

Video Notes

We hope you can get a better understanding of the environment Lux and the strategies in our paper by watching the video and the accompanying text illustration.

Game 1 (0:00-0:24)

In this game, you can see that despite a tie in the first 60 rounds, our policy quickly built up an advantage over TB after that, which is realized by the strategies we describe in our paper:

- Carefully choose locations when building citytiles.
- Accurate planning for resource collection and consumption.
- The special Five-diagonal shape as in Turn 199-240 (0:12-0:17 in this video)

Game 2 (0:24-0:46)

In this game, you can see our policy masters both offensive and defensive strategies and are good at balancing them for the final victory.

- In the beginning, our policy and TB's choose to build Great Wall in the middle of the board to protect resources from the opponent.
- However, while successfully guarding our own resources, our policy also sends one worker for the opponent's resources. At turn 68 (0:27 in the video), this worker reaches TB's coal and manages to build a CityTile and more workers. We think this difference mainly contributes to the overall victory.
- This video illustrates that the core of this game is about resource occupation, the winner must protect his own resources and also grab the opponent's resources.

Game 3 (0:47-1:10)

This game on a 32x32 map illustrates the efficiency of transfer learning over different map sizes. With only around 1 million game frames, our policy can achieve a 95% win-rate v.s. TB from a 10% win rate in the beginning, which demonstrates the efficiency of this transfer learning.

Game 4 (1:10-1:36)

This game demonstrates our policy's desperate all-in-for-the-win strategy. At turn 335, the number of our and TB's citytiles is 234:261, but after only 15 turns, as our policy begins to build as many cities as possible, at turn 350, the numbers of citytiles become 293 and 269. It is amazing that our agents build 60 citytiles within 15 turns, which is key to the final result..

Game 5 (1:37-3:05)

This game shows the generalization to even bigger maps. In this game, the policy we use is the policy model trained on map size 32 without any fine-tuning. The policy shows the amazing ability of building large-scale cities, at turn 357(3:01) there are 976 and 733 CityTiles of each team, which demonstrates the generalization of our proposed environment and model.