The Leaderboard

The Zero Robotics Leaderboard has adapted over the years, and is now based on the Whole History Rating (WHR) system. The WHR approach tracks rank throughout each phase of the competition in order to rank players on the leaderboard based on score.

Score is equal to mean rank. However, this rank is not the ranking you see on the leaderboard, but rather a numerical indicator of your game performance. We will refer to this as "true rank." As you improve game performance, your true rank will *increase*, while your leaderboard rank will *decrease* to approach 1. Mean rank is simply the average of all of your true ranks after every submission.

After every submission, your team's algorithm, or "player", will be matched against 20 other teams' submitted players. These 20 teams are the teams with true ranks closest to yours. Based on your current true rank and the true ranks of the opposing teams, the likelihood of you winning each match is calculated. This is known as your "win probability." Your new true rank is then computed as a function of match history and rank history, taking every submission into account. In short, if you perform better than expected, your true rank will increase, or if worse than expected, fall.

Thus, not all wins are equal. A win over a well-ranked team is more valuable than a win over a beginner team. Similarly, a loss to a beginner team will hurt you more than a loss to an experienced team. Wins and losses are valued by their relation to the projected win probability. Your score corresponds to a 50% win probability, meaning you are expected to win against teams with lower scores than yours, but not against teams with higher scores. Unexpected match results, either favorable or unfavorable, will leave the most noticeable impacts on your score.

Although this process is based on probability, it is wholly reliable. The rank computation algorithm corrects errors in rank history throughout the competition, and optimization and stabilization safeguards have been added to calculate rank with extreme precision. Also, the leaderboard is designed to allow data to propagate throughout the system. If two players don't

run against each other because their rank difference is more than 20 slots, overlapping matches update both teams' scores. Between submissions, at night, and at the end of every competition phase, batch optimizations are made to update all scores and ensure that every team's rank is relative to all other teams' ranks.

The foundation of the leaderboard rests on the principle that player skill improves with time. As the competition progresses, ranks become increasingly susceptible to change. Allowing scores to increasingly diverge as time progresses effectively separates players by skill. In the latter phases of the competition, including the 3D and Alliance rounds, the leaderboard allows significantly more divergence to help identify the best teams.

Revisiting the score equation, the team with the highest score will be rank 1 on the leaderboard. The best way to stabilize and even improve your rank is the most logical way: keep working at your algorithms. There are no surefire alternatives. Do not try to maintain your rating by repeatedly submitting the same player, nor should you aim for ties. Identical submissions are grouped and will yield no rank change, and there is a penalty for tying a game. Also, don't let fear of a bad match ruining your team's prospects keep you from submitting frequently. Even though every submission is factored into your match history, results of past competitions demonstrate that teams that make many submissions tend to perform better. Also, as another safeguard, the leaderboard protects against large rank fluctuations. Finally, it is important to note that after every competition phase, all ranking data is refreshed and teams start from scratch.