Build-a-Rule Tournament

Teams compete to write the shortest rule that correctly sorts a training set — then the rule gets tested on cases it has never seen. A playful, competitive way into overfitting and generalization.

AGE BAND	DURATION	GROUP SIZE	CONCEPT
Ages 10–14	30 minutes	Teams of 3–4	Learning

Learning objective

Students can explain why a rule that perfectly fits the examples it was built from can still perform poorly on new examples.

What you need

- A 'training set' of 12 example cards per team, each with 3–4 traits and a yes/no label
- A sealed 'test set' of 8 new cards per team, same format, held back until round two
- Pencils and scratch paper

Why this works

A rule that explains every example in front of you can feel like the best possible rule. It is often just the rule most tightly stitched to the examples you happened to have. This tournament makes that distinction visible and scored, which is a more durable lesson than being told about it.

How it works

1. Each team receives 12 training cards — fictional characters with 3–4 traits and a “yes/no” label (for example, “let this character into the club: yes/no”).
2. Teams write the shortest rule that correctly predicts every training card’s label, then present it to the class.
3. Score round one on brevity and how confidently each team can defend their rule.
4. Reveal the sealed set of 8 new cards. Teams apply their rule and score round two on how many labels they get right.
5. Compare standings. The shortest, most confident round-one rule does not always win round two.

Talk about it

- Did the team with the best-looking rule in round one also win round two? What happened?
- What’s the risk of writing a rule that fits your training examples perfectly?
- Real automated systems are graded the same way this tournament is — on data they haven’t seen yet. Why does that matter more than how well they explain the data they trained on?