

Deep Barca: A Probabilistic Agent to Play the Game *Battle Line*

S. McCulloch

Ohio Wesleyan University

stmccull@owu.edu

Daniel Bladow

Gonzaga University

dbladow@zagmail.gonzaga.edu

Tom Dobrow

Google, Inc

tomdobrow@gmail.com

Haleigh Wright

University of California,

Santa Barbara

wright@umail.ucsb.edu

ABSTRACT

Recent years have seen an explosion of interest in “modern” board games. These differ from the “classic” games typically seen in Artificial Intelligence research (e.g. Chess, Checkers, Go) in that the modern games often have a large component of randomness or non-public information, making traditional game-tree methods infeasible. Often, these modern games have an underlying mathematical structure that can be exploited. In this paper, we describe an intelligent agent to play the game *Battle Line*, which uses elements of theorem-proving and probability to play intelligently without utilizing game trees. The agent is superior to the only other known computer player of the game and plays at a level competitive with top human players.

CCS CONCEPTS

- Computing methodologies ~ Probabilistic reasoning
- Computing methodologies ~ Game tree search

ADDITIONAL KEYWORDS

Modern Board Games, Battle Line

1 INTRODUCTION

Historically, the main work in developing intelligent agents to play board games has focused on “classic” games such as Chess [5], Checkers[12], or Othello[4]. The agents that play these games often focus on variants and improvements to game-tree search. While modern agents have many clever approaches and optimizations, there is still an exponentially-sized game tree to search. Even the recent “AlphaGo” program by Google [13] uses neural networks, but does so in order to evaluate a game tree more efficiently.

In recent years, an influx of new games have been created of varying levels of complexity. These games often depend on non-public information (such as a secret hand of cards or tiles) or randomness, making the size of a typical game tree infeasibly large.

There have been efforts to design agents to play these games. Some of these efforts have involved designing variants of classical search methods [8] [10] [11], some have applied multiple agents to the problem, and some have kept a simpler model of the game to reduce the complexity [14]. Another approach is to develop a rule-

based system derived from specific rules and strategies of a given game. These rules can then either be applied directly to choose a move, or serve as the basis of an evaluation function in a search [7] [8].

While these methods can produce competitive players, any approach that uses a variant of classical search will eventually become confounded by the exponential growth of the game tree—often sooner in modern games than in more classical games because of the need to model randomness or hidden information. Efforts to reduce the game’s complexity may remove important information. Rule-based systems can only perform as well as the rules in which they are given, and since those rules are typically programmed in by humans, they do not capture the autonomous nature that we desire in our agents.

Our approach is to notice that these games often have interesting underlying mathematical structure and to exploit that structure. If we can recast a game’s strategy as a mathematical problem, then we have the ability to use theorems and approaches from mathematics to design an agent that makes moves based on an underlying mathematical model. We feel that an agent made according to this paradigm is more autonomous than one that follows a rule-based system or an evaluation of leaves of a game tree because its strategy comes from the mathematical underpinnings of the game itself.

We have used this approach in the past on a different game, called *Football Strategy* [8], where the rules of the game can be viewed as a normal-form game, in the game-theoretic sense. By applying concepts of Nash Equilibria, our agent created a mixed strategy that made it competitive against top human players.

In this paper we focus on the two-player strategy card game *Battle Line*, published by GMT Games in 2000 [1]. We have developed an agent named “Deep Barca” that plays this game by exploiting the underlying probabilistic and logical nature of the card game.

2 DESCRIPTION OF THE GAME

In *Battle Line* there are two decks of cards: Troop cards and Tactic cards. Most of the game is played using the Troop cards, which is a deck of 60 cards consisting of the numbers 1-10 in six different colors. On their turn, players play a Troop card from their hand on one of nine flags. Each player will eventually place three cards on each of nine flags. The three card hand that is formed is called a

formation. (See Figure 1 for an example of what flags and formations look like). The player with the better formation will claim the flag. Winning the game requires claiming 5 of the 9 flags or any 3 flags in a row.

2.1 Evaluating Formations

The values of the formations correspond to three-card poker hands. The order in which cards are played does not matter; the cards in the final formation are evaluated collectively. The ranks of the different formations are as follows (the actual game uses military terminology for these ranks, but we have translated them into terminology that is closer to that used in poker to aid understanding).

- A straight flush is the highest possible hand. It consists of three cards, all the same color, in consecutive order (for example, the 7, 8, and 9 of Blue).
- A three of a kind is the next highest hand, consisting of three cards of the same number (for example, the red, blue, and green 8).
- A flush is three cards of the same color, but with no relationship among the numbers (for example, the Yellow 1, 4, and 8).
- A straight is three cards in consecutive order, but with no relationship among the colors (for example, a Red 2, Orange 3, and Purple 4).
- A nothing is the lowest formation, and is any collection of cards that does not fit one of the above rankings. (for example, a Red 2, Orange 4, and Blue 6). Note that a single pair (for example, a Green 2, Blue 2, and Orange 7) counts as a nothing.

When both players have completed formations, if one player’s formation is a higher rank, that player wins. If two players have the same rank, then the highest total wins (so a three of a kind with 7’s beats a 3 of a kind with 4’s). If two players have the same rank and the same total, whoever finishes their hand first wins the flag. In practice among skilled players, most flags are won with a straight flush or three of a kind.



Figure 1: Some examples of formations

Figure 1 illustrates some of these situations. In the leftmost column, the top player has a straight flush, which beats the opposing three of a kind. In the second column, both players have flushes, but the top one totals to 18 and will beat the bottom’s total of 12.

The third column of Figure 1 shows a situation where the top player has played the best possible hand-- the highest possible straight flush (10, 9, 8) that is finished first. Once the player plays the final card of this formation, it cannot be beaten. When this situation occurs, a player is allowed to claim a flag at the beginning of their next turn, even before the opponent has completed their three card hand on the flag. This prevents the opponent from playing any more cards on the flag and is a powerful move, as it reduces the available number of places for the opponent to play cards.

It is also possible in other situations to use the state of the board to prove that a player will not be able to beat a completed formation no matter what Troop cards are played. For example, in the fourth column of Figure 1, the top player has completed a straight totaling 27. The best formation the bottom player could possibly complete is a straight totaling 9. In this case, the player is also allowed to claim the flag at the start of their next turn.

This situation can also arise when a needed card exists elsewhere on the board. In the rightmost formation of Figure 1, the bottom player has two cards towards a straight flush, which if completed would win the flag. However, the cards that can complete the straight flush have already been played (the 8 in the third column and the 5 in the same column by the top player). In this case, the best formation that can be completed by the bottom player is a flush, which will lose to the top player’s three of a kind, and so the top player can claim this flag at the start of their turn.

2.2 Tactics Cards

The deck of 10 Tactics cards provides “special” cards that can bend the rules of the game in various ways. Some are wild cards (and thus can take on a variety of colors or numbers depending on what makes the best possible formation), some change the resolving rules of a flag (for example, making the hands consist of 4 cards instead of 3, or turning all hands into “nothings” so only the highest total wins), and some change the positions of the cards (allowing a player to move their cards between flags, or steal a card from an opponent’s formation). The possibility of an opponent responding to a completed formation with a Tactic card play is the main reason why flags are claimed at the start of the claimers turn.

Tactic cards operate under two important restrictions. First, if a player chooses to draw a Tactic card, they do so instead of drawing a Troop card for that turn. This reduces the player’s hand of playable Troop Cards. Second, a player is not allowed to have played more than one Tactics card beyond what their opponent has played. So, if Player A had played two Tactics cards, and Player B had only played one, Player A would not be able to play another Tactics card until Player B has played a second one.

3 DESIGN OF THE AGENT

Our goal in designing Deep Barca was to make the agent as autonomous as possible, giving it a set of underlying mathematical and probabilistic principles from which moves will be devised. The agent models these principles in several different ways, and these models generate rankings of the possible moves that the agent can make.

3.1 Enforcing the Rules

Before intelligently deciding on the best moves, it is necessary to develop a program that enforces the rules of the game. Many of those rules are straightforward. The most complex involves when a flag can be claimed.

As stated in the previous section, flags can be claimed by a player at the start of their turn if it can be proven that no possible set of cards the opponent can play will beat the player's formation. To determine this, when a formation is finished, the program determines the set of superior formations: formations that can beat the formation that was just completed. Then, the set of possible cards that can complete the formation are examined. These are cards that are currently not on the board (though they may be in player's hands). If any superior formation can be completed by adding these cards to it, then that formation is feasible. A flag can only be claimed if all superior formations are infeasible.

3.2 The Probabilistic Model for Decision Making

Deep Barca uses a general probabilistic model for decision making to play *Battle Line*. Each turn, for each card in its hand and for each flag that card could be placed on, the agent evaluates the probability of winning the game if the card being considered is placed on the flag being considered. It then chooses the option that maximizes this probability. Calculating the probability of winning the game for each move is a three-part process.

First, for each flag the agent calculates the top four formations that each player could make on the flag. These are the four strongest formations that could possibly be made, given the Troop Cards remaining in the agent's hand and the cards left in the deck of Troop Cards.

Next, the agent calculates the probability of winning each flag given the move under consideration. We categorize each flag in one of the following three ways:

- If neither player has played a card towards a formation on the flag, we simply say the probability of winning the flag is 50%.
- If both players have played at least one card towards a formation on the flag, we calculate the probability of winning the flag via what we call the *Multiple Formation Approach*, or *MFA*, described in the next section.
- If one player has played at least one card towards a formation on the flag, but the other player has yet to play a card, we use a mix of the MFA from above and the *Best Single Formation Approach*, or *BSFA*, also described in the next section.

Lastly, we calculate the probability of winning the game as a function of the probabilities of winning each of the nine flags. Since we have two different win conditions (winning five out of nine, and

winning three adjacent flags) we take the weighted average of the probabilities of winning via each of these two conditions as a function of the number of cards left in the Troop Deck. Thus, in the beginning of the game, the agent is far more concerned with simply winning the majority of the flags since this goal in practice entails doing as well as possible on as many flags as possible. Near the end of the game, the agent is more concerned with winning three adjacent flags and will even willingly sacrifice a fringe flag if it means winning three in a row somewhere else.

3.3 The MFA and BSFA Methods

The Multiple-Formation Approach (MFA) is designed to estimate the best hands that can be created once both players have committed to a hand. It is computed for each of the agent's top four formations computed above. For each potential final formation, the agent finds the probability of winning via that formation (the probability of making that formation multiplied by the probability of that formation not being beaten by the opponent). The probability of the agent winning that flag is the union of the probabilities of winning via each of its top formations.

Table 1 shows what these probabilities would look like in the situation where the agent has played a 9 on a flag, and the opponent has played a 4 of a different color. For this table, we assume that no cards in any useful formation exists in the player's hand or on any other flag. The top two formations the agent can make using the 9 are straight flushes (10-9-8 and 9-8-7). If these are successfully made, the opponent cannot possibly beat them using the 4 that was committed to its flag. The next highest formation is three nines. This involves the agent drawing any two of the five remaining nines in the deck, but can be beaten if the opponent draws any of the three possible straight flushes using the 4. The fourth highest formation is a flush totaling 26 (10-9-7). This loses not only to any straight flush, but also to any three 4's.

Formation	Probability of Making	Probability of Opponent beating	Chance to win flag
SF 10-9-8	.25	0	.25
SF 9-8-7	.25	0	.25
9-9-9	.82	.58	.35
Flush 26	.25	.94	.02

Table 1: A sample table used in the MFA for a 9 vs a 4 on an empty board. For details of how the probabilities are computed, see section 3.4

In our experiments, we found that the MFA approach did not work well on flags where one player had not played any cards. This was because on an empty flag all of the top formations in the game are still theoretically possible. Once a card is played on a flag, the space of potential formations is drastically reduced-- instead of being able to create any feasible formation, now only formations that include the played card can be considered. This had the effect of making the agent undervalue playing cards on empty flags. The reason for this is that the MFA does not take into account the need of a player to play a card on a flag (usually reducing the potential value of the formations that can be made on the flag) each turn. Instead, the probability is based on whether the formation can be made by the end of the game.

Thus, to make the agent less pessimistic, in situations where one of the two flags had no cards played, we implemented the Best

Single Formation Approach (BSFA). In the BSFA, we find the single best possible formation for the player with no cards and then evaluate the chances of winning the flag if that formation was played vs the top four formations of the player who has played cards. This simulates the potential “goal” of the opponent who sees a card played and is trying to beat what our agent is trying to do.

The BSFA on its own yields very narrow decision making because the agent will not even consider many decent or adequate formations (since they are not the single best formations that could be made). Often even subpar formations are enough to yield a victory on a particular flag. The MFA on its own is weak as well since it is heavily biased in favor of the player who has yet to play a card on the flag. This is due to the fact that all of the top formations generated for that player are simply the very strongest formations in the game since the player has yet to commit any card that would narrow his or her options. Using a mix of these two approaches, our agent makes much better decisions than it would using either approach by itself.

3.4 Evaluating Formations Probabilistically

In both the MFA and BSFA models described above, the agent needs to determine the likelihood of whether one formation will defeat another, especially in a situation where the formations are currently incomplete, and the cards needed to complete the formation will need to be drawn from the deck. We use the following probabilities to aid our calculations:

- If a card is present on the board, and is in the formation, it has a 100% probability of being part of the final formation.
- If a card is present on the board but is part of another flag, it has a 0% probability of being part of the final formation.
- If a card is not present on the board but is in the agent’s hand, then the agent gives a 100% chance of that card being a part of the formation if the desired formation belongs to the agent and a 0% chance of the card being part of the formation if the desired formation belongs to the opponent.
- If the card is not present in the board or in the agent’s hand, the card must be drawn to be useful. The cards in the Troop deck are assumed to be split evenly among both players. The probability of the agent drawing a needed card is their share of the Troop deck divided by the total number of unseen cards (both in the Troop deck and the opponent’s hand). The complement of that probability is the probability that the opponent will draw (or has drawn) that card. Notice that the agent will do this calculation even if the card is actually in the opponent’s hand, since the agent has no legitimate way of knowing this fact.

3.5 Discarding on a Lost Flag

The above probabilistic behavior forms the bulk of our agent, and thus it chooses moves largely autonomously within the models it is given. However, our models do not cover some areas. Thus, there is a need to develop approaches for these specific situations. Even here, we attempted to make the decisions as autonomous as possible, avoiding low-level pre-programmed rules.

The first situation arises in the turn before claiming a flag. Recall that a flag is claimed at the *start* of a turn, meaning that in the normal course of events, the opponent has a turn to react between the playing of a card that will win a flag and the actual claiming of it. This turn exists to give the opponent a chance to play a Tactics card that may result in changing the resolution of a flag. But even if the player cannot or does not wish to play a Tactics card, they must take a turn. This turn often has rich strategic depth.

Since a flag that is claimed is ineligible to be played upon by either player, even if the formation is incomplete, it usually behooves the player about to lose the flag to play a card on the flag that is about to be lost. This “discarding” action has several benefits:

- Since playing a card on a flag often decreases the possible formations that can be formed, it is usually beneficial to delay committing to a flag as long as possible. Playing a card on the flag that is about to be claimed helps delay decisions about other flags.
- Since proving a flag is unbeatable often involves using cards already played on the board as a reference, discarding is a good chance to play a card that will help a later proving effort.
- Since once the flag is claimed the flag will become unplayable, *not* playing on the flag before it is claimed will mean the owner of the losing flag will have less total spaces to play on in the game (and thus have less total options in the game). This loss of tempo is especially noticeable at the end of the game, since it is possible for one player to have no legal places to play a card and be forced to pass.

Deep Barca handles the situation in which it is about to lose a flag in a different way than it handles the normal course of the game. If it recognizes (via its internal proving mechanism) that it is about to lose a flag, it attempts to find a card to discard. The decision of which card to dump is made by creating a set of “wishlists”-- cards that each possible formation (for each player) are hoping to draw or to play on a flag. The agent will choose a card that appears most frequently on the opponent’s wishlist and least frequently on its own. The hope is that this discarding action will thus not adversely affect any of the top formations the agent is trying to make but possibly will aid in proving that a formation the opponent is trying to make is impossible.

These wishlists are also used by the agent to resolve conflicts where a single card is one of the top formations on multiple flags. By noting what cards are desired by the agent in other places, we reduce the probability that such a card will be available on the conflicted flag.

3.6 Drawing and Playing Tactics Cards

In addition to the Troop deck, Battle Line has a deck of Tactics Cards. These cards do many different things, and if played at the right time, can swing the ownership of a flag from one player to another, for example by playing a wild card to complete a formation that would be proven impossible using played Troop cards. The agent has two decisions regarding Tactics cards. The first decision is deciding when to draw one. The agent will be drawing a Tactics

card instead of a Troop card, and since there are limitations on how many Tactics cards can be played, it is not a good idea to have an excess. The second decision is when to play one. Since there are limitations on the frequency of playing these cards, the timing of the play should be carefully considered.

Since the different Tactics cards are so different and the situations in which they can be played are so different, we had to resort to several ad-hoc rules:

- If playing a Tactic card could win the agent the game, play it.
- If the agent is about to lose a flag, and the loss can be prevented by the play of a Tactics card, play it if it raises the agent's chance of winning the flag to over 50%, or if losing the flag would lose the game.

Other than that, except for very specific situations relating to specific cards, the agent will not play a Tactics card. This has the result of the agent being rather conservative with their Tactics cards, which mirrors the way they tend to be used by expert-level humans.

Similarly, the agent will draw Tactics cards defensively. If it determines that the opponent is one card away from winning a flag that can only be beaten by a Tactics card, it will draw one. It will draw a second if the opponent has two or more cards or if the one tactic card the agent possesses is one of the "less useful" ones. In either case, the agent will not draw a Tactics card that it cannot play.

4 RESULTS

4.1 Evaluating the Quality of Deep Barca

We evaluated the quality of our agent against computer and human players. To the best of our knowledge, there is only one other computer agent, which is the IOS App *Reiner Knizia's Battleline*, by Gourmet Gaming [6]. We played a best-of-three match, in which our agent handily won 2-0. In one of the games, our agent won without losing a flag. Against human players, we tested against a group of 5 experienced human players, and the agent won about 50% of the time. One of the authors is a former three-time world champion at *Battle Line* [2] at the World Boardgaming Championships [3]. We played a 10 game series, in which Deep Barca won three of the games.

With these results we can confidently say that our agent can competitively play against even the very best human *Battle Line* players. Also, our agent was able to make its decision for each turn in between a tenth of a second and one second, which is far faster than a human player could play. This is encouraging, since it leaves the option for more complex decision-making without those decisions taking an undue amount of time.

4.2 Emergent Behavior

Since our agent generally followed a general mathematical model for decision making, we often found that it would make moves that went against our immediate human intuition. In our experiments,

we found some emergent behaviors that were particularly noteworthy.

Eight > Ten: We noticed that early in the game, on an empty flag, the agent often chose to play an 8 on the flag when it could also play a 10. Most high level human players would prefer a Troop Card with value 10 to a Troop Card with value 8, all else equal. The 10 is a higher value, and thus potentially can make hands with higher totals than the 8. However, the 8 is eligible for more straight flushes (10-9-8; 9-8-7; 8-7-6 as opposed to just 10-9-8), which are the strongest formations. Because of this, our agent's preference of an 8 over a 10 is not necessarily bad, and many human players even have this preference. We were pleased to see this behavior emerge, as it was unexpected and not caused by direct human intervention.

Playing low-value viable formations: If the agent has two cards toward a straight flush (for example, the 5 and 6 of red), and the opponent has two cards towards a three of a kind (for example, two 7's), the agent would often give up on the straight flush, going all the way down to a straight (for example, playing a blue 4). This was especially common in situations where many of the cards of the opposing three of a kind were already played (and thus making the three of a kind less likely) or in situations where there was just one possible card that would complete the straight flush.

At first glance, it appeared that the agent was making a mistake giving up a flag so easily. But in reality, once two cards of the same value are played on a flag, the only possible outcomes are a three of a kind or a nothing. The agent has recognized this fact and decided to create a formation that can *only* be beaten by the opponent drawing the card to complete their three of a kind.

While the actual usefulness and timings of these moves against top players can be debated, the ability to find these kinds of moves without having the strategies directly programmed in validates our approach to make the agent largely autonomous and based on underlying mathematical principles.

5 FUTURE WORK

There are many possible improvements we could make to the agent to strengthen quality of play and reduce ad hoc decision making even further.

We would like to improve the way the agent draws and plays Tactics cards to be less ad-hoc.

We would like to improve the agent's ability to make longer-term decisions. Since we evaluate the quality of each move based only on the game state one turn in the future, the agent can often miss moves involving multiple Tactic Cards used together, since the play of any of the cards by itself is weak.

The agent's ability to handle the scenario in which multiple flags are being claimed at the same time is very weak. We would like to be able to evaluate which of the claimed flags are worth fighting for and which of the flags might be too detrimental to lose, but currently we do not.

Our agent keeps the top four possible formations in its MFA model of rating a flag, mainly for efficiency reasons. We would like to examine what would happen if that number was increased.

Since our agent evaluates its moves so quickly, it would be interesting to see if it would be possible to incorporate some lookahead search into the agent, even if just for one or two moves.

ACKNOWLEDGMENTS

This work was partially sponsored by the National Science Foundation's Research Experience for Undergraduates Program, Grant # 1262850

REFERENCES

- [1] BGG. Battle Line. Available online at <https://boardgamegeek.com/boardgame/760/battle-line> (2017)
- [2] BPA. Battle Line Event History. Available online at <http://boardgamers.org/eventhistory/bat.html> (2017)
- [3] BPA The World Boardgaming Championships. Available online at <http://boardgamers.org/wbcindex.html> (2017)
- [4] Buro, M. From Simple Features to Sophisticated Evaluation Functions. In *1st International Conference on Computers and Games*, 126-145 (1999).
- [5] Campbell, M., Hoane, A. & Hsu, F. Deep Blue. *Artif. Intell.* 134,57–83 (2002)
- [6] Gourmet Gaming. Reiner Knizia's Battleline. Available online at <http://www.gourmetgaming.com/mobileBattleline.php> (2017)
- [7] Hall L.; Gordon A.; James R.; and Newall L. A Lightweight Rule-Based AI Engine for Mobile Games. In *Proceedings of the 2004 ACM SIGCHI International Conference on Advances in Computer Entertainment Technology*, 284-289. New York, NY: Association of Computing Machinery. (2009)
- [8] Heyden C. Implementing a Computer Player in Carcassone. Master's Thesis. Department of Knowledge Engineering, Maastricht University, Maastricht, Netherlands. (2009)
- [9] McCulloch, S. A Game-Theoretic Intelligent Agent for the Board Game Football Strategy. *Proceedings of the 26th Modern AI and Cognitive Science Conference*, pp. 121-125. (2015)
- [10] Schadd, M and Winands M. Quiescence Search for Stratego. In *Proceedings of the 21st Benelux Conference on Artificial Intelligence*. Eindhoven, the Netherlands. (2009)
- [11] Schadd, M, Winands M, Bergsma M, and Uiterwijk J. Solving Fanorona. *ICGA Journal*, 30(1):43-45. (2007)
- [12] Schaeffer, J, Culberson, J, Norman Treloar, Brent Knight, Paul Lu, Duane Szafron. A world championship caliber checkers program. *Artif. Intell.* 53,273–289 (1992)
- [13] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwiese, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel and Demis Hassabis. Mastering the Game of Go with Neural Networks and Deep Tree Search, *Nature* 529, 484-489 (2016).
- [14] Thomas, R. Real-time Decision Making for Adversarial Environments Using a Plan-based Heuristic. PhD diss., Department of Computer Science, Northwestern University, Evanston IL. (2003)