Competition Highlights
Canadian Mathematical Olympiad
and
Junior Olympiad
(CMO/CJMO)

Paweł Prałat* James Rickards†

August 2025

The Canadian Mathematical Olympiad (CMO) is an annual, invitational, proof-based competition for Canadian students. It is Canada's premier national advanced mathematics competition. Students attempt to solve 5 problems in three hours, with each problem graded on a scale from 0 to 7. In 2020, the CMS introduced the Canadian Junior Mathematical Olympiad (CJMO), also by invitation only, a variant specifically for students in grade at most 10. These 3-hour competitions are held each March at a selected time and date (by default, the second Thursday of March). All official participants write at the same time and are proctored by their local school faculty or staff. For more information visit `https://cms.math.ca/competitions/cmo/`.

The CMO is an important contest for students with international aspirations, as a good performance leads to the Canadian Team Selection Test, and then onto the International Mathematical Olympiad (IMO) itself. Qualification for the C(J)MO is primarily via the Canadian Open Mathematics Challenge (COMC), an open contest written in late October.

In total, the 2025 CMO was written by 95 students, with 93 official entrants. The CJMO was written by 17 students, all official entrants. Five Canadian provinces were represented, with the number of contestants as follows:

CMO: AB (8), BC (13), ON (47), QC (8), SK (1)
CJMO: AB (1), BC (1), ON (12), QC (1)

(note that Canadian citizens residing outside of Canada can also officially write the C(J)MO, accounting for the discrepancy in numbers).

*Department of Mathematics, Toronto Metropolitan University, Toronto, ON, Canada; e-mail: `pralat@torontomu.ca`

†Department of Mathematics and Computing Science, Saint Mary's University, Halifax, NS, Canada; e-mail: `james.rickards@smu.ca`

Grading for both contests went relatively smoothly, with a team of 15 mathematicians, including professors, students, and former contestants, contributing their time. The top score on the CMO was 27, achieved by Warren Bei, and the mean score was 8.1. The Matthew Brennan Award for best solution went to Warren Bei for an excellent solution to problem 3. This was a difficult problem with most solutions being fairly technical. Warren's writeup was very short and clean. On the CJMO, a top score of 23 was achieved by Warren Maximilian Lin, and the mean score was 10.1. A full breakdown of the marks assigned problem by problem is in Table 1.

Score	P1	P2	P3	P4	P5
7	0	34	8	5	5
6	3	9	4	0	2
5	4	9	1	0	0
4	2	6	0	0	0
3	27	10	0	0	0
2	14	8	4	1	1
1	26	4	1	0	2
0	19	15	77	89	85
Avg	1.91	4.33	0.99	0.39	0.54

(a) CMO

Score	P1	P2	P3	P4	P5
7	0	2	0	1	0
6	15	0	0	0	1
5	1	0	1	0	0
4	0	0	0	2	0
3	1	0	3	4	0
2	0	0	3	1	0
1	0	2	2	1	0
0	0	13	8	8	16
Avg	5.76	0.94	1.29	1.76	0.35

(b) CJMO

Table 1: C(J)MO score breakdown by problem.

An intrepid reader may note that CMO problem 1 (which was also CJMO problem 3) was significantly harder than usual, with no perfect 7's awarded. This was due to an unintentional clerical error. The question entailed a voting procedure among n hockey players, which took place in rounds. The question was intended to read "Prove that eventually, all players will unanimously vote for the same person." Instead, the final version asked to prove that "after n rounds, all players will unanimously vote for the same person." Despite several checks, this subtle change went unnoticed until after the contest.

Fortunately, the version with n rounds is still a correct problem, just more appropriately placed as the fourth or fifth problem on the contest. The increased difficulty had a large knock-on effect, with overall scoring averages being about a problem lower than last year's C(J)MO. It also offers a good lesson to future olympiad contestants: do not always trust the ordering of the problems! It is very common for test-setters to misestimate the relative difficulty of a contest, leading to misnumbered problems. Even a contest like the IMO occasionally has problems out of order (with the most famous example being 2011's "windmill problem").

Problem 5 on the CMO was a neat problem about an ant traveling around a rectangle.

Problem 1. *A rectangle R is divided into a set S of finitely many smaller* *rectangles with sides parallel to the sides of R such that no three rectangles in S share a common corner. An ant is initially located at the bottom-left corner of R. In one operation, we can choose a rectangle $r \in S$ such that the ant is currently located at one of the corners of r, say c, and move the ant to one of the two corners of r adjacent to c.*

Suppose that after a finite number of operations, the ant ends up at the top-right corner of R. Prove that some rectangle $r \in S$ was chosen in at least two operations.

A natural approach to solving the problem is to prove the contrapositive: assume that the ant does not choose a rectangle twice, and prove that it cannot travel from the bottom left to the top right.

A good next step is to draw a few sample dissections S, and attempt to walk the ant to the top right following the moves. For example, see Figure 1.

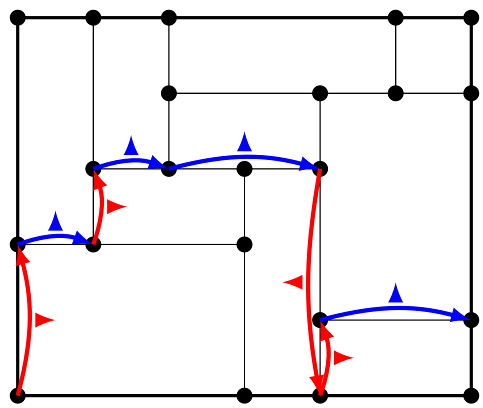

Figure 1: A possible path by the ant. Arrows indicate direction of travel, and point to the chosen rectangles.

One will find that the ant can travel pretty far, in fact, it seems like it can visit most vertices in S, and it can get very close to the top right corner. Furthermore, if we drop the repeated rectangle rule, it is very easy to walk to the top right. This signals that there must be some sort of *invariant* that is blocking the ant.

Such an invariant is clearly broken when choosing the same rectangle twice. Let us examine the given conditions a bit more closely: there is the clause that "no three rectangles in S share a common corner." This is clearly very important, as the problem would be false without it! See Figure 2 for a simple path that reaches the top-right corner of R, and does not choose a rectangle twice.

What is special about forbidding such a three corner intersection? This implies that all corner intersections are "T-intersections": they form a pattern like $\vdash$, $\dashv$, $\perp$, or $\top$. In particular, to reach an intersection, there are exactly two possible rectangles that can be chosen. One rectangle must be chosen to reach

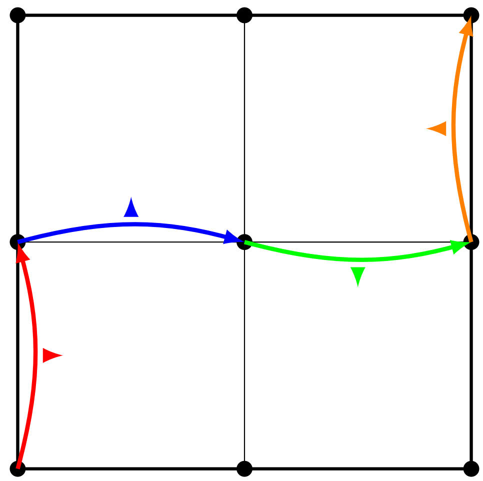

Figure 2: How the ant can win, if three corners may intersect. (Note that three corners intersecting implies that a fourth must also intersect.)

the intersection point, and the other is therefore chosen to leave it, since we cannot repeat rectangles.

In fact, this is exactly where the non-repeating condition gets used. It also implies a strengthening of the problem: in order for the ant to get to the top right, not only did the same rectangle have to be used twice, but it had to be used twice *consecutively*!

Heading back to the problem, it seems like we are close. We have identified a key property of the setup, which limits the possible moves, and need to form this into an invariant that is somehow preserved. There is very little to work with here, other than the location of the ant's path relative to the chosen rectangle.

In particular, we can track if the ant is traveling vertically or horizontally, and if the chosen rectangle is to the left or to the right of the ant as they travel. This makes four combinations, and we can keep track of them as we walk (an analogous invariant would be tracking the cardinal directions, e.g. SW corner to NW corner). Colour the ant's moves as follows:

- **Red**: vertical, with chosen rectangle on the right (with respect to the ant's travel);
- **Orange**: vertical, with chosen rectangle on the left;
- **Green**: horizontal, with chosen rectangle on the right;
- **Blue**: horizontal, with chosen rectangle on the left.

In Figure 2, we made it to the top-right corner of R, and we see that all four colours were used. On the other hand, in Figure 1, there are no intersections of three rectangles, we followed the no rectangle repeats rule, and only red and blue were used!

At this point it is clear that this observation must lead to a solution. Indeed, note that the first move (from the bottom left) must choose the bottom left rectangle, and either be vertical and right (red), or horizontal and left (blue). By analyzing the different T-intersection possibilities, we find that if the ant walks a blue or a red path, their next move is again blue or red, hence all future moves are as well. See Figure 3 for a demonstration of the possibilities after a move north or east (180 degree rotation gives the south and west cases).

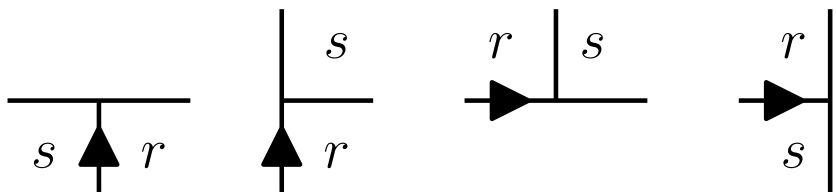

Figure 3: Possible ant moves going north or east. The ant enters on rectangle r, and leaves on rectangle s. There are two possible choices, and each one remains blue or red.

Finally, we must go back to the original problem. Why can the ant not reach the top right corner? Well, the final move must be to choose the upper-right rectangle, which is left of the vertical move to make it there, and right of the horizontal move. These moves are coloured green and orange, which cannot occur after a blue or red move. This provides the contradiction we require!

A takeaway of this problem is that even very difficult questions can sometimes be solved by a series of small observations. Nowhere in this solution did we need to come up with a difficult and clever idea: it all fell apart naturally from a slow and methodical investigation into how the ant can travel.