

The Algorithm Will See You Now

When there is no specialist for two hundred kilometres, what counts as good enough?

7 min read

Before you read

- **triage** — sorting patients by urgency when resources cannot treat everyone at once
- **false negative** — a result that misses a condition that is in fact present
- **threshold** — the cut-off point at which a system changes its decision
- **accountability** — being answerable for a decision, including when it is wrong
- **deferral** — passing a decision upward rather than making it
- **counterfactual** — what would have happened instead, had the thing not been done



A clinical officer reads a retinal screening result, rural Kenya.

¹ The clinic sits at a junction four hours from the nearest referral hospital. On a Tuesday morning the queue holds perhaps ninety people, and there is one clinical officer to see them. He has eleven minutes each if nobody is seriously ill, which somebody always is.

² Since last year he has had help. A screening tool on a tablet reads retinal photographs and flags signs of diabetic eye disease, a condition that causes blindness slowly, silently, and entirely preventably if caught. The tool takes under a minute per patient. There is no ophthalmologist within two hundred kilometres.

³ This is the setting in which arguments about medical artificial intelligence become serious, and it is not the setting in which they are usually conducted. Most debate assumes a well-staffed hospital where the question is whether a machine should replace an expert. Here no expert is being replaced. The comparison is not machine against specialist. It is machine against nothing.

⁴ That reframing is powerful, and it is also the point at which careful thinking tends to stop. It is worth pushing further, because several of the objections survive it.

⁵ Start with accuracy. These systems are trained on large collections of images, and performance depends on how closely new patients resemble the training population. A model built largely on images from one region may perform measurably worse elsewhere — different equipment, different lighting, different patterns of disease. The failure is not dramatic. It shows up as a slightly higher rate of false negatives: people told they are fine who are not.

⁶ A false negative in a screening programme is a particular kind of harm. The patient does not merely fail to receive care; the patient receives reassurance, and reassurance changes behaviour. Someone who has been told their eyes are healthy does not make the four-hour journey next year.

⁷ Against that stands the counterfactual, and it is unforgiving. Without the tool, that patient is not examined at all. A screening system that misses one case in ten still finds nine that nobody was going to find. Refusing to deploy it until it performs as well as a specialist is, in practice, choosing the worse outcome in order to keep the cleaner conscience.

⁸ So far the argument is about numbers, and numbers can be improved. The harder objection is not about accuracy at all.

⁹ Consider what happens when the tool is wrong. A specialist who misses a diagnosis can be questioned, can explain the reasoning, can be held to a professional standard and, where necessary, can be found at fault. A model that misses a diagnosis produces no reasoning. The clinical officer who accepted its output did not make the judgement; the engineers who built it never saw the patient; the ministry that approved it evaluated a statistic, not a case. Responsibility has not been assigned to anyone. It has been dissolved.

¹⁰ This matters beyond fairness to the individual patient. Accountability is how systems improve. A fault that can be traced generates a review, a change in practice, a correction. A fault that belongs to no one generates a note in a log.

¹¹ There is a second effect, slower and harder to measure. A tool that is right most of the time trains the people around it to stop checking. Deferral is rational for an overloaded officer: the machine is usually right, and there are eighty-nine people still waiting. But a judgement that is never exercised does not stay sharp, and the skill that was meant to be supplemented is quietly replaced.

¹² None of this amounts to an argument against deploying the tool. The counterfactual is still unforgiving, and nine found cases are nine people who keep their sight. It is an argument against the way such deployments are usually justified — as a straightforward improvement, with the question closed once the accuracy figure clears a threshold.

¹³ The threshold is a beginning, not a conclusion. The questions that follow it are whether the people who bear the cost of an error had any say in choosing that threshold; whether a wrong result can be traced to anyone at all; and whether the judgement being supplemented will still exist in ten years, when the tool is withdrawn, or updated, or quietly fails.

¹⁴ Ninety people are waiting at a junction four hours from the hospital. The tool should almost certainly be used. That is not the same as saying the matter is settled.

Part A Multiple choice

Circle the letter next to the best answer.

1. The writer argues in paragraph 3 that most debate about medical AI is flawed because it
 - A overestimates how accurate these systems can become
 - B assumes a setting in which an expert is available to be replaced
 - C ignores the cost of the equipment involved
 - D treats patients as statistics rather than individuals
2. According to paragraph 6, a false negative is especially harmful because
 - A it is harder to detect than a false positive
 - B it produces reassurance that changes the patient's future behaviour
 - C it occurs more often than any other kind of error
 - D it cannot be corrected once recorded
3. When the writer says responsibility has been "dissolved" (paragraph 9), he means that
 - A the fault lies with the clinical officer alone
 - B no individual can be held answerable for the error
 - C the error is too small to require investigation
 - D the ministry has formally accepted liability
4. The final two paragraphs indicate that the writer's overall position is that the tool
 - A should be withdrawn until accuracy improves
 - B should be used, but should not be treated as settling the question
 - C is acceptable only in well-staffed hospitals
 - D poses risks that outweigh its benefits

Part B Written response

Answer in two or three sentences. Refer to the passage.

5. Explain the reframing the writer performs in paragraph 3, and why he immediately warns that it is where "careful thinking tends to stop".
6. In your own words, explain why a model may perform worse in one region than another, and what form that failure takes.
7. Explain the writer's claim in paragraph 10 that accountability matters "beyond fairness to the individual patient".

8. Paragraph 11 describes deferral as rational. Explain why the writer calls it rational and why he nonetheless treats it as a problem.

9. The passage ends with a sentence that separates two things. Explain what is being separated and why the writer chooses to close on that distinction rather than on a recommendation.

Part C Think deeper

There is no single right answer here. Take a position and support it with evidence.

10. The writer argues that withholding an imperfect tool means "choosing the worse outcome in order to keep the cleaner conscience". Evaluate that claim. Is it fair to those who object?

11. Who should decide the threshold at which a screening tool is considered good enough to deploy — and should the people who bear the cost of errors have a vote? Argue your position.

12. The passage identifies a long-term risk that skills will quietly disappear. Does this risk justify limiting a tool that saves sight now? Set out the strongest opposing case before answering.

13. Suppose the tool were required by law to produce an explanation for every result. Predict two consequences, one for patients and one for the system, and assess whether the requirement is worth its costs.

Answer key

For parents and teachers

Part A

1. B — assumes a setting where an expert exists to replace | 2. B — reassurance changes behaviour | 3. B — nobody can be held answerable | 4. B — use it, but the question is not closed

Part B

5. The comparison shifts from machine-versus-specialist to machine-versus-nothing, which makes deployment look obviously right. The writer warns against stopping there because several objections — accuracy drift, accountability, skill loss — are untouched by the reframing.

6. Performance depends on how closely patients resemble the training data; differences in equipment, lighting and disease patterns degrade it. The failure is undramatic, appearing as a raised rate of false negatives rather than visible malfunction.

7. Accountability is the mechanism by which errors produce review and change. Without a traceable fault there is no correction, so the system stops improving — a harm to future patients, not only the present one.

8. Rational because the machine is usually right and the queue is long, so checking costs time the officer does not have. A problem because unexercised judgement decays, so a tool meant to supplement expertise ends up replacing it.

9. It separates the practical decision — use the tool — from the claim that the ethical question is answered. He closes there because his argument is about how deployments are justified, not whether this one should proceed.

Part C

These answers vary. Mark the reasoning, not the conclusion.

10. Strong answers notice the phrase is rhetorically loaded, since it characterises opponents' motives rather than their reasons, and decide whether the substance survives that. Agreement or disagreement both creditable.

11. Look for a decision procedure, not just a decision-maker. The best answers acknowledge the practical difficulty of consulting a population before deployment and propose something workable anyway.

12. The opposing case must be stated at full strength — nine people keep their sight now, and a future risk is speculative. Mark the quality of that statement as heavily as the conclusion.

13. Two consequences, correctly split between patient and system, plus a cost assessment. Reward students who recognise that an explanation requirement may slow or block deployment entirely, which has its own victims.

EyeQBoost builds thinking skills in young learners through cognitive training, reading and gamified online classes. More free reading sets at eyeqboost.com/resources.