Firefighter spatial reasoning: study guide

Public safety / fire service · suite apt-057-firefighter-spatial-reasoning · generated 2026-09-15T15:26:06.090Z

Detail	Value
Title	Firefighter spatial reasoning: study guide
Generated	2026-09-15T15:26:06.090Z
Fixture/version	apt-057-firefighter-spatial-reasoning
Sector	Public safety / fire service
Guide version	v2

> Private practice result. Not an official exam certificate, employer decision, hiring signal, admissions decision, or guaranteed outcome. Scores stay on this device unless you export them.

> Answer keys and scoring logic stay server-side and are never included in any download or export.

How to use this guide

This file contains the whole study outline for this suite: every skill it draws on, the full lesson for each of those skills, worked examples, practice tips, a glossary, and where each piece of material comes from. Nothing here is a summary of a page you still have to visit.

1. Read the skills section end to end once, without timing yourself.
2. Work the guided practice mode for the suite, using the practice tips as a checklist.
3. Move to the mini-test only when guided practice feels unhurried.
4. Sit the full simulation last, once, in the conditions you expect on the day.

Practice attempts are stored on the device you used, never on an account. Exporting a result is the only way anything leaves that device.

Practice modes and durations

Practice modes for Firefighter spatial reasoning

Mode	Duration	What it is for
Guided practice	15 minutes	Untimed, with feedback after every item.
Mini-test	18 minutes	A short timed set for checking pace.
Full simulation	45 minutes	Full length and full time, in one sitting.

- Start guided practice: https://learn.novusstreamsolutions.com/aptitude/practice/new?bank=apt-057-firefighter-spatial-reasoning&mode=guided
- Start mini-test: https://learn.novusstreamsolutions.com/aptitude/practice/new?bank=apt-057-firefighter-spatial-reasoning&mode=mini
- Start full simulation: https://learn.novusstreamsolutions.com/aptitude/practice/new?bank=apt-057-firefighter-spatial-reasoning&mode=full

Skills covered, in full

This suite draws on 5 skill constructs. Each one below carries its complete lesson.

Spatial reasoning

Rotation, folding, views, orientation, maps, and three-dimensional visualization.

Spatial reasoning is the ability to turn an object in your head and be right about the result: to tell a rotation from a mirror image, fold a flat net into a solid and know which faces meet, read plan and elevation views back into a three-dimensional shape, and hold your own orientation steady while a route turns underneath you. It is a heavily weighted construct in pilot and aircrew selection, firefighter and military technical batteries, architecture and design admissions, and trades screening where a drawing has to be built. Much of it can be trained by replacing imagination with a checkable procedure. Practice is device-local unless you export it.

What you should be able to do after this lesson:

1. Distinguish a rotation from a reflection by checking whether the clockwise order of three marked features is preserved or reversed.
2. Read a cube net and name which faces end up opposite, using the rule that squares two apart in a straight strip are opposite.
3. Reconstruct a solid from plan, front and side views, count its unit cubes, and identify when the views leave the count undetermined.
4. Predict the hole pattern produced by unfolding a punched sheet, treating each unfold as a reflection across the fold line rather than a rotation.
5. Work with bearings and relative direction, converting between a map frame and a traveller frame without losing track of which is which.
6. Name the cross-section produced by cutting a cylinder, cone or cube on a given plane, including the oblique cuts people default to calling circles.

Worked examples and pitfalls

Rotation preserves handedness; reflection reverses it: A figure carries three distinguishable marks (say a red band, a blue band and a yellow band), which read red, blue, yellow going clockwise around its centre. Rotating that figure in the plane, by any angle, never changes that clockwise order: it still reads red, blue, yellow clockwise from whatever mark you start at. Mirroring it reverses the order, so a reflected copy reads red, yellow, blue clockwise. That single check answers the whole is-this-a-rotation-or-a-reflection family without imagining any motion. It works on letters, which is why no in-plane rotation of the letter R produces a backwards R, and no rotation of a lower-case b produces a d: b and d are mirror images across a vertical line, and turning one will never give you the other. Contrast b and q, which are related by a half turn and therefore genuinely are rotations of each other. The distinction is exactly what the clockwise-order test measures. In three dimensions the equivalent is handedness at a corner: pick three edges meeting at a vertex, and a rotation preserves whether they form a right-handed or left-handed set while a reflection swaps it. This is exactly why most Western dice, viewed at the corner where the 1, 2 and 3 faces meet, show those numbers running anticlockwise, and a die showing them clockwise is a mirror-image die rather than a differently rotated one. The one honest limitation: a figure with its own mirror line looks identical after reflection, so the test decides nothing for symmetric shapes, which is precisely why item writers use asymmetric ones.

Cube nets: two apart in a strip means opposite: A net made of six squares. Four of them, call them A, B, C and D, sit in a horizontal strip left to right. Square E sits directly above B, and square F sits directly below B. Fold it. The strip of four wraps around the cube as a band, so squares two apart in that strip end up on opposite faces: A is opposite C, and B is opposite D. E and F are the remaining pair, and they become the top and bottom, opposite each other. The two rules that solve almost every net item are exactly these. Squares separated by one square in a straight line end up opposite. Squares sharing an edge on the net end up adjacent on the cube, and can never be opposite. Everything else is applying them carefully. A die adds one more constraint worth carrying: opposite faces of a standard die sum to seven, so 1 faces 6, 2 faces 5 and 3 faces 4, and a net showing 1 and 6 two apart in a strip is consistent while a net showing 1 and 2 two apart is not a standard die. The classic error on these items is trying to fold the net in your head all at once. Do not. Pick one square as the base, identify its opposite by the two-apart rule, then repeat for a second pair, and you have determined the cube in two cheap steps rather than one expensive act of visualisation.

Three views, one solid, and the cubes you cannot see: A solid is built from unit cubes on a base one cube deep and three cubes wide. The column heights, left to right, are 3, 1 and 2. The front view is therefore a staircase profile of heights 3, 1, 2. The side view, looking along the row from the right, is a single column three cubes tall, because the tallest column hides the others behind it. The plan view is a three-by-one rectangle. Total cubes: 3 plus 1 plus 2, which is 6, and here the three views pin it down exactly because the plan view tells you the solid is only one cube deep. Now change one thing. Suppose the plan view is three by two, so the solid is two cubes deep. The front view can still read 3, 1, 2 while the count is anything from 6, one row of cubes with a hollow back, up to 12, a solid block matching the front profile in both rows. The views constrain the shape but do not determine it, and a well-written item will either supply the count or ask for the minimum and maximum rather than a single figure. That is the trap in this family: candidates who assume every hidden position is filled will over-count, and candidates who assume nothing is hidden will under-count. Read the plan view first to establish the footprint, then use the elevations to cap each column, then ask explicitly whether the item wants the minimum, the maximum, or a determined value.

Punched and unfolded, with coordinates: Take a square sheet 8 units across, with the origin at its centre so the corners are at plus or minus 4 in each direction. Fold the right half onto the left half across the vertical centre line, then fold the bottom half up onto the top half across the horizontal centre line. What is left visible is one quarter, the region where x runs from minus 4 to 0 and y runs from 0 to 4. Punch a single hole at the point minus 3, 3. Now unfold, one step at a time, and treat every unfold as a reflection. Undoing the horizontal fold reflects across the line y equals 0 and adds a hole at minus 3, minus 3. Undoing the vertical fold reflects both existing holes across x equals 0 and adds holes at 3, 3 and 3, minus 3. The finished sheet has four holes, at all four combinations of plus or minus 3. Two rules generalise from this. The number of holes is one punch multiplied by two for each fold the punch passed through, so two folds give four holes and three folds give eight, provided the punch went through every layer, which it does when the folds are through the middle. And the positions are mirror images across each fold line, never rotations. That second point is where the marks go: candidates who place the fourth hole by rotational symmetry get three holes right and one wrong, which is exactly the answer option the item includes.

Bearings, and the difference between the map and the walker: A walker leaves a gate, goes 300 m due north, then 400 m due east. The straight-line distance back to the gate is the hypotenuse of a 300 by 400 right triangle, which is 500 m: the 3-4-5 triangle, so no calculator is needed. The bearing from the gate to the walker is measured clockwise from north, and it is the angle whose tangent is 400 over 300, which is about 53 degrees, so the bearing is 053. The return bearing is that plus 180, which is 233. Bearings are always three digits and always clockwise from north, so 53 degrees is written 053 and a bearing of 350 is very slightly west of north, not south. Now the other half of this construct, which trips up more people than the trigonometry. You are facing south and you turn left. Which way are you now facing? East. Your left hand points east when you face south, and the standard wrong answer, west, comes from reading left as left-on-the-map while forgetting that the map is drawn north-up and you are not. The general fix is to state your heading before every turn and rotate your frame with the traveller: facing north, left is west; facing south, left is east; facing east, left is north; facing west, left is south. On route items, physically rotate the page so the direction of travel points up before working out a turn, and rotate it back only to read off the final compass direction.

What a cut actually reveals: Cross-section items ask what shape a plane makes as it slices a solid, and the reflex answer of circle is wrong most of the time. A cylinder cut perpendicular to its axis gives a circle. Cut parallel to the axis, it gives a rectangle. Cut obliquely so the plane crosses only the curved surface, it gives an ellipse; cut obliquely so the plane also crosses one flat end, it gives an ellipse with one side sliced flat, a D-shape. A cone is richer, because its sections are the conic sections that give the family its name: perpendicular to the axis gives a circle; an oblique cut that crosses every side of the cone gives an ellipse; a cut parallel to one slanted side gives a parabola; a cut parallel to the axis gives a hyperbola branch; and a cut through the apex gives a triangle. A cube is the one that surprises people. A cut parallel to a face gives a square, obviously. A cut through three vertices that are mutually adjacent to one corner gives an equilateral triangle. And a cut through the midpoints of six edges, perpendicular to a long diagonal of the cube, gives a regular hexagon: a six-sided section from a six-faced solid with no curves anywhere. For a flat-faced solid, work these by counting faces rather than by visualising: the section of a convex polyhedron gains exactly one straight edge for every face the plane passes through, so a cut crossing six faces of a cube must close into a six-sided figure and a cut touching only three cannot have more than three sides. For curved solids the question becomes whether the plane meets the curved surface all the way round, which gives a closed circle or ellipse, or only partway, which leaves an open arc closed off by a straight chord.

How to practise this skill

- Replace visualisation with a rule wherever one exists. Two-apart-means-opposite on nets, and clockwise-order-preserved for rotations, both give a checkable answer in seconds and are reliable when mental imagery gets tired late in a section.
- Physically rotate the paper or the device on route and orientation items so the direction of travel points up, work out the turn, then rotate back to read the compass direction. Fighting the north-up convention in your head is the single largest source of left-right errors.
- On view and cube-count items, read the plan view first to fix the footprint, then use the elevations to cap each column, then decide explicitly whether the question wants a minimum, a maximum, or a determined count.
- Treat every unfold as a reflection across the fold line. Say the word reflection out loud as you draw each new hole; the rotational version of the pattern is a deliberately included wrong answer in almost every folding item.
- Build a physical reference once. Fold a paper cube net, mark the faces, and hold it: ten minutes with a real cube fixes the opposite-face rule more firmly than an hour of diagrams, and the memory transfers to items you have never seen.
- Practise untimed first and check whether your errors cluster in rotation, folding, views or orientation, because the four subskills train differently. Attempts are stored on this device only, so an honest error log costs nothing and is what tells you which drill to repeat.

Glossary

Chirality (handedness): The property that distinguishes a figure from its mirror image. Rotation preserves it and reflection reverses it, which is what makes the clockwise-order test decisive for asymmetric figures.

Net: A flat arrangement of faces that folds into a solid. On a cube net, squares two apart in a straight strip become opposite faces and squares sharing an edge become adjacent ones.

Plan view: The view looking straight down on an object. It fixes the footprint of a stacked-cube solid, which is why it is the first view to read.

Elevation: A view looking horizontally at an object, usually given as front and side. Each elevation caps the heights it shows but hides anything directly behind.

Orthographic projection: A drawing convention showing an object through separate parallel views with no perspective. Together the views constrain a solid but do not always determine it.

Bearing: A direction measured clockwise from north and written with three digits, so due east is 090 and a heading slightly west of north is 350. The reverse bearing differs by 180 degrees.

Egocentric versus allocentric frame: The difference between directions relative to a traveller who is turning and directions relative to a fixed map. Left means different compass headings in the two frames, which is why headings must be restated before every turn.

Cross-section: The shape revealed where a plane cuts a solid. It gains one straight edge for each flat face the plane crosses and one curve for each curved surface, which turns the question into a count.

Study this next

- Spatial reasoning skill hub: https://learn.novusstreamsolutions.com/aptitude/skills/spatial-reasoning
- Spatial reasoning practice bank: https://learn.novusstreamsolutions.com/cognitive-skills/spatial-reasoning
- Mechanical comprehension lesson: https://learn.novusstreamsolutions.com/aptitude/skills/mechanical-comprehension/lesson
- Firefighter spatial reasoning suite: https://learn.novusstreamsolutions.com/aptitude/tests/firefighter-spatial-reasoning
- Architecture and design aptitude suite: https://learn.novusstreamsolutions.com/aptitude/tests/architecture-and-design-aptitude
- Spatial reasoning lesson on Novus Learn: https://learn.novusstreamsolutions.com/aptitude/skills/spatial-reasoning/lesson

Where this material comes from

- Worked items written for Novus Learn from standard introductory geometry: cube-net opposite-face rules, orthographic view reconstruction, reflection across fold lines, right-triangle distances and bearings, and the conic and polyhedral cross-sections. The 3-4-5 distance, the 053 and 233 bearings, and the four hole coordinates were each computed and checked.
- Terminology checked against the public Wikipedia articles 'Net (polyhedron)', 'Multiview orthographic projection', 'Chirality', 'Bearing (navigation)' and 'Conic section'. Definitions only; every figure described above is original.
- Novus Learn aptitude construct registry (catalog seed) for the construct scope and related suite mapping.
- Public educational framing only: not affiliated with any official exam board, publisher or employer, and no copyrighted test item is reproduced.

> Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Mechanical comprehension

Forces, motion, gears, pulleys, levers, fluids, pressure, and basic machines.

Mechanical comprehension is the ability to predict what a physical system will do (which way a gear turns, how hard you have to pull, how far the load actually rises) from forces, moments and the simple machines, rather than from a memorised formula sheet. It carries real weight in apprenticeship entry batteries, military technical selection, and screening for maintenance, machine operation and rigging roles. The same reasoning is daily work on site: sizing a jack, choosing a block-and-tackle arrangement, deciding why a belt slips under load but not at idle. Practice here is device-local: no account, and nothing leaves this device unless you export it.

What you should be able to do after this lesson:

1. Balance moments about a pivot (effort x effort arm = load x load arm) and name whether a given tool is a first-, second- or third-class lever.

2. Work out the speed, torque and rotation direction of a driven gear from tooth counts, including what an idler gear does and does not change.
3. Count the rope parts that support a moving block to get the mechanical advantage of a pulley system, then state how much rope must be pulled per metre of lift.
4. Apply pressure = force / area to a hydraulic jack and predict the piston-travel penalty that comes with the force gain.
5. Predict rpm across a belt or chain drive from pulley or sprocket diameters, and say when slip breaks the prediction.
6. State the trade-off that governs every simple machine: multiplied force always costs distance, and friction always makes the delivered advantage smaller than the ideal one.

Worked examples and pitfalls

Levers: balance the moments, then name the class: A wheelbarrow carries 60 kg whose centre of mass sits 0.4 m from the wheel axle, and you grip the handles 1.2 m from that same axle. Take moments about the axle: lift force x 1.2 m = 60 kgf x 0.4 m, so the lift force is 60 x 0.4 / 1.2 = 20 kgf, about 196 N. The mechanical advantage is simply the arm ratio, 1.2 / 0.4 = 3. Two traps live in this item. The first is measuring the load arm from your hands instead of from the fulcrum; the arm is always the perpendicular distance to the pivot, and the pivot here is the wheel contact, not the barrow body. The second is calling it a first-class lever because that is the one everyone pictures. A wheelbarrow is second class: fulcrum at one end, load in the middle, effort at the far end, which is why its mechanical advantage is always greater than one. Compare tweezers or a pair of tongs, where the effort sits between fulcrum and load: that is third class, mechanical advantage below one, and you are deliberately trading force away to buy speed and control at the tip. Your own forearm lifting a weight is the same arrangement, which is why a 5 kg dumbbell loads the biceps far more than 5 kg.

Gear trains: tooth counts set speed, meshes set direction: A 12-tooth driver turning at 300 rpm meshes directly with a 36-tooth gear. The ratio is driven teeth over driver teeth, 36/12 = 3:1, so the output turns at 300/3 = 100 rpm and, ignoring friction, carries about three times the torque. One external mesh reverses rotation, so the output turns opposite to the driver. Now drop a 20-tooth idler between them. Step it through: 300 x 12/20 = 180 rpm at the idler, then 180 x 20/36 = 100 rpm at the output. The overall ratio is unchanged at 3:1, the idler's tooth count cancels, but there are now two external meshes, two reversals, so the output turns the SAME way as the driver. That is the whole reason idlers are fitted. The tempting wrong answer treats the idler as another reduction stage and reports 180 rpm or some product of both ratios. The check that never fails: only the first and last gear in a simple train affect the ratio, and the direction depends on whether the number of external meshes is odd (reversed) or even (same). An internal or ring mesh, as in a planetary set, does not reverse at all.

Pulleys: count the rope parts that carry the load: A 200 kg load, about 1,962 N. Hung from a single pulley bolted to a beam, the pulley only changes the direction you pull; both rope parts still meet at a fixed axle, the mechanical advantage is 1, and you pull the full 200 kgf. Hang the pulley on the load instead, with one rope end anchored above and the other in your hands, and two rope parts now support the moving block: mechanical advantage 2, effort about 100 kgf or 981 N, but you must pull 2 m of rope for every 1 m the load rises. Build a tackle with four parts supporting the moving block and the effort falls to 200/4 = 50 kgf, about 490 N, at the cost of 4 m of rope pulled per metre of lift. The trap is counting sheaves instead of supporting parts. A three-sheave arrangement gives three parts if the dead end is made off to the fixed block and four if it is made off to the moving block, and the answer differs by 33 percent. The second trap is treating the figure as delivered force: real sheaves lose a few percent each to bearing and rope friction, so quoted mechanical advantage is the ideal velocity ratio, and the effort you actually feel is higher.

Hydraulics: pressure is shared, force and travel are traded: A jack has a 2 square centimetre input piston and a 50 square centimetre output ram. Push the small piston with 100 N and the pressure in the fluid is 100 / 2 =

50 N per square centimetre, which is 500 kPa. Pascal's principle says every part of the confined fluid sees that same 500 kPa, so the large ram feels 50 N/cm2 x 50 cm2 = 2,500 N. Mechanical advantage 25. Nothing is free: fluid is effectively incompressible, so volume in equals volume out. A 25 cm stroke on the small piston moves 2 x 25 = 50 cubic centimetres, and 50 cubic centimetres spread over a 50 square centimetre ram raises it 1 cm. Twenty-five centimetres of pumping buys one centimetre of lift: exactly the force-times-distance bargain a lever strikes. Candidates lose this item by assuming both pistons travel the same distance, or by concluding the jack manufactures energy. A related item asks about a tank: the pressure at the bottom depends on the height of fluid above and its density, not on the width of the vessel, so a narrow 3 m standpipe reads the same bottom pressure as a 3 m deep swimming pool.

Belts and chains: same belt speed, different rpm: A motor pulley 100 mm in diameter runs at 1,450 rpm and drives a 250 mm pulley through an open V-belt. The quantity that is genuinely shared is belt speed, not rpm, so the driven pulley turns at 1,450 x 100/250 = 580 rpm. Note the direction of the ratio: with gears you divide by driven teeth, with belts you divide by driven diameter, and the arithmetic looks the same only because both are proportional to circumference. An open belt drives both pulleys the same way; a crossed belt reverses the driven pulley, which is the entire point of the classic crossed-belt diagram item. The trap is slip. A flat or V-belt under an overload slips, so the driven speed falls below 580 rpm while the motor holds 1,450: the symptom of a glazed belt or a weak tensioner. A chain or a toothed timing belt physically cannot slip a tooth without damage, which is why camshaft and indexing drives use them and why a belt-driven answer and a chain-driven answer to the same question can legitimately differ.

Springs in series and parallel: the answer most people reverse: Two identical coil springs, each 20 N/mm. Mount them side by side, both carrying the same load, and they act in parallel: stiffness adds to 40 N/mm, so a 400 N load compresses the pair 400/40 = 10 mm. Now stack them end to end, in series. Each spring carries the whole 400 N, force is not split down a chain, so each deflects 400/20 = 20 mm and the total is 40 mm. The combined stiffness is 400/40 = 10 N/mm, half of one spring. Most candidates reverse the two because 'series' sounds like the case where things add. Two reliable checks: in parallel the arrangement is always stiffer than one spring, and in series it is always softer than the softest spring. That is also why suspension designers add a helper spring in parallel to raise rate under load, and why a long slender bolt clamps more forgivingly than a short stubby one carrying the same preload.

How to practise this skill

- Sketch the free body and mark the pivot before writing a single number. Mechanical items are lost by choosing the wrong fulcrum far more often than by mis-multiplying.
- Fix the three lever classes to a tool you can picture (pliers and a see-saw for first class, a wheelbarrow and a nutcracker for second, tweezers and your own forearm for third), and say the class out loud before answering.
- Sanity-check every answer against the work rule: if a machine multiplies force it must divide distance. An answer where the effort both moves less and is smaller than the load is wrong somewhere.
- Convert units before you multiply, not after. A moment computed with millimetres on one side and metres on the other is off by a thousand and looks perfectly plausible.
- For any direction question, walk the mesh or the belt with a finger: each external gear pair reverses, an internal or ring mesh does not, an open belt keeps direction and a crossed belt flips it.
- Work untimed until you can explain the reasoning aloud, then add the clock. Attempts are recorded on this device only, so a messy first pass through gear trains costs you nothing but time.

Glossary

Moment (torque): A force multiplied by the perpendicular distance from its line of action to the pivot, in newton-metres. A body is in rotational balance when clockwise moments equal anticlockwise moments about the same point.

Fulcrum: The pivot a lever turns about. Its position relative to the load and the effort defines the lever class and therefore whether the tool multiplies force or multiplies movement.

Mechanical advantage: Output force divided by input force. Ideal mechanical advantage is the geometric figure; actual mechanical advantage is what you get after friction, and it is always lower.

Velocity ratio: Distance moved by the effort divided by distance moved by the load. In a lossless machine it equals the ideal mechanical advantage, which is why force gain always costs travel.

Gear ratio: Driven teeth divided by driver teeth. A ratio above one reduces speed and multiplies torque; only the first and last gear of a simple train affect it.

Idler gear: A gear placed between driver and driven purely to add a mesh. It reverses the output direction and bridges a distance without altering the overall ratio.

Pascal's principle: Pressure applied to a confined fluid is transmitted undiminished to every part of it, which is what lets a small piston at high pressure produce a large force on a bigger ram.

Head: The height of a fluid column above a point. Pressure at depth is density x gravity x height and is independent of the vessel's width or shape.

Study this next

- Mechanical comprehension skill hub: https://learn.novusstreamsolutions.com/aptitude/skills/mechanical-comprehension
- Millwright and industrial mechanic suite: https://learn.novusstreamsolutions.com/aptitude/tests/millwright-and-industrial-mechanic
- Spatial reasoning lesson: https://learn.novusstreamsolutions.com/aptitude/skills/spatial-reasoning/lesson
- Skilled trades, maintenance and repair careers: https://learn.novusstreamsolutions.com/careers/families/skilled-trades-maintenance-and-repair
- Encyclopedia search: simple machine: https://learn.novusstreamsolutions.com/search?q=Simple%20machine
- Mechanical comprehension lesson on Novus Learn: https://learn.novusstreamsolutions.com/aptitude/skills/mechanical-comprehension/lesson

Where this material comes from

- Worked examples written for Novus Learn from standard introductory mechanics: moments about a pivot, gear ratios from tooth counts, rope-part counting for tackle, and Pascal's principle.
- Terminology checked against the public Wikipedia articles 'Simple machine', 'Lever', 'Gear train' and 'Pascal's law'. Definitions only; every number and item above is original.
- Novus Learn aptitude construct registry (catalog seed) for the construct scope and related suite mapping.
- Public educational framing only: not affiliated with any official exam board, publisher or employer, and no copyrighted test item is reproduced.

> Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Attention and concentration

Sustained focus, selective attention, and accuracy under time pressure.

Attention and concentration is the skill of still noticing on minute forty of a checking block as reliably as you did on minute two, and of pointing the noticing at the right thing when two things compete. It has three distinguishable parts - selective attention, sustained attention or vigilance, and divided attention - and assessments load them differently: cancellation and checking tasks load vigilance, conflict tasks load selection, and dispatch-style monitoring loads division. It is the construct that decides whether a records clerk, a dispatcher, an air-side controller or a quality inspector catches the one wrong digit in a shift. Practice here is device-local, with no account and nothing uploaded.

What you should be able to do after this lesson:

1. Count occurrences of a target in a dense string without double-counting or losing your place, using a fixed scan path rather than free looking.
2. Separate hits, misses and false alarms in your own results and work out whether you are set too eager or too cautious for the scoring rule in force.
3. Predict where in a long block your error rate will rise, and schedule micro-resets against that curve instead of pushing through it.
4. Recognise a transposition error, the class that survives a same-characters gist check and is therefore missed most often.
5. Explain why an automatic process such as reading interferes with a controlled one, and use that to anticipate which distractors will actually cost you time.
6. Run a two-stream monitoring task by alternating on a fixed cadence rather than attempting genuine simultaneity.

Worked examples and pitfalls

A cancellation count you can actually check: Count every 7 in this row: 4 7 1 7 7 3 9 7 2 8 7 5 7 6 0 7. Working left to right and tapping once per hit: hits at the second, fourth, fifth, eighth, eleventh, thirteenth and sixteenth positions. That is seven sevens. Two error modes produce nearly all the wrong answers. The first is the adjacent pair - the 7 7 at positions four and five gets counted once, because the eye takes a repeated character as one perceptual object. The second is losing the place after the 9, where the visually similar 9 and 7 force a moment of re-checking and the scan restarts a character early, producing an over-count of eight. The defence for both is a fixed scan path with a physical anchor: a fingertip or cursor moving one character at a time, never jumping back. Re-scanning to check is the thing that creates the double-count, so if you must verify, verify by counting a second time from the RIGHT-hand end and comparing totals, never by re-reading part of the row.

Hits, misses and false alarms: the eager candidate loses: A checking batch contains 300 records, 40 of which are genuinely faulty. Candidate A flags 46 records and 36 of them are truly faulty. Hits 36, misses 40 minus 36 equals 4, false alarms 46 minus 36 equals 10. Candidate B flags 38 and 34 are truly faulty: hits 34, misses 6, false alarms 4. On hit rate alone A wins, 36 out of 40 which is 90 percent against B's 34 out of 40 which is 85 percent. Now apply the scoring rule that most checking tasks actually use, where a false alarm cancels a hit: A scores 36 minus 10 equals 26, B scores 34 minus 4 equals 30. B wins by four despite catching two fewer faults. This is why 'flag anything that looks odd' is bad advice on a scored checking task, and why the first thing to read on a checking item is whether wrong flags are penalised. Your response criterion - how much evidence you demand before flagging - is adjustable, and it should be set from the scoring rule, not from your temperament.

The transposition that passes every gist check: Compare these two lines and decide whether they match. Invoice 4820-7391-06. Invoice 4820-7931-06. They do not: the middle group reads 7391 in the first and 7931 in the second, with the 3 and the 9 swapped. Transpositions are the most-missed error class in record checking for a structural reason - the character SET is identical, the length is identical, the first and last characters of the group are identical, so every fast check the visual system runs comes back clean. Substitutions and omissions change the character inventory and get caught; transpositions do not. Two habits raise the catch rate. Read digits in fixed groups of two rather than as a whole number, so 73-91 against 79-31 becomes a mismatch at the first group instead of a subtle difference somewhere in a four-digit blur. And check groups in a deliberately non-natural order - last group, first group, middle group - because reading left to right lets the confirmation you built at the start carry you through the middle, which is exactly where the swap is usually planted.

Conflict: why reading fights you: The classic demonstration is Stroop's, published in 1935: the word RED printed in blue ink, with the instruction to name the ink colour. Naming takes measurably longer, and errors go up, because reading a familiar word is automatic and cannot be switched off, so the automatic response has to be suppressed before the controlled one can be produced. The same conflict has a numeric version you can test on yourself in a second: how many characters are in the string 4 4 4? The answer is three, and the digit 4 pulls at you the entire way. In an assessment this appears wherever the salient feature and the asked-for feature come apart - a chart where the tallest bar is not the answer to the question, a form where the highlighted field is not the one being verified, a row where the bold total is not what the stem requested. The practical move is to name the target feature out loud before you look - 'ink colour', 'character count', 'the value for March' - because pre-loading the target biases selection before the automatic reading response gets a chance to win.

Where the errors actually appear in a 45-minute block: Errors in a long checking block are not spread evenly. Mackworth's 1948 clock-watching study established the pattern that gives the effect its name: detection declines over a prolonged watch, with the sharpest deterioration early rather than at the very end. Practically, on a 45-minute self-timed checking block, expect your per-minute error rate in minutes 20 to 45 to run visibly above minutes 1 to 20 even though nothing about the material changed, and expect the subjective sense of effort to lag the actual decline, so it will not feel like you are getting worse. Two things work against it. Break the block into three fifteen-minute segments with a ten-second reset between them - look away, unfocus, breathe out - which costs thirty seconds of a 45-minute block, about one percent of the time, and buys back more than that in caught errors. And score your practice by segment rather than as one number, because a single overall accuracy figure hides exactly the information you need: whether your problem is skill, which shows as flat error rate, or endurance, which shows as a rising one.

Two streams: alternate on a cadence, do not try to merge: A dispatch-style monitoring task: keep a running total of the numbers announced on channel one while watching channel two for the code word AMBER. Genuine simultaneity is not available - the two tasks compete for the same control resource - so the choice is not whether to alternate but whether to alternate deliberately or accidentally. Deliberate looks like this: fix the arithmetic to a rhythm, updating the total only at each announcement and holding a single number between updates, which frees the gaps for channel two. If the total is 34 and the next announcement is 7, you spend under a second reaching 41 and then you are free again. Accidental looks like re-deriving the running total from the beginning because you did not trust it, which locks up the whole window and is when AMBER goes past unheard. The measurable failure of divided attention is almost never a failure to hear the target; it is a failure to have any spare capacity at the moment it arrived. The related phenomenon worth knowing is inattentional blindness, illustrated by Simons and Chabris in 1999: an unexpected and perfectly visible event is missed entirely when attention is committed to a counting task.

How to practise this skill

- Fix your scan path before you start, and never re-scan backwards to double-check. Backward re-scanning is what produces both double-counts and lost places; if you need to verify a count, run it again from the other end and compare.
- Read the scoring rule before the first item. Whether false alarms are penalised changes the correct strategy completely, and a criterion set for the wrong rule costs more marks than slow reading does.
- Check numeric strings in fixed groups of two or three and in a deliberately shuffled group order. Transpositions are the errors that survive whole-string comparison, and they are the ones planted on purpose.
- Time your practice in segments and log accuracy per segment, not per session. A flat error rate means you need speed; a rising one means you need endurance and breaks, and the two problems have opposite fixes.
- Name the target feature out loud before each conflict item. Pre-loading the relevant dimension is the cheapest available defence against the salient-but-wrong answer.

- Practise with an actual distractor present rather than in silence. A test hall has movement, coughing and page-turning, and attention practice in perfect quiet trains a condition you will not meet.

Glossary

Selective attention: Prioritising one source or feature while suppressing competing ones. It is what fails when the visually salient element of a display captures the response instead of the requested one.

Sustained attention (vigilance): Maintaining detection performance over a long, low-event period. It is the capacity that checking, monitoring and inspection tasks are really sampling.

Vigilance decrement: The decline in detection performance over a prolonged watch, typically steepest early in the block. Mackworth's 1948 clock test is the standard public reference.

Divided attention: Handling two streams that compete for control. Performance is best modelled as fast alternation with a switch cost, not as genuine parallelism.

Hit, miss, false alarm, correct rejection: The four outcomes of any detection decision: flagging a real fault, missing one, flagging a clean record, and correctly leaving a clean record alone. Any honest accuracy claim needs at least the first three.

Response criterion: How much evidence you require before flagging. Shifting it trades misses against false alarms without changing your underlying sensitivity, which is why it must be set from the scoring rule.

Stroop interference: The slowing that occurs when an automatic response, such as reading a word, conflicts with the requested response, such as naming its ink colour. Named for Stroop's 1935 experiments.

Inattentional blindness: Failing to notice a fully visible, unexpected event because attention is committed elsewhere - the effect Simons and Chabris demonstrated in 1999 with a counting task.

Study this next

- Attention and concentration skill hub: https://learn.novusstreamsolutions.com/aptitude/skills/attention-concentration
- Cognitive Skills Lab: attention practice: https://learn.novusstreamsolutions.com/cognitive-skills/attention-concentration
- Emergency dispatch and 911 communications suite: https://learn.novusstreamsolutions.com/aptitude/tests/emergency-dispatch-and-911-communications
- Error detection lesson: https://learn.novusstreamsolutions.com/aptitude/skills/error-detection/lesson
- Perceptual speed lesson: https://learn.novusstreamsolutions.com/aptitude/skills/perceptual-speed/lesson
- Attention and concentration lesson on Novus Learn: https://learn.novusstreamsolutions.com/aptitude/skills/attention-concentration/lesson

Where this material comes from

- Character strings, invoice comparisons, batch counts and monitoring scenarios above are written for Novus Learn. All figures are original and no published test item is reproduced.
- Terminology and directions of effect follow widely published work: Mackworth (1948) on the vigilance decrement, Stroop (1935) on response conflict, Simons and Chabris (1999) on inattentional blindness, and the standard hit/miss/false-alarm framing from signal detection theory. Concepts only; no result figures are attributed to those studies.
- Novus Learn aptitude construct registry (catalog seed) for construct scope and suite mapping.
- Public educational framing only - not affiliated with any official exam board, publisher or employer, and not a clinical, diagnostic or attention-disorder screening measure.

> Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Map and route reasoning

Directions, coordinates, schedules, routes, and navigation constraints.

Map and route reasoning is spatial thinking with a clock attached: bearings and turns, map scale, grid references, and the journey-time arithmetic that decides whether a route is feasible at all. It appears in police, fire and ambulance selection, driver and dispatch testing, military navigation batteries, and any logistics or field-service role where a wrong turn is a missed window. The items are rarely hard mathematically; they are lost by rotating the map instead of the number, by averaging two speeds, or by adding a walking leg to the wrong end of a connection. Your practice attempts are stored on this device and go nowhere unless you export them.

What you should be able to do after this lesson:

1. Convert any facing into degrees clockwise from north, apply a sequence of left and right turns arithmetically, and convert the result back to a compass point.
2. Compute a back bearing and use it to check that an out-and-return leg is consistent.
3. Convert map distance to ground distance from a ratio scale such as 1:25,000, and say why the ratio has no units.
4. Give and read four-figure and six-figure grid references in the correct order, and state which square or point a reference identifies.
5. Compute door-to-door journey time from distance, average speed and fixed stops, and work backwards from a fixed deadline to a latest departure.
6. Compare candidate routes against stated constraints and identify which constraint actually binds on the vehicle, the load or the time window.

Worked examples and pitfalls

Bearings: turn the number, not the map: You are facing north. Turn 90 degrees right and you face 090, east. Turn a further 135 degrees right: 90 + 135 = 225, south-west. Now turn 90 degrees left: 225 - 90 = 135, south-east. The method that survives any number of turns is to hold the facing as a number of degrees clockwise from north, add for right turns, subtract for left turns, take the result modulo 360, and convert back at the end using the eight-point table: 0 north, 45 north-east, 90 east, 135 south-east, 180 south, 225 south-west, 270 west, 315 north-west. Candidates who instead picture themselves rotating lose the left-right convention the moment they are facing south, because on the mental image left now appears on the right of the page. The companion idea is the back bearing: the reverse of a leg is the forward bearing plus or minus 180 degrees, so a leg travelled out on 040 comes home on 220, and a leg out on 250 comes home on 070. Checking that every outbound and return pair differs by exactly 180 catches a dropped turn immediately.

Scale: from the sheet to the ground: On a 1:25,000 sheet, one centimetre of paper represents 25,000 centimetres of ground, which is 250 metres. A path measured at 6.4 cm is therefore 6.4 x 250 = 1,600 m, or 1.6 km. Move to a 1:50,000 sheet and the same 6.4 cm becomes 3.2 km, because each centimetre now stands for 500 m. Two mistakes recur. The first is dividing rather than multiplying, which turns a 1.6 km walk into a 26 m one and should be caught instantly by asking whether the answer is plausible for a footpath. The second is forgetting that a ratio scale carries no units: 1:25,000 means one inch to 25,000 inches just as much as one centimetre to 25,000 centimetres, which is exactly why the printed scale bar is safer than a ruler when a sheet mixes conventions. Finally, map distance is horizontal distance. A route that climbs 300 m over that 1.6 km is meaningfully longer and much slower on foot, which is why walking-time estimates add time for ascent separately rather than inflating the distance.

Grid references: along the corridor, then up the stairs: On a numbered map grid, the four-figure reference 3572 means easting 35 and northing 72. It identifies a whole one-kilometre square, and specifically the square whose south-west corner sits at that intersection: not its centre. To locate a point inside it, divide the square into tenths in each direction: a feature five tenths east and five tenths north of that corner is 355725 in six figures, which pins the position to roughly 100 m. The two failures are reliable enough to be worth naming. The first is giving northings before eastings, so 7235 gets quoted for a feature that is actually at 3572, pointing at an entirely different square or at nothing at all. The second is treating the reference as the centre of the square, which shifts every reported position by half a kilometre to the north-east. The mnemonic that fixes both is 'along the corridor, then up the stairs': you walk along before you go up, and eastings always come first.

Journey time with a stop, and the average-speed trap: A van covers 84 km at an average 56 km/h. Driving time is 84/56 = 1.5 hours, or 90 minutes. Add a 20-minute unload part way and the door-to-door time is 110 minutes, so a 08:20 departure arrives at 10:10. If the receiving bay closes at 10:00, the total budget is only 100 minutes; with the 20-minute stop unavoidable, driving must fit into 80 minutes, and the required average becomes 84 / (80/60) = 63 km/h. That is the shape of most feasibility items: find the fixed time, subtract the fixed stops, and solve for what is left. The trap sits in the word average. Suppose the same 84 km is driven as 42 km at 70 km/h and 42 km at 50 km/h. The mean of the two speeds is 60 km/h, and that is the wrong answer. Add the times instead: 42/70 = 0.60 h and 42/50 = 0.84 h, total 1.44 h, so the true average is 84/1.44 = 58.3 km/h. Average speed is always total distance over total time, never the mean of the speeds.

Routing under constraints: find the one that binds: Depot to site, three options. Route A is 18 km and crosses a bridge posted at 7.5 tonnes. Route B is 14 km but the final 300 m is one-way against you, forcing a 1.2 km loop, so it is really 15.2 km. Route C is 16 km and includes 800 m of school zone limited to 20 km/h between 08:00 and 09:00, where you would otherwise do 40 km/h. Your loaded van weighs 3.5 tonnes and you leave at 08:30. Price the constraints. The bridge limit does not bind at all: 3.5 tonnes is well under 7.5, so route A's restriction costs nothing and A is simply the longest at 18 km. The school zone does bind, but barely: 0.8 km at 20 km/h takes 2.4 minutes against 1.2 minutes at 40, a penalty of 1.2 minutes, leaving C at 16 km plus about a minute. Route B, at 15.2 km including its loop, is still the shortest. The trap is choosing A because it is the option with no restriction that applies, a constraint that does not bind is not a cost, and noticing that is precisely what the item measures.

Timetables: work backwards from the fixed event: Buses leave the depot at five and thirty-five minutes past each hour. The ride to the station takes 25 minutes, then a 6-minute walk to the platform. Trains depart at 09:12, 09:42 and 10:12. To catch the 09:42, you must be on the platform by 09:42, so the latest boarding time is 09:42 minus 25 minus 6 = 09:11. The 09:05 bus qualifies: it reaches the station at 09:30 and the platform at 09:36, six minutes clear. The 08:35 bus is on the platform at 09:06 and would also make the 09:12 with a little to spare. The 09:35 bus reaches the platform at 10:06 and can only make the 10:12. The two errors that cost marks are subtracting the walking time from the wrong end (adding it to the train departure instead of the bus arrival, which manufactures twelve minutes of slack that does not exist), and reading a timetable's arrival column as though it were the connection's departure. Fix the deadline first, list every fixed time in one column, then choose.

How to practise this skill

- Convert every facing to degrees clockwise from north before you do anything else. Arithmetic on a number survives a question that turns you round four times; mental rotation does not.
- Never rotate the paper. Rotating the map reverses your left and right the instant the route heads south, and that single habit accounts for a large share of lost bearing marks.
- Write the scale conversion once at the top of the page, '1 cm = 250 m', and reuse it for the whole set. Recomputing it per item is where factor-of-ten errors are born.
- For timetable items, work backwards from the fixed deadline and write every fixed time in a single column before you look at the options. Forward-adding invites you to forget the walk or the changeover.
- Never average two speeds. Add the times, then divide total distance by total time. The arithmetic mean of the speeds is almost always one of the offered answers, and it is there deliberately.
- When a route question lists a restriction, ask explicitly whether it binds on this vehicle, this load and this

departure time. Practise untimed until spotting a non-binding constraint is automatic; attempts stay on this device, so repetition costs nothing.

Glossary

Bearing: A direction given as degrees measured clockwise from north, conventionally written with three digits so that east is 090 and south-west is 225.

Back bearing: The bearing of the return leg, equal to the forward bearing plus 180 degrees if it is under 180, or minus 180 if it is over. A leg out on 040 returns on 220.

Ratio scale: A unitless map scale such as 1:25,000, meaning one unit measured on the sheet represents 25,000 of the same units on the ground, whatever unit you choose.

Easting and northing: The horizontal and vertical grid coordinates of a position. Eastings are always quoted first, which is what the 'along the corridor, then up the stairs' mnemonic encodes.

Six-figure grid reference: A four-figure square reference refined by tenths in each direction, locating a point to roughly 100 m rather than to a one-kilometre square.

Dead reckoning: Fixing a position from a known starting point plus the recorded headings, speeds and elapsed times, without any external fix to correct the accumulated error.

Average speed: Total distance divided by total elapsed time. When a question asks door to door, the elapsed time includes every stop, not only the time in motion.

Contour interval: The fixed height difference between adjacent contour lines. Tightly spaced contours mean steep ground, and therefore a journey time longer than the flat map distance suggests.

Connection time: The interval between an arrival and the next scheduled departure, including any walk or transfer. It is the constraint that actually decides most timetable items.

Study this next

- Cognitive Skills Lab: map and route practice: https://learn.novusstreamsolutions.com/cognitive-skills/map-route-reasoning
- Map and route reasoning skill hub: https://learn.novusstreamsolutions.com/aptitude/skills/map-route-reasoning
- Route and timetable reasoning suite: https://learn.novusstreamsolutions.com/aptitude/tests/route-and-timetable-reasoning
- Supply chain, transport and logistics careers: https://learn.novusstreamsolutions.com/careers/families/supply-chain-transport-and-logistics
- Spatial reasoning lesson: https://learn.novusstreamsolutions.com/aptitude/skills/spatial-reasoning/lesson
- Map and route reasoning lesson on Novus Learn: https://learn.novusstreamsolutions.com/aptitude/skills/map-route-reasoning/lesson

Where this material comes from

- Navigation and scheduling examples computed for Novus Learn from first principles: clockwise bearing arithmetic, ratio-scale conversion, grid-reference convention, and distance-speed-time with fixed stops.
- Terminology checked against the public Wikipedia articles 'Bearing (navigation)', 'Scale (map)', 'Ordnance Survey National Grid' and 'Dead reckoning'. Definitions only; every route, timetable and figure above is original.
- Novus Learn aptitude construct registry (catalog seed) for the construct scope and related suite mapping.
- Public educational framing only: not affiliated with any official exam board, publisher or employer, and no copyrighted map extract or timetable is reproduced.

> Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Safety judgment

Hazard recognition, safe sequencing, escalation, and risk controls.

Practise recognizing hazards and choosing a safe, proportionate response: protect people, control immediate exposure, use the correct reporting path, and verify that the control is effective.

What you should be able to do after this lesson:

1. Separate a hazard from the likelihood, severity, and exposure that shape its risk.
2. Choose an immediate control that stays within the person's training and authority.
3. Recognize when work should pause and when a supervisor, emergency process, or qualified specialist is needed.

Worked examples and pitfalls

Worked scenario: a damaged machine guard: A guard is loose before a scheduled run. Do not operate the machine or improvise a repair beyond your authorization. Keep people away, isolate or label the equipment only as procedure permits, report the defect to the responsible person, and wait for an approved inspection or repair. A deadline does not remove the hazard.

Risk-triage questions: Ask: what can cause harm, who is exposed now, how severe could the outcome be, what control is available, and who has authority to apply it? The safest answer is not always the most dramatic option; it is the option that controls the real exposure without creating a new hazard.

How to practise this skill

- Treat warnings, permits, isolation rules, and personal protective equipment as parts of a system, not interchangeable shortcuts.
- Do not ask an untrained person to investigate a hazard simply because they are nearby.
- In real work, site procedures and qualified safety direction override any general preparation heuristic.

Glossary

Hazard: A source or situation with the potential to cause harm.

Risk: A judgement about possible harm that considers likelihood, severity, and exposure.

Control: A measure that removes a hazard or reduces exposure to it under an approved process.

Study this next

- Occupational safety and health: https://learn.novusstreamsolutions.com/search?q=Occupational%20safety%20and%20health
- Hazard: https://learn.novusstreamsolutions.com/search?q=Hazard
- Risk assessment: https://learn.novusstreamsolutions.com/search?q=Risk%20assessment
- Safety judgment lesson on Novus Learn: https://learn.novusstreamsolutions.com/aptitude/skills/safety-judgement/lesson

Where this material comes from

- Novus Learn original safety-critical, hazard-recognition, and operational suite scenarios.
- Novus educational framework: identify the hazard, protect people, control exposure, report, and verify.

> Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Recommended preparation

- Build spatial reasoning: https://learn.novusstreamsolutions.com/aptitude/skills/spatial-reasoning/lesson
- Practice mechanical comprehension: https://learn.novusstreamsolutions.com/aptitude/skills/mechanical-comprehension/lesson
- Learn safety judgement: https://learn.novusstreamsolutions.com/aptitude/skills/safety-judgement/lesson

Answer keys, scoring and privacy

> Answer keys and scoring logic stay server-side and are never included in any download or export.

Formal suite answer keys and scoring logic stay on the server and are not part of any download, in any format. Downloadable keys exist only for the open, untimed practice material (the practice packs on the puzzles, cognitive-skills and reasoning practice lab pages), where the answers are already public teaching content.

Novus Learn needs no account. Your practice history lives in this browser's storage on this device, and is sent to a server only if you create an account and switch on backup. This file contains no attempt, session or result link, so it is safe to share.

Private practice result. Not an official exam certificate, employer decision, hiring signal, admissions decision, or guaranteed outcome. Scores stay on this device unless you export them.

Answer keys and scoring logic stay server-side and are never included in any download or export.

Source: https://learn.novusstreamsolutions.com/aptitude/tests/firefighter-spatial-reasoning/study-outline