

# Learning to Play like a Human: Case Injected Genetic Algorithms for Strategic Computer Gaming

Chris Miles    Sushil J. Louis    Nicholas Cole  
Evolutionary Computing Systems LAB  
University of Nevada  
Reno - 89557  
{miles,sushil,ncole}@cs.unr.edu

John McDonnell  
SPAWAR Systems Center  
San Diego, CA  
john.mcdonnell@navy.mil

March 11, 2004

## 1 Introduction

We use case injected genetic algorithms to learn how to competently play computer strategy games that involve long range planning across complex dynamics [1]. Imperfect knowledge presented to players requires them adapt their strategies in order to anticipate opponent moves. We focus on the problem of acquiring knowledge learned from human players, in particular we learn general routing information from a human player in the context of a strike force planning game. By incorporating case injection into a genetic algorithm, we show methods for incorporating general knowledge elicited from human players into future plans. In effect allowing the GA to take important strategic elements from human play and merging those elements into its own strategic thinking. Results show that with an appropriate representation, case injection is effective at biasing the genetic algorithm toward producing plans that contain important strategic elements used by human players.

Our research focuses on a strike force asset allocation game which maps to a broad category of resource allocation problems in industry and the military. Genetic algorithms can be used in our game to robustly search for effective allocation strategies, but these strategies may may not necessarily approach real world optima as the game is an imperfect reflection of reality. Humans with past experience playing the real world game tend to include external knowledge (gained through their experience) when producing strategies for the real-world problem underlying the game. Acquiring and incorporating knowledge from the way these humans play should allow us to carry over some of this external knowledge into our own play. Results show that case injection combined with a flexible representation biases the genetic algorithm toward producing strategies similar to those learned from human players. Beyond playing similarly on the same mission, the genetic algorithm can use strategic knowledge across a range of similar missions to continue to play as the human would. The left side of Figure 1 shows a screen shot of the game we are designing while the right side shows the scenario considered in this abstract.

## 2 The Game

Strike force asset allocation consists primarily of allocating a collection of strike assets to a set of targets. We have implemented this game on top of a professional game engine making it more interesting than a pure optimization problem, therefore enticing more people to play the “game” and so that we can acquire player knowledge from them. The game involves two sides: Blue and Red, Blue allocates a set of platforms (aircraft) to attack Red’s targets (buildings) while Red has defensive installations (threats). These threats complicate complicate Blue’s planning, as does the varying effectiveness of Blue’s weapons against each




Figure 1: Left: Game Screen-shot. Right: Mission scenario

target. Potential new threats and targets can also "pop-up" on Red's command in the middle of a mission, requiring Blue to be able to respond to changing game dynamics. Both players seek to minimize the damage they receive while maximizing the damage dealt to their opponent. Red plays by organizing defenses in order to best protect its targets and Red's ability to play popups can affect its planning. For example, feigning vulnerability can lure Blue into a pop-up trap, or keep Blue from exploiting a weakness out of fear of such a trap. Blue plays by allocating platforms and their assets (weapons) as efficiently as possible in order to destroy the targets while minimizing risk. Many factors determine risk. The platform's route, the effect of accompanying wingmen, weather, and the presence of threats around chosen targets. GAP develops strategies for the attacking strike force, including flight plans and weapon targeting for all available aircraft. When confronted with popups, GAP responds by replanning with the case injected genetic algorithm to produce a new plan of action. Beyond producing near optimal strategies we would like to bias it toward producing solutions similar to those it has seen used by humans playing Blue in the past. We do this so that GAP can be a better and more versatile player who incorporates strategies taken from a range of human experts. Specifically we want to show that when using case injection the genetic algorithm more frequently produces plans similar to those a human has played in the past on the same mission. We show that GAP can continue to use information taken from the human, even when confronted with a different mission, to continue to play strategies closer to the style of that human than GAP's non injected counterpart.

Previous work in strike force asset allocation has been done in optimizing the allocation of assets to targets, the majority of it focusing on static pre-mission planning [2, 3, 4, 5]. A large body of work exists in which evolutionary methods have been applied to games [6, 7, 8, 9, 10]. However the majority of this work has been applied to board, card, and other well defined games which have many differences from popular real time strategy (RTS) computer games such as Starcraft, Total Annihilation, and Homeworld[11, 12, 13]. Entities in our game exist and interact over time in continuous three dimensional space. Sets of algorithms control these entities, seeking to fulfill goals set by the player leading them. This adds a level of abstraction not found in those traditional games. In most of these computer games, players not only have incomplete knowledge of the game state, but even identifying the domains of this incomplete knowledge is difficult. John Laird [14, 15, 16] surveys the state of research in using Artificial Intelligence (AI) techniques in interactive computers games. He describes the importance of such research and provides a taxonomy of games. Several military simulations share some of our game's properties [17, 18, 19], the purpose of our game is not to provide a perfect reflection of reality however, but to both provide a platform for research in strategic planning and to have fun.

The mission being played is shown in on the right side of Figure 1. This mission was chosen to be simple, to have easily analyzable results, and to allow the GA to learn external knowledge from the human. As many games show similar dynamics, this mission is a good arena for examining the general effectiveness of using case injection for learning from humans. The mission takes place in Northern Nevada and California, Lake Tahoe is visible near the bottom of the map. Blue possesses one platform which is armed with 8

assets (weapons) and the platform takes off from and returns to the lower left hand corner of the map. Red possesses eight targets distributed in the top right region of the map, and six threats that defend them. The first stage in Blue’s planning is determining the allocation of the eight assets. Each asset can be allocated to any of the eight targets, giving  $8^8 = 2^{24}$  allocations. The second stage in Blue’s planning involves finding routes for each of the platforms to follow during their mission. These routes should be short and simple but still minimize exposure to risk. We categorize Blue’s possible routes into two categories. Yellow routes fly through the corridor between the threats, while green routes fly around. The evaluator has no direct knowledge of potential danger presented to platforms inside the corridor area. Because of this, the evaluator optimal solution is the yellow route, as it is the shortest. The human expert however, understands the potential for danger in the corridor because of likelihood of a pop-up trap. Knowing this the green route is the human optimal solution. Our goal is now to bias the GA to produce the green route, while still optimizing the allocation. Teaching GAP to learn from the human and produce green strategies even though yellow strategies have higher fitness is the goal of this research.

### 3 Results

Parameterizing our router allows us to produce routes of different types, incorporating this parameter into the chromosome then allows individuals to contain information about not just what weapons to use and which targets to attack, but general information about how to reach those targets. Specifically, referring to Figure 1 (right), we want to learn the parameter for green routes overcome the algorithm’s preference for short routes in order to avoid traps. By injecting cases that incorporate a preference for green routes we can change the proportion of green routes produced by our system, despite such solutions having lower than evaluator-optimal fitness. Furthermore by artificially inflating the fitness of injected solutions and their descendants in the system’s population, we can further control the proportion of routes produced that avoid potential traps. Figure 2 shows these proportions in terms of our routing parameter,  $rc$ .

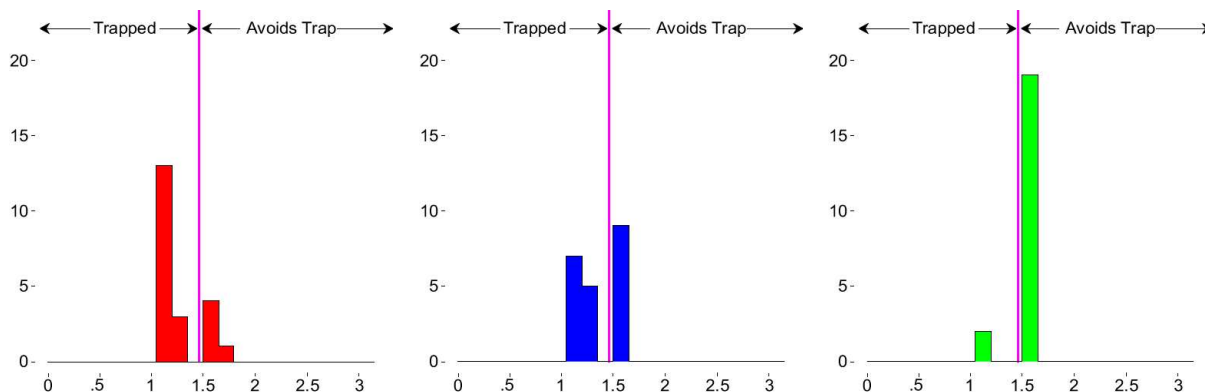


Figure 2: Distribution of RC values. Left: GA alone. Middle: With case injection. Right: Case Injection and fitness biasing.

We wish to investigate knowledge acquisition and incorporation (from humans) to bias genetic algorithm search to model human decision making. Knowledge acquisition is generally considered a difficult Artificial Intelligence (AI) problem and we tackle this using case-injected genetic algorithms. Case-injected genetic algorithms work by saving individuals from the population of a GA, and later introducing them into a GA solving a similar but different problem. Louis showed that case injection improves convergence speed and the quality of solutions found by biasing the current search toward promising regions identified from previous problem solving experience [1, 20]. In this work, the system *acquires cases from humans* for injection into GAP’s population. The idea is to automatically acquire cases by instrumenting the game interface to record all human decision making during game play. Our goal is not only to improve the GA’s performance, but

to bias the search by acquiring and using knowledge that is external to the actual evaluation and fitness of an individual plan – knowledge being expressed by expert humans in their formation of game plans. The genetic algorithm searches for an optimal strategy while case injection biases it to contain elements from strategies used by humans. Used in conjunction we hope to produce near optimal strategies that incorporate important information external to the evaluation function.

We would also like to incorporate dynamics into the game without the expense of simulating them. Consider a game involving armies doing battle, both players periodically rotate out their front line troops to let them rest. Whether or not the game simulation has an accurate model of fatigue the result is much the same, playing in anticipation of this dynamic (fatigue) is equivalent to implementing it. If GAP can incorporate these game dynamics from watching humans play, it should deepen the feeling of strategy involved in playing the game.

Our Genetic Algorithm Player (GAP) should be able to function both as a decision aid and as a trainer. GAP should be able to function as a trainer, a player who plays not just to win but to teach their opponent how to better play the game, in particular to prepare them for future play against human opponents. This would allow us to use GAP for acquiring knowledge from human experts and transferring that knowledge to less adept human players without the expense of individual training with experts. We want to use GAP for decision support, whereby GAP provides suggestions and alternative strategies to humans actively playing the game. Producing strategies that are both near optimal strategies and compatible with those being considered by human players, should help maximize the positive effect on the decision making process.

These roles require GAP to play with objectives in mind besides that of winning – objectives that would be difficult to quantify inside the evaluator. As humans function effectively in these regards, learning from them should help GAP better fulfill these responsibilities. Preliminary results indicate that case injection seems to be a promising approach towards acquiring knowledge and biasing genetic search toward human preferred solutions.

## Acknowledgment

This material is based upon work supported by the Office of Naval Research under contract number N00014-03-1-0104.

## References

- [1] Sushil J. Louis. Evolutionary learning from experience. *Journal of Engineering Optimization*, To Appear in 2004.
- [2] B. J. Griggs, G. S. Parnell, and L. J. Lemkuhl. An air mission planning algorithm using decision analysis and mixed integer programming. *Operations Research*, 45(5):662–676, Sep–Oct 1997.
- [3] V. C-W. Li, G. L. Curry, and E. A. Boyd. Strike force allocation with defender suppression. Technical report, Industrial Engineering Department, Texas A&M University, 1997.
- [4] K. A. Yost. A survey and description of usaf conventional munitions allocation models. Technical report, Office of Aerospace Studies, Kirtland AFB, Feb 1995.
- [5] Sushil J. Louis, John McDonnell, and N. Gizzi. Dynamic strike force asset allocation using genetic algorithms and case-based reasoning. In *Proceedings of the Sixth Conference on Systemics, Cybernetics, and Informatics*. Orlando, pages 855–861, 2002.
- [6] David B. Fogel. *Blondie24: Playing at the Edge of AI*. Morgan Kaufman, 2001.
- [7] Christopher D. Rosin and Richard K. Belew. Methods for competitive co-evolution: Finding opponents worth beating. In Larry Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 373–380, San Francisco, CA, 1995. Morgan Kaufmann.

- [8] Jordan B. Pollack, Alan D. Blair, and Mark Land. Coevolution of a backgammon player. In Christopher G. Langton and Katsunori Shimohara, editors, *Artificial Life V: Proc. of the Fifth Int. Workshop on the Synthesis and Simulation of Living Systems*, pages 92–98, Cambridge, MA, 1997. The MIT Press.
- [9] Graham Kendall and Mark Willdig. An investigation of an adaptive poker player. In *Australian Joint Conference on Artificial Intelligence*, pages 189–200, 2001.
- [10] A. L. Samuel. Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3:210–229, 1959.
- [11] Blizzard. Starcraft, 1998, [www.blizzard.com/starcraft](http://www.blizzard.com/starcraft).
- [12] Cavedog. Total annihilation, 1997, [www.cavedog.com/totala](http://www.cavedog.com/totala).
- [13] Relic Entertainment Inc. Homeworld, 1999, [homeworld.sierra.com/hw](http://homeworld.sierra.com/hw).
- [14] John E. Laird. Research in human-level ai using computer games. *Communications of the ACM*, 45(1):32–35, 2002.
- [15] John E. Laird and Michael van Lent. The role of ai in computer game genres, 2000.
- [16] John E. Laird and Michael van Lent. Human-level ai’s killer application: Interactive computer games, 2000.
- [17] Gil Tidhar, Clinton Heinze, and Mario C. Selvestrel. Flying together: Modelling air mission teams. *Applied Intelligence*, 8(3):195–218, 1998.
- [18] Graeme Murray Serena. The challenge of whole air mission modeling, 1995.
- [19] D. McIlroy and C. Heinze. Air combat tactics implementation in the smart whole air mission model. In *Proceedings of the First International SimTecT Conference, Melbourne, Australia, 1996.*, 1996.
- [20] Sushil J. Louis and John McDonnell. Learning with case injected genetic algorithms. *IEEE Transactions on Evolutionary Computation*, To Appear in 2004.