

Rewind the Robot

Train it, save every version, and roll back when the data goes bad.

OBJECTIVE

You're training a helper robot. Each round you add one dataset to a version of its brain and save the result as a new version. Datasets raise or lower four skills: Spot, Chat, Walk, Safe. Then an Event card can reveal that a dataset was bad, and every version built with it changes. After 6 rounds, release your best version. Highest total skill wins.

QUICK START – READ THIS, START PLAYING

- 1 Your Version Log starts with v0, the base model: every skill 1, total 4. Lay 3 Dataset cards face up and shuffle the 6 Event cards face down.
- 2 On your turn, take one face-up Dataset. Pick any version on your log, copy its datasets into the next row, add the new one, and write which version you built from. That's a commit.
- 3 Add up each skill: 1 + every dataset's number. Skills stay between 0 and 5. Refill the face-up row.
- 4 When everyone has played, flip one Event. If it changes a dataset, fix every version that has it. After 6 rounds, circle the version you release. Highest total wins.

YOUR TURN

- 1 COMMIT: take one face-up Dataset card. Choose any version on your log as the parent: the newest one, or an older one if you'd rather branch.
- 2 Write the next version number, its parent, and every dataset in it: the parent's datasets plus the new card's code.
- 3 Work out each skill: 1 + the numbers from every dataset in that version. A skill can't go below 0 or above 5. Write the total.
- 4 OR MERGE instead of taking a card: pick two versions. The new version gets every dataset from both, counted once each. Work out its skills the same way.
- 5 Turn a new card face up so there are 3 again. Keep the Dataset card in front of you so everyone can see what you trained on.

HOW TO SCORE

- 1 Test Day: circle ONE version on your log to release. It can be any version without an X, even v0.
- 2 Your score is that version's total: Spot, Chat, Walk, Safe added up (20 is perfect). Highest score wins. Ties share the win.
- 3 Solo: score 15 or more, then try to beat your best with a new shuffle.

TIPS

- 1 Release your best version, not your newest. In 500 practice games, picking the best version on the log beat shipping the newest one 303 times and never did worse.
- 2 Roll back. When an Event hurts your newest version, build on an older one instead. Doing that beat always building on the newest 336 times and lost 5.
- 3 A dataset that adds a lot to one skill can quietly take from another. Try it on a branch and keep the old version.
- 4 Every Event card is public. Count which datasets are still under suspicion before you lean on them.

WHAT THE GAME TEACHES

- 1 A model is its data. The same recipe with different datasets gives a different robot, so teams write down exactly what went into every version.
- 2 Problems in training data often show up later: mislabeled photos, rude chats, data nobody agreed to share. One bad dataset changes every version built with it.
- 3 Teams keep old versions so they can roll back. Newer isn't always better.