomes heavily depend on the implementation of the variation operators as well as the individual of the EMOA. Nonetheless, the multi-objective approach allows analyzing and comparing different results and emerged strategies of an evolutionary game, which were created by optimizing rule changes.

References

Abar, S., Theodoropoulos, G. K., Lemarinier, P., and O'Hare, G. M. (2017). Agent based modelling and simulation tools: A review of the state-of-art software. *Computer Science Review*, 24:13–33.

Adami, C., Schossau, J., and Hintze, A. (2016). Evolutionary game theory using agent-based methods. *Physics of Life Reviews*, 19:1–26.

Coelho, C. G. C. and Ralha, C. G. (2022). Mase-egti: An agent-based simulator for environmental land change. *Environmental Modelling & Software*, 147:105252.

Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6:182–197.

Hilbe, C. and Traulsen, A. (2016). Only the combination of mathematics and agent-based simulations can leverage the full potential of evolutionary modeling. *Physics of Life Reviews*, 19:29–31.

Izquierdo, L. R., Izquierdo, S. S., and Sandholm, W. H. (2019). An introduction to abed: Agent-based simulation of evolutionary game dynamics. *Games and Economic Behavior*, 118:434–462.

Moya, I., Chica, M., and Cordon, O. (2021). Evolutionary multi-objective optimization for automatic agent-based model calibration: A comparative study. *IEEE Access*, 9:55284–55299. Interesting.

Sandholm, W. H. (2020). *Evolutionary Game Theory*, pages 573–608. Springer US.

Wölfl, B., te Rietmole, H., Salvioli, M., Kaznatcheev, A., Thuijsman, F., Brown, J. S., Burgering, B., and Staňková, K. (2021). The contribution of evolutionary game theory to understanding and treating cancer. *Dynamic Games and Applications*.