Catan: A Probabilistic Approach

Eric Gan, Jialin Li, Dhvani Patel, Samuel Tay

May 23, 2020

Description

Our overall strategy tries to minimize the expected number of turns required to win based on the cost and benefit of the next action. We take into consideration the additional resource generation of building a settlement or city at each possible vertex, and choose the action that minimizes the hitting time to victory if we only purchased cards the rest of the way until winning. This semi-greedy strategy serves as an effective trade-off between generating many resources and quickly gaining victory points by buying cards. This strategy gave us strong results, often taking fewer than 100 moves to win in single player mode, and winning more than 90% of the time in multiplayer mode against the Staffbot on Gradescope.

Assumptions

In the one player game, we determine the optimal next action by assuming that we will only draw cards the rest of the way to win. This assumption simplifies the number of states we need to consider and serves as a reasonable heuristic towards winning. In the two player game, we make similar assumptions to determine our best next action. To simplify our computations, we assume that the opponent is non-adversarial. We maintain legal actions by ensuring our opponent's moves do not conflict with our desired actions at each turn.

Attempted Strategies

Greedy Start Single Player

Our strategy for single player was exactly as we outlined in the proposal. That is, we had a greedy start where we force our Player to buy 2 settlements (and a port if necessary to balance resources) and then upgrade those 2 settlements to a city. These are the initial greedy assumptions we're making and then the rest of the game, we implement actions (buy cards till we win, or buy another settlement + cards till we win, or upgrade to a city + buy cards till we win) chosen to minimize the hitting time to winning. The way we choose a vertex to build a settlement on is by choosing the vertex that generates the most resources based on the probability of landing on that square. The way we choose a settlement to upgrade to a city is by choosing the settlement that has the most probabilistic resource generation.

Safe and Greedy 2P

Our strategy for multiple players is for the most part similar to our single player strategy since our initial strategy is greedy. However, there are some differences since the goal is to first prioritize winning, and then focus on reducing the time to win. With multiple players, building more settlements provides some additional value that put our Player at an advantage versus the opponent. Thus, the threshold between building a settlement and the other actions change based on the following additional values settlements have.

1. **Time:** Multiple player games generally last longer than single player games, and additionally, while it is the opponent's turn, the only thing our Player can do is gain resources. Thus, settlements become a more worthwhile investment with longer games, since our Player cannot make a decision every turn, but they can still gain resources.

2. **Location:** Settlements also have the additional value of denying opponents from resources and preventing them from having access to trading ports.

Final Strategy

Hitting Time

We are now at some state, X_i and we're looking to get to state X_{i+1}. We will get to the next state using the action (choices are actions C, S, and T) that gives us the lowest hitting time. Each hitting time is calculated accordingly: Action C = Buy a card till we reach 10 victory points. Action S = Build a new settlement at a location with lowest hitting time, and then buy cards till we reach 10 victory points. Action T = Upgrade a settlement (if available) to a city at a location with lowest hitting time, and then buy cards till we reach 10 victory points. Our optimal action, A, is then equal to: A = $\min_{C,S,T}(E[C], E[S], E[T])$, and we use action A to get from state X_i to state X_{i+1}. Following our proposed strategy, we first considered using Markov Chains to solve for the hitting time. However, considering the number of states the Markov chain has (depending on our current points and resources of each type) is very large, we seek the alternative of calculating the hitting time entirely based on our current resource production rate using the probability distribution of dices. In this way, calculating the hitting time takes constant time. The difference between our method and using Markov chain should be insignificant since the production rate is constant, and we are calculating the hitting time for not only the next point, but to the end of the game, so the number on average should be large, making the error more negligible.

Then we considered the resources we need to build a settlement as the combined resources for building two roads and one settlement. If this yields the lowest hitting time, we start building the roads leading to the settlement, then the settlement. We do not calculate the heuristic again until we finish building the settlement. While this strategy works well with a single player, it causes an infinite loop when playing against an adversary, since the opponent can beat us to building the settlement we're aiming to build over the course of our next few turns.

To prevent the infinite looping, we decided that instead of combining the process of building two roads and a settlement, we now consider only one step further: either building a road or a settlement. We modified the method of calculating hitting times to accommodate the number of roads needed (0, 1, or 2). In this way, even when a potential settlement is first claimed by the opponent, we can continue with other options.

Ports and Trading

Once we determined our next action given the hitting times, we would then try to accumulate the resources to accomplish that action. Since rolls were probabilistic, there was a decent chance we would obtain a surplus of resources we did not need and could potentially lose them to the robber. Thus, if we had a surplus of resources that we didn't need for the next action, we would do a resource balance where we would attempt to trade those resources we didn't need in for the resources we actually needed to achieve that action faster.

When deciding to go for a trading port, there were two considerations we took in. We first ensured that the port remained a Manhattan distance of 2 away from one of our current settlements, since otherwise it would not be worth building such a long road to obtain it. Then if the port was a 2:1 port, then we would ensure that our resource production levels were above a certain threshold (for our final action we decided on 0.5) before deciding it was worth buying that port. This is because otherwise we wouldn't generate enough of that specific resource required to make use of the 2:1 deal, and it would not be worth building a settlement that only gets resources on one specific die roll. If the port was the special 3:1 port, then we decided to first have 2 settlements built before buying that port. This is because we wanted to build as far inside as possible to prevent being blocked off by the opponent, and ports really only become useful once the resource generation is high enough to make use of the trade.

Results

Our final results for multiplayer (against the staff bot) are: 126 bot has a win rate of 0.0, with average 0 turns per won game and student has a win rate of 1.0, with average 107.18 turns per won game (resulting in a 10/10 autograder score!). For our single player, we get an average number of 79.926 turns per game, for 10 trials.