

CLINICAL THINKING

UNDER PRESSURE

THE

RADIAL

PLAYBOOK



DR JAMES WALLACE



Copyright © 2026 Dr. James Wallace

Published in the United Kingdom.

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, except by a reviewer who may quote brief passages in a review to be printed in a magazine or newspaper. The contents of this book may not be used to train large language models or other artificial intelligence products without written permission from the copyright holder.

All trademarks, service marks, trade names, product names, and logos appearing in this publication are the property of their respective owners.

Important note

This playbook is an educational resource. It exists to support clinical reasoning, reflection, supervision and teaching, and it does not replace local policies or pathways, current clinical guidelines, senior advice, specialist input or professional judgement.

Clinical decisions must be made in context. Use the most current local guidance, seek help when you need it, and escalate whenever risk is visible or uncertainty remains.

RADIAL has not been validated as a diagnostic instrument or a clinical prediction rule, and this resource should never be used as a sole basis for clinical decision-making. It is an author-developed thinking framework, drawing together established principles from acute care, clinical reasoning, medical education, human factors and patient safety.

About the clinical cases

The cases in this playbook are composite teaching examples. They reflect patterns that recur across acute care, but every detail has been altered, combined or generalised, and none describes an identifiable patient, member of staff or organisation.

Each case exists to make the decision process visible: what created the risk, what didn't fit, what needed to happen, and what changed once the patient was reassessed. None is offered as a template for a single correct investigation or treatment pathway.

How to use this playbook

This isn't a book to read once and shelve. Start with Part I to understand why capable clinicians think differently under pressure. Part II will help you recognise the pressure pattern that most often distorts your own decisions. Part III sets out the RADIAL loop itself, and Part IV applies it to realistic clinical and operational cases. Parts V and VI turn the framework into a repeatable practice — for assessment, escalation, handover, supervision and reflection.

Don't try to memorise every paragraph. Learn the six questions and notice which one you tend to skip.

Contents

Important note

Contents

Introduction — decisions happen before certainty

PART 1 - The decision-making layer beneath medicine

Chapter 1 — Knowledge is necessary, but not sufficient

Chapter 2 — Fast thinking, slow thinking and the clinical checkpoint

Chapter 3 — Bandwidth, cognitive load and the pressured mind

PART II - Four clinical pressure patterns

Chapter 4 — The Overloaded Processor

Chapter 5 — The Premature Closer

Chapter 6 — The Hesitator

Chapter 7 — The Intuitive Sprinter

PART III - The RADIAL Clinical Thinking Loop

Chapter 8 — A loop, not an algorithm

Chapter 9 — R: Recognise instability

Chapter 10 — A: ABCDE prioritisation

Chapter 11 — D: Define worst-case risk

Chapter 12 — I: Identify contradictions

Chapter 13 — A: Act and escalate

Chapter 14 — L: Loop and reassess

PART IV - Clinical cases under pressure

Chapter 15 — The breathless older patient

Chapter 16 — The patient who appears intoxicated

Chapter 17 — Two diagnoses in the same patient

Chapter 18 — The story that explains too much

Chapter 19 — The quiet child

Chapter 20 — The older trauma patient who “just fell”

Chapter 21 — The patient between teams

Chapter 22 — The task that disappeared at handover

Chapter 23 — The abnormal result the computer did not emphasise

PART V - Using RADIAL in real clinical work

Chapter 24 — The ten-second and one-minute versions

Chapter 25 — Escalation as a decision threshold

Chapter 26 — Handover, ownership and closed loops

Chapter 27 — Senior visibility and verification

Chapter 28 — Results review, discharge and the senior sweep

Chapter 29 — RADIAL as team language

PART VI - Building a personal clinical thinking system

Chapter 30 — Your danger phrases and corrective questions

Chapter 31 — Deliberate practice with real cases

Chapter 32 — Confidence without false certainty

RADIAL quick reference

The four patterns at a glance

Evidence base and limitations

References

Next step — from framework to practice

About the author

Introduction — decisions happen before certainty

Emergency medicine is full of moments when the answer isn't yet clear, but the patient still needs you to move. That tension is what this playbook is built around.

You may not have the full history. The observations may be borderline rather than catastrophic. The first test may be reassuring while the patient remains wrong. A relative may add one sentence that changes the whole case. A nursing colleague may say, "I am not happy with them," before the numbers offer a neat explanation. The department may be crowded, the handover incomplete, and several people may be waiting on your decision at the same time.

In those moments, the question is not simply what the diagnosis is. It is whether this patient is safe now, what the risk is, and what needs to happen next.

Diagnosis matters. Knowledge matters. Guidelines matter. But none of them removes the need to think while information is incomplete and the situation is still moving.

Clinical thinking becomes visible at the point of commitment: deciding to wait or act, investigate or discharge, reassure or challenge, accept a label or reopen it, manage independently or escalate. It is visible in how you judge whether a patient is stable, compensating or beginning to deteriorate — and visible again when the patient doesn't respond the way you expected.

The diagnostic process is not a solitary act inside one clinician's head. It is a sequence running through patients, relatives, nurses, paramedics, doctors, advanced practitioners, specialists, results, records, handovers and the system around them. The National Academies described diagnosis as a collaborative process that unfolds over time, rather than a single act of naming disease [1].

That is why a good framework cannot only help generate a differential. It must help manage risk, direct attention, communicate uncertainty, and bring the clinician back to the case when new information appears. RADIAL is designed for that layer of practice.

R — Recognise instability: stable, unstable or compensating?

A — ABCDE prioritisation: what kills first?

D — Define worst-case risk: what cannot be missed?


I — Identify contradictions: what does not fit?

A — Act and escalate: what needs action now?

L — Loop and reassess: what has changed?

RADIAL does not tell you the diagnosis. It tells you what your thinking needs to do next.

The framework is deliberately simple — not because acute care is simple, but because a framework that cannot survive fatigue, noise, interruption and uncertainty will not be used when it is needed most. The aim is not to slow you down; it is to make you safer at speed. The aim is not to remove uncertainty; it is to help you act while uncertainty remains. And the aim is not to manufacture confidence in the first diagnosis; it is to keep the patient safe while that diagnosis evolves.



PART 1 - The decision-making layer beneath medicine

Chapter 1 — Knowledge is necessary, but not sufficient

When a clinical decision goes wrong, the first explanation offered is usually knowledge. Did the clinician know enough? Did they remember the relevant disease, follow the guideline, order the right test?

Sometimes that is exactly the problem. Medicine cannot be practised safely without knowledge, technical competence and familiarity with current guidance. But many of the hardest decisions arise when the clinician already knows several plausible diagnoses and still struggles to decide what matters right now.

The patient does not present as a textbook case. The observations are abnormal but not dramatically so. The story is familiar except for one discordant feature. The likely diagnosis is benign, but the dangerous alternative would change the whole pathway. The patient has already been reviewed, and yet no one clearly owns the next decision. A result has come back but has not been read against the patient's trajectory.

None of this is a knowledge deficit. It is the decision-making layer beneath medicine — the layer where a clinician who already knows that pulmonary embolism, sepsis, acute coronary syndrome, heart failure and pneumonia can all cause breathlessness still has to decide which risk deserves attention first, which data are decisive, which contradictions matter, what can safely wait, and what needs to happen before the picture is complete. It is the layer where a clinician who knows a quiet vomiting child can conceal serious pathology still has to notice when a familiar label has become too reassuring, and where a clinician who knows an anticoagulated older adult can have occult injury after a fall still has to act on a pattern of modest abnormalities before any single number looks dramatic.

The National Academies' work on diagnostic safety put this well: diagnostic error emerges from the interaction of cognition, communication, teamwork, workload, technology and organisational systems [1]. That reframes the question. Instead of asking who failed to know the answer, it asks what made the safe answer hard to reach, hard to communicate, or hard to act on. This playbook takes that wider view — clinical reasoning does not happen apart from the environment it happens in.

Five pressures that distort otherwise good decisions

The first is information without prioritisation. Modern clinical environments generate more data than any one clinician can weigh equally — observations, results, messages,

referrals and alerts all competing for attention. Thoroughness is valuable; indiscriminate attention is not the same thing as safety.

The second is familiarity without challenge. A recognisable presentation allows fast, often correct pattern recognition. The risk is that familiarity quietly suppresses the search for evidence that would disconfirm it.

The third is caution without a decision threshold. Waiting can be entirely appropriate, or it can be an unexamined default. The real question is never whether uncertainty remains — it almost always does — but whether enough risk is already visible to justify acting anyway.

The fourth is action without reassessment. A correct first move can still be followed by an unsafe second phase if no one checks how the patient responded. Treatment, new results and changing physiology should all update the working model.

The fifth is handover without ownership. Information can pass between teams while responsibility stays vague. A patient can be “known to” three services and still have no one making the next decision. Being seen is not the same as being safe.

None of these are character flaws. They are predictable vulnerabilities in complex work, and a clinical thinking framework does not make clinicians perfect — it simply gives good clinicians a way to catch a predictable, pressure-related failure before it becomes harm. The clinician’s job is not always to know the final answer immediately. It is to keep the patient safe while moving towards it.

Evidence note

The diagnostic safety literature supports a systems view of diagnosis rather than reducing error to individual knowledge or bias alone [1]. Research on clinical reasoning also warns against the simple claim that intuitive thinking is dangerous and analytical thinking is always safer — expertise depends on both rapid pattern recognition and the ability to shift into slower, deliberate analysis when a case is difficult, contradictory or high risk [2].

Chapter 2 — Fast thinking, slow thinking and the clinical checkpoint

Clinicians reason in more than one mode. Sometimes a diagnosis arrives quickly — the pattern is familiar, the features cohere, and the next step feels obvious. At other times the clinician compares alternatives, tests hypotheses, hunts for contradictions and consciously weighs probability against consequence.

These are often described as fast, non-analytical reasoning and slower, analytical reasoning. The distinction is useful but easily misused. Fast thinking is not inherently careless — experienced clinicians depend on pattern recognition to spot the unwell patient and act before every detail is available. Slow thinking is not inherently accurate either; a clinician can analyse a case at length and still start from the wrong assumption, overweight an irrelevant detail, or simply delay a necessary action.

So, the safer question is not “am I thinking fast or slowly?” It is “does this decision need a checkpoint?”

A checkpoint is justified when the patient is physiologically vulnerable; when the likely diagnosis is benign but the plausible alternative is dangerous; when one feature does not fit the current explanation; when the patient is not responding as expected; when the decision involves discharge, transfer or reduced monitoring; when you feel unusually certain unusually early, or unable to act despite visible risk; when the case is being passed between people or teams; or when fatigue, interruption or workload has quietly reduced your confidence in the information in front of you.

Deliberate reflection has been shown to improve diagnostic accuracy in complex or unusual cases specifically, while offering little advantage in straightforward, routine ones [3]. Structured reflection can also counter the specific biases created by recent experience or by a salient but misleading feature [4,5]. That supports a selective approach: don’t force every case into prolonged analysis but build in a brief interruption point for the cases where being wrong is costly or the pattern isn’t secure.

RADIAL is designed as that interruption point. It doesn’t ask you to abandon your first impression. It asks whether the patient is safe, whether the physiology is secure, whether a dangerous risk remains, whether anything contradicts the story, what action is required, and whether the situation has changed — fast thinking, with fast checks.

The problem with “debiasing” as a slogan

Clinical education often teaches lists of biases — anchoring, availability, confirmation bias, search satisficing, framing, premature closure. Knowing the labels can raise awareness, but awareness does not reliably change behaviour at the bedside, and trials of generic cognitive-forcing education have produced mixed or disappointing results [7,8].

So RADIAL is not sold here as a magic debiasing checklist. Its value is more modest, and more useful: it creates a shared set of questions, makes risk and uncertainty easier to communicate, places reassessment inside the decision rather than after it, supports brief structured reflection at high-risk moments, and integrates into handover, supervision and escalation. A framework survives practice when it is attached to real workflow.

Evidence note

Studies of deliberate reflection support its use in selected difficult cases, but the evidence for broad cognitive-forcing training is inconsistent [3–8]. That is why RADIAL is presented as a practical prompt and shared language, not as a validated intervention guaranteed to prevent error.

Chapter 3 — Bandwidth, cognitive load and the pressured mind

A clinician can have the knowledge required to solve a case and still be unable to use it well in the moment. The mind has limited working capacity, and cognitive load theory distinguishes between the inherent complexity of a task and the extra load created by how information is presented or the environment it is processed in [6]. In acute care, these loads combine.

A single patient may generate competing diagnoses, several observations, multiple results, medication risks, safeguarding concerns, family information and an uncertain trajectory — while the clinician is simultaneously interrupted by calls, requests, handovers, bed pressures, documentation, staffing gaps and other patients whose needs are just as urgent. The result is not simply feeling busy. It is a loss of bandwidth.

When bandwidth narrows, priorities blur, unfinished tasks are easier to lose, and the clinician may lean too heavily on the first coherent story. New information gets heard but not integrated. The same possibilities are revisited without progress. Escalation is delayed because the case hasn't been mentally organised, and reassessment becomes opportunistic rather than planned.

This is why senior clinical maturity is not just knowing more. It is knowing when to rely on a colleague's summary and when to verify the patient personally; when to keep moving and when to protect a minute of uninterrupted thought; when to write a task down rather than trust memory; and when workload has narrowed the safety margin enough to need help.

Cognitive load is not solved by trying harder

When the mind is full, effort alone can increase strain without improving prioritisation. What tends to help instead is reducing the active problem to the top two or three risks, separating immediate threats from issues that can wait, saying or writing the working model aloud, assigning named owners to actions, setting a defined reassessment point, moving to a quieter space for a high-consequence decision where that is possible, asking a colleague to cross-check the case, and returning to physiology when the differential becomes unmanageable. This is not simplification through carelessness — it is simplification through structure.

Fatigue and interruption: use precise language

Fatigue is associated with poorer clinical reasoning in some studies, particularly after prolonged shifts, though the size and consistency of the effect vary across settings [9]. Experimental work on interruptions has not shown a uniform reduction in vignette-based diagnostic accuracy either [10].

So, the practical conclusion has to stay measured. It is too strong to claim that every interruption causes diagnostic error. It is equally unrealistic to pretend that fatigue, divided attention and task-switching carry no operational consequence at all. The danger is rarely a dramatic collapse in intelligence — more often it is a small reduction in reserve: one result not reviewed, one task not handed over, one contradiction not revisited, one patient not reassessed, one concern not properly heard, one plan left without an owner. High-reliability practice is built around making those small failures easier to catch.

A bandwidth check

Before a high-consequence decision, it is worth asking: do I have the information I think I have? Have I heard and integrated the concern that was just raised? Is there anything I need to verify myself? What task would be lost if I don't write it down now? And who else should share this decision? None of this removes workload — it protects the decision from it.

PART II - Four clinical pressure patterns

The four patterns in this playbook are not diagnoses, personality types or fixed identities. They are recognisable pressure responses. You may see yourself strongly in one or move between all four across the same shift. Experience does not remove them — senior clinicians can overload, close early, hesitate and over-trust instinct just as readily as anyone else. The content of the pattern may change with expertise, but the underlying vulnerability stays human.

Naming a pattern is not about labelling the clinician. It is about identifying the missing counterweight.

Chapter 4 — The Overloaded Processor

The Overloaded Processor is thorough, detail-aware and risk-sensitive. They often see more than other people do — comorbidities, medication interactions, atypical presentations, subtle abnormalities, the ways a case could deteriorate. In complex patients this is a genuine strength.

The risk appears when every possibility is given equal weight. Pneumonia, pulmonary embolism, heart failure, sepsis, acute coronary syndrome, medication effect, delirium, dehydration, anaemia, COPD, frailty, anxiety — all may be plausible, but they are not all equally important at the same moment. When the mind becomes crowded this way, the Overloaded Processor may revisit the same differential repeatedly, delay action while trying to complete the whole picture, struggle to say what the core concern actually is, order investigations without a clear hierarchy, and lose track of tasks because too many stay active at once.

The missing counterweight: prioritisation

The Overloaded Processor does not need to care less. They need to sort information by consequence and time.

The Top Three Rule helps: name the three active risks that matter most right now — not every theoretical diagnosis, not every issue that might matter later, but the three possibilities or complications most likely to change immediate management. From there, ask what must not be missed now, what can safely wait, and what information would change the next action.

It also helps to externalise the working model — say it or write it down: “My leading diagnosis is pneumonia. My immediate risks are sepsis, respiratory failure and cardiac

pathology. The confusion and borderline blood pressure are why I am not treating this as a simple chest infection.” A model becomes easier to test once it is visible. And when the differential keeps expanding, it helps to return to physiology and ask the simplest question of all — stable, unstable, or compensating? Physiology decides urgency before diagnosis decides destination.

Danger phrase: “There are too many possibilities.”

Corrective question: “Which three risks change what I do now?”

Clinical motto: Not everything is equally urgent.

Chapter 5 — The Premature Closer

The Premature Closer is decisive, efficient and action oriented. Acute care needs clinicians who can form a plan and move a patient forward, and decisiveness is not itself a flaw. The risk appears when the first coherent diagnosis becomes fixed before the whole risk has been tested.

The presentation feels familiar. The patient appears stable. The first result isn't alarming. The symptoms are improving. A label arrives early and starts organising everything that follows — this is diagnostic momentum. Once a case has a name, new information tends to be read in a way that protects the name rather than tests it. The Premature Closer may discount data that don't fit, stop generating meaningful alternatives, treat “currently stable” as “safe to close,” accept a referral label without rebuilding the case themselves, explain away a contradiction rather than look into it, or discharge a patient after improvement without checking the trajectory that led there.

The missing counterweight: interruption

The Premature Closer does not need to slow down across the board. They need one short checkpoint before commitment — three questions, asked deliberately: What does not fit this diagnosis? What plausible diagnosis would make this plan unsafe? And is the patient becoming safer over time, or have I simply accepted the initial label?

Danger phrase: “This is probably just...”

Corrective question: “What would make this explanation incomplete?”

Clinical motto: Pause before it becomes fixed.

Chapter 6 — The Hesitator

The Hesitator is cautious, reflective and alert to consequence. Caution protects patients from arrogance and reckless certainty, and it encourages clinicians to seek advice and recognise their own limits. The risk appears when caution turns into passive waiting while the patient already shows enough risk to justify action.

The Hesitator may want one more observation set, one more blood result, one more scan, one more opinion, one more piece of certainty before committing — and sometimes that is exactly right. But it can also delay analgesia, monitoring, escalation, treatment or a move to a safer environment. The underlying problem is not uncertainty itself, which is normal. It is the absence of a decision threshold.

The missing counterweight: action on visible risk

The Hesitator does not need to become reckless. They need to define what evidence is sufficient to act. Ask what decision has to be made in the next three to five minutes, and separate reversible action from irreversible: closer monitoring, senior review, repeat observations, analgesia, moving to a higher-acuity area, targeted investigation or starting a relevant pathway are all low-regret moves that can be adjusted later. Waiting is not neutral — it is itself a decision, with its own consequences, and it's worth naming what those are. And escalation does not require a complete diagnosis: “I am concerned because the physiology is drifting and the current explanation does not account for it” is enough to say out loud.

Danger phrase: “I will just wait for...”

Corrective question: “Is there enough visible risk to make the patient safer now?”

Clinical motto: Visible risk justifies action.

Chapter 7 — The Intuitive Sprinter

The Intuitive Sprinter recognises patterns quickly and is comfortable acting on incomplete information. This can be a genuinely advanced clinical skill — experienced clinicians often notice the quiet child, the grey patient, the change in behaviour, or the story that doesn't match the room, well before they can fully articulate why. The risk is not intuition itself. It is letting a useful first impression harden into an untested conclusion.

The Intuitive Sprinter may act before checking the full physiology, struggle to explain their reasoning to colleagues, underweight a slower or less familiar diagnosis, skip deliberate reassessment because the initial response felt convincing, or confuse confidence with calibration.

The missing counterweight: testing

The Intuitive Sprinter does not need to suppress instinct — they need to make it visible and testable. Name the first impression out loud: “My first impression is sepsis.” “This child looks high risk.” “This feels like heart failure, but I need to account for the fever and pleuritic pain.” Then ask what finding would prove it wrong and try to explain the reasoning to someone else — if the logic can't be communicated, the decision probably isn't yet formed enough for the team to safely act on. Finally, build in the second look: when will you check whether the instinct was right?

Danger phrase: “I know what this is.”

Corrective question: “What would prove my instinct wrong?”

Clinical motto: Trust it, then test it.

PART III - The RADIAL Clinical Thinking Loop

Chapter 8 — A loop, not an algorithm

RADIAL is a structured clinical thinking loop for recognising risk, prioritising decisions and acting safely under pressure. It is not a protocol, not a clinical prediction rule, and not a substitute for disease-specific pathways, local policy or senior judgement. It won't tell you which antibiotic to prescribe, which scan to request, or which specialty should accept the patient — it sits one layer beneath those decisions. It is the thinking layer.

The six questions are: recognise instability — stable, unstable or compensating? ABCDE prioritisation — what kills first? Define worst-case risk — what cannot be missed? Identify contradictions — what does not fit? Act and escalate — what needs action now? Loop and reassess — what has changed?

The order matters, but it isn't rigid. In a visibly deteriorating patient, recognition, ABCDE, action and escalation happen almost simultaneously. In a stable but diagnostically uncertain patient, the most valuable steps may be defining the worst-case risk and identifying contradictions. After treatment or new results, the loop simply begins again — which is why it should be used as a set of cognitive priorities, not a form to fill in.

RADIAL in one breath: Safe now? First threat? Worst case? What does not fit? What needs doing? What has changed?

Why a loop?

Clinical decisions are often talked about as endpoints — diagnose, admit, discharge, refer. They are sequences. A diagnosis starts as a hypothesis. A plan carries assumptions. Treatment generates new information, and the patient either follows the expected trajectory or doesn't. A safe decision is therefore not only a plausible choice at one moment — it is a choice with a built-in mechanism for detecting when it has become wrong. That final letter, Loop, is what separates RADIAL from a static checklist.

The RADIAL loop at a glance

The diagram below sets out the loop the way it would appear as a circular diagram — six steps, run continuously, with the reminder at its centre: do not process everything equally.

RADIAL: The Clinical Thinking Loop

A practical framework for thinking clearly under pressure



• *Do not process everything equally.* •

Chapter 9 — R: Recognise instability

Ask: stable, unstable or compensating? This comes first because physiology decides urgency. Don't begin by asking only what the diagnosis is — begin by asking whether this patient is safe right now.

Instability can be obvious — airway compromise, severe respiratory distress, shock, major haemorrhage, reduced consciousness, seizures, peri-arrest physiology. But many dangerous patients aren't dramatically unstable on first assessment. They are compensating: a rising respiratory rate before the oxygen saturation falls, tachycardia with a technically acceptable blood pressure, new confusion in an older adult, a quiet child who has stopped interacting normally, a patient who looks grey, sweaty, exhausted or frightened, pain or distress out of proportion to the apparent diagnosis, or a trajectory that keeps quietly worsening despite reassuring snapshots.

NEWS2 standardises the recording of key physiological parameters and explicitly includes new confusion as a sign of acute illness [12]. It is a useful common language and escalation aid, but it is not a substitute for clinical judgement — a score describes measured physiology at a single moment, not reserve, trajectory, appearance, or the weight of a clinician's concern.

Stability is a moment. Safety is a trajectory. A patient can be stable at two o'clock and unsafe an hour later. A patient can sit within acceptable ranges while compensating at the very edge of their reserve. A treatment can improve one number without touching the underlying risk. Recognition therefore needs three views at once: the current physiology, the direction of travel — improving, static, or deteriorating — and how much physiological reserve this patient has left to compensate with. Older adults, frail patients, children, pregnant patients, immunosuppressed patients and anyone on medication that blunts the normal physiological response may not display risk in the textbook way; the point isn't to memorise a list of exceptions, but to avoid false reassurance from a single number or label.

The escalation signal that has no number

A nursing colleague, paramedic, relative or other clinician may simply say, "They are not right." That is not a diagnosis, but it is still data — often reflecting a change in behaviour, appearance or trajectory observed over time. Treat it as a prompt to reassess, not as an inconvenience to be translated into a more acceptable number before it can be taken seriously.

Practice prompt

Before leaving the bedside, ask: what suggests this patient is compensating? What would deterioration look like next? And when, and by whom, will that change be detected?

Chapter 10 — A: ABCDE prioritisation

Ask: what kills first? ABCDE is sometimes taught as a sequence to recite. In practice it is a prioritisation discipline. The Resuscitation Council UK advises a complete initial ABCDE assessment, treating life-threatening problems before moving on, calling for help early, using the whole team, and reassessing regularly to see whether interventions are working [11].

The value of ABCDE isn't that it replaces diagnosis — it stops the diagnostic story from distracting you from physiology. A patient may have pneumonia, but the immediate problem may be oxygenation or shock. A child may have vomiting, but the immediate concern may be circulation, conscious level or glucose. An older patient may have a hip fracture, but the urgent threats may include occult bleeding, hypothermia, head injury or the medical cause of the fall itself. A patient may be seizing after travel, but airway protection, oxygenation, glucose, temperature and electrolyte disturbance take priority over settling on a final aetiological label.

ABCDE is not a ritual

A mechanical run through the letters can itself become a form of premature closure if each one is declared normal without asking whether the patient is safe. Used dynamically, it asks: what is the active physiological problem, what could deteriorate first, what intervention can happen now, what response should that intervention produce, and what finding would demand immediate help?

In a team, ABCDE supports parallel work — one person assessing and communicating while others attach monitoring, obtain access, prepare treatment, repeat observations or call for senior support. What matters is that the team shares the same current model of the problem and the same next priorities; that is where clinical thinking becomes visible, not just in how quickly tasks are performed but in whether the team knows why those tasks matter right now.

Practice prompt

If the diagnosis disappeared from the notes, could you still describe the patient's physiological problem and the immediate priorities?

Chapter 11 — D: Define worst-case risk

Ask: what cannot be missed? The most likely diagnosis and the most dangerous plausible one are not always the same, and emergency and acute care often require holding both models at once — what is most likely, and what would be most harmful to miss.

Worst-case thinking doesn't mean imagining every rare catastrophe. It means identifying the plausible diagnosis, complication or trajectory that would change immediate management. In central chest pain, the likely explanation may be benign, but acute coronary syndrome, aortic pathology, pulmonary embolism or pneumothorax may all need consideration depending on the presentation. In a quiet, vomiting child, gastroenteritis is common, but sepsis, diabetic ketoacidosis, appendicitis, obstruction, intussusception or torsion may change the plan. In an older fall, the visible fracture can dominate the picture while head injury on anticoagulation, occult haemorrhage, hypothermia, rhabdomyolysis or a medical precipitant remain hidden underneath it.

Define, do not merely list

A long differential can look thorough without sharpening the decision. Defining worst-case risk means answering a tighter set of questions: which dangerous possibility is credible in this patient, which feature makes it credible, what would you expect to find if it were present, what action, investigation, monitoring or escalation would change because of it, and when can it be considered sufficiently addressed? That turns a fear into a plan.

Sometimes the worst case isn't a different diagnosis at all — it's a complication or a future state. What could this become in the next hour? What could fail if the patient stays in this environment? What result would demand urgent action? What would you regret not having anticipated? This matters particularly in sepsis-like presentations, evolving respiratory failure, high-risk intoxication, frailty, anticoagulated trauma, and patients waiting for transfer.

Practice prompt

Complete the sentence: "The leading diagnosis is _____, but the risk I am actively protecting against is _____ because _____."

Chapter 12 — I: Identify contradictions

Ask: what does not fit? Contradictions are often where the important information actually lives — no diarrhoea in the child labelled gastroenteritis; exertional breathlessness and peripheral oedema in the patient labelled uncomplicated diabetic ketoacidosis; confusion and borderline blood pressure in the patient labelled a simple COPD exacerbation; focal neurology in the patient assumed to be intoxicated; a prolonged wait with no named decision-maker for a patient who has already been “reviewed.”

A contradiction doesn’t automatically prove the first diagnosis wrong. It may mean the diagnosis is incomplete, a second process is present, the patient is deteriorating, or the plan simply needs a stronger safety net.

Plausible is not the same as complete

One of the most dangerous features of a plausible diagnosis is that it explains enough. Once several findings fit, the mind can quietly stop asking what remains unexplained — especially when the label has arrived ready-made from triage, an ambulance handover, a referring clinician or an earlier review. It helps to deliberately split the case into two columns: what fits the working diagnosis, and what doesn’t — the findings that are unexplained, disproportionate, later than expected, physiologically discordant, or simply inconsistent with the trajectory. The second column shouldn’t be allowed to disappear just because the first one is convincing.

Contradictions can be operational, too

Not every contradiction is diagnostic. The patient is described as safe for discharge, but no one has reviewed the result. The plan says senior review, but no senior is named. The handover says a team will attend, but the team hasn’t accepted ownership. The observations are called stable, but they haven’t been repeated since treatment. The patient is in a lower-acuity area, but their monitoring needs have quietly increased. These can matter every bit as much as a conflicting symptom.

Practice prompt

Before committing to a plan, ask: what fact would be inconvenient if my current explanation were wrong?

Chapter 13 — A: Act and escalate

Ask: what needs action now? Thinking without action does not make a patient safer. You don't always need the final diagnosis before taking the next safe step — that step may be treatment, monitoring, analgesia, repeat observations, targeted investigation, a move to a safer clinical area, senior review, a specialty discussion, or simply clarifying who owns the patient. The action should be proportional to the visible risk and consistent with local guidance.

When a case feels complex, ask what single action would most reduce immediate risk while you keep thinking. That single question guards against two opposite errors: waiting for complete certainty before doing anything and launching an indiscriminate response without a clear risk model.

Escalation is a decision threshold

Escalation is not an admission of incompetence. It is the point at which the consequence of managing alone outweighs the benefit of further independent work. Escalate when physiology is unstable or drifting; when the patient isn't responding as expected; when the worst-case risk remains insufficiently controlled; when the diagnosis and the physiology disagree; when a decision exceeds your authority or experience; when system barriers are delaying safe care; when the patient needs a level of monitoring, treatment or expertise the current area can't provide; or when you've simply reached the limit of your own bandwidth. The GMC requires clinicians to act promptly where patient safety may be seriously compromised and expects leaders to build an environment where concerns can be raised and acted on safely [19].

Structure the escalation

A useful escalation covers four things: what is happening, what you're worried this could become, what has changed or doesn't fit, and what you need from the person you're speaking to. That last part matters — a description without a requested action can leave responsibility just as ambiguous as before. This is not panic. It is structured risk communication.

Ownership is an action

Sometimes the most important intervention isn't a drug or a test. It is naming who owns the next decision, and by when. A patient can remain in a physically safe place while clinically drifting simply because no one has accepted responsibility for the next step — resolving that ambiguity is itself part of patient care.

Practice prompt

If the plan is to wait, complete the sentence: "We are waiting for _____. During that time, _____ owns the patient, _____ will be monitored, and the plan changes if _____."

Chapter 14 — L: Loop and reassess

Ask: what has changed? A clinical decision is rarely finished when the first plan is made. The patient changes. The observations change. Treatment changes the physiology. Results shift the probabilities. A relative supplies the missing piece of history. The nurse reports a new concern. The patient looks better, worse, or simply different. Reassessment isn't administrative repetition — it is active risk detection.

Reassess the model, not only the numbers

A repeat set of observations is useful, but it is not the whole reassessment. Ask whether the patient has improved the way you expected, which findings remain unexplained, whether a new result has changed the worst-case risk, whether the diagnosis still accounts for the physiology, whether the response to treatment has made the original plan safer or less safe, whether the current environment is still appropriate, and whether the patient is now suitable for discharge, transfer or reduced monitoring.

The second look

Many errors are not failures of the first assessment — they are failures to return. A haemoglobin falls. A potassium comes back dangerously high. An ECG evolves. Pain recurs. Confusion persists. The patient who was compensating begins to tire. A plan that was entirely reasonable thirty minutes ago becomes unsafe simply because the case has moved on. A diagnosis should become safer over time. It should not just be accepted at the outset and then protected from challenge.

Make reassessment explicit

Every significant plan should be able to answer: what are we expecting to happen, when will we review, who will review, what change triggers escalation, and what result must be seen before the case is closed? That turns an intention into a closed loop.

Evidence note

The ABCDE approach explicitly includes regular reassessment and evaluation of treatment effect [11]. Structured handover interventions such as the I-PASS bundle have reduced medical errors and preventable adverse events by improving how illness severity, action lists, contingency planning and receiver synthesis are transferred between teams [13]. These findings support a wider principle running through this playbook: plans become safer when actions, ownership, contingencies and reassessment are all made explicit rather than assumed.

PART IV - Clinical cases under pressure

The cases in this section are deliberately imperfect. Each contains a plausible explanation, incomplete information, and a point where thinking can drift. Don't read them for one correct answer — read them for the questions that keep the patient safe while the answer is still forming.

Chapter 15 — The breathless older patient

The case

A 76-year-old man presents with increasing breathlessness and mild confusion. Temperature 38.1°C, respiratory rate 26, oxygen saturation 93% on air, heart rate 112, blood pressure 104/68. His history includes COPD, heart failure and type 2 diabetes. The triage nurse says, "He just does not seem right."

Several explanations arrive at once: pneumonia, COPD exacerbation, heart failure, sepsis, pulmonary embolism, acute coronary syndrome, delirium, dehydration, medication effect. This is the Overloaded Processor's case — every possibility is plausible enough to demand attention.

The trap

The clinician tries to process every diagnosis equally. The list grows, but the decision doesn't get any clearer.

Running RADIAL

He is not peri-arrest, but he is physiologically vulnerable — tachypnoea, tachycardia, borderline blood pressure and new confusion form a pattern. The question isn't whether any single number is catastrophic; it's whether he is compensating with limited reserve. Breathing and circulation need immediate attention: work of breathing, oxygenation, perfusion, conscious level and temperature, assessed and treated in line with clinical findings and local pathways.

The immediate risks aren't the whole differential — they are the processes that would change management now: sepsis, respiratory failure, and significant cardiac or thromboembolic pathology. If this is "just COPD," why is he confused and febrile? If it's "just infection," why is the blood pressure borderline? If it's "just heart failure," what explains the fever? These contradictions don't need ten more diagnoses — they identify which features must stay active in the risk model.

The next actions should reduce immediate danger and clarify the dominant risks, with escalation if physiology deteriorates, the patient needs higher-acuity care, or the working diagnosis stays discordant with the clinical picture. Reassessment then asks whether the respiratory rate has improved, the oxygen requirement is changing, the confusion is resolving, the blood pressure is stable, and whether blood gas, ECG, imaging or laboratory results have shifted the model.

The shift

Without structure: “There are too many possibilities.” With RADIAL: “These are the risks that matter first. This is the next safe action. This is when I will reassess.”

Field lesson

Complexity isn’t reduced by ignoring detail. It’s reduced by ranking detail according to risk and time.

Chapter 16 — The patient who appears intoxicated

The case

A patient arrives with slurred speech, altered behaviour and an initial description that they appear intoxicated. They are awake and haemodynamically stable, and the label feels plausible — but something doesn't fit. The speech disturbance is focal. The history is uncertain. The behaviour doesn't match the amount of alcohol reportedly consumed.

The trap

The first label is socially and clinically familiar, and once “intoxication” enters the case, everything that follows risks being read through it.

Running RADIAL

Current stability doesn't remove neurological risk — airway, breathing, circulation, conscious level, glucose and the pattern of neurological deficit all still need assessing, alongside securing the immediate physiological priorities and identifying any time-critical neurological feature. Asking what cannot be missed here points squarely at acute stroke or another intracranial process as the dangerous plausible diagnosis. What doesn't fit intoxication is a focal deficit, an abrupt onset, asymmetry, a history that doesn't match the examination, or a trajectory that doesn't follow expected intoxication.

The action here is urgent escalation and the relevant local stroke or neurological pathway where indicated — the point isn't to prove the diagnosis at the bedside before acting, it's to avoid losing time while a time-sensitive diagnosis remains plausible. Reassessment then tracks the neurological findings, timing, collateral history and investigation results, updating the working diagnosis as the evidence develops.

The shift

The easy label may be correct. It still has to survive the worst-case question.

Field lesson

Framing begins before you meet the patient. Rebuild the case from physiology, history and examination rather than inheriting the label untested.

Chapter 17 — Two diagnoses in the same patient

The case

A patient presents thirsty, dehydrated, losing weight and acutely breathless. Initial testing shows hyperglycaemia, ketonaemia and metabolic acidosis — diabetic ketoacidosis provides a coherent explanation. But the patient also describes several weeks of exertional breathlessness and has developed lower-limb swelling. The breathlessness seems disproportionate to the dehydration, and the story contains both fluid loss and possible congestion.

The trap

The first diagnosis is real, and because it's real, it becomes harder to look beyond it. This is search satisfaction: finding one important abnormality and unconsciously treating the search as complete.

Running RADIAL

The immediate physiological consequences of the metabolic disturbance and any cardiorespiratory compromise need assessing together, with immediate threats treated while gathering the information needed to understand the breathlessness and circulation. Worst-case thinking asks what dangerous process could coexist with, or have precipitated, the metabolic emergency — significant acute cardiac pathology must stay active as a differential diagnosis, wherever the history and examination support it. What doesn't fit uncomplicated diabetic ketoacidosis is the longer history of exertional symptoms, the peripheral oedema, and a cardiorespiratory pattern that remains unexplained.

Local pathways for the metabolic emergency should run alongside investigation and escalation of the cardiac risk — one diagnosis should never delay treatment of the other. Reassessment tracks physiology, fluid response, ECG findings, cardiac markers and evolving symptoms, since treating one problem can change the risk associated with the other.

The shift

A correct diagnosis can still be an incomplete diagnosis.

Field lesson

The presence of one genuine abnormality doesn't guarantee it explains the whole patient.

Chapter 18 — The story that explains too much

The case

An unwell patient is brought to a resuscitation area after a seizure. The history includes recent foreign travel and diarrhoea, and an infectious or metabolic explanation forms quickly. A clinician pauses — the seizure seems insufficiently explained. The history is revisited with a relative, the blood results are re-examined, and the timeline is reconstructed. Further information reveals heavy alcohol consumption while away, followed by abrupt cessation on return.

The trap

Travel and diarrhoea are vivid, memorable details, and they build a compelling frame before the whole physiology and timeline have been integrated.

Running RADIAL

A seizure with altered physiology needs immediate ABCDE assessment, temperature, glucose and evaluation for ongoing neurological or metabolic threat, with life-threatening problems treated before the cause is refined further. The dangerous possibilities to define include ongoing seizures, severe metabolic disturbance, intracranial pathology, infection, and toxicological or withdrawal syndromes, depending on context. What doesn't fit the initial travel-related story is the seizure pattern, the timing, the missing exposure history, unexplained autonomic or neurological features, or blood results that don't match the first account.

Treatment and escalation should follow the active risks while the history is clarified — collateral information can be decisive when the patient can't give a reliable account themselves. In this case the patient's partner details a sudden cessation in significant alcohol intake whilst abroad, necessitating management of acute withdrawal. As the history sharpens: a diagnosis is safer when it explains the timeline, the physiology and the response to intervention, all three together.

The shift

Context is useful. It can also become a frame that hides the cause.

Field lesson

When a story feels unusually neat, ask which features it explains and which it merely distracts from.

Chapter 19 — The quiet child

The case

A six-year-old presents with abdominal pain, vomiting and reduced oral intake. Temperature 37.9°C, heart rate 142, respiratory rate 24, capillary refill time three seconds. There has been no diarrhoea. The child is quiet, lying still, and reluctant to move. The parent says, “I think it is a stomach bug.”

The trap

Vomiting is common, gastroenteritis is familiar, and the parent has already supplied a diagnosis. The Intuitive Sprinter recognises the pattern and moves quickly towards reassurance.

Running RADIAL

Tachycardia, a prolonged capillary refill and altered behaviour are concerning even without dramatic hypotension or hypoxia — a quiet child can be more concerning than a distressed, interactive one. Airway, work of breathing, circulation, conscious level, glucose, hydration and exposure all need assessing to decide whether the child is safe to remain in the current pathway. Depending on history and examination, the worst-case list includes serious infection, diabetic ketoacidosis, appendicitis, obstruction, intussusception, torsion, or safeguarding concerns.

What doesn't fit a simple stomach bug: no diarrhoea, reluctance to move, localised pain, persistent tachycardia, prolonged refill, bilious vomiting, abnormal behaviour, pain out of proportion. The point isn't to investigate every child for every diagnosis — it's to stop a familiar label overriding physiology and behaviour. Appropriate assessment, analgesia, monitoring and senior or paediatric review should follow according to risk and local guidance, keeping the pathway open while surgical or metabolic disease remains plausible. On reassessment: does the heart rate improve, does the child become more interactive, does the pain localise, can they move normally, and do any new symptoms appear? After referral to an inpatient surgical team, imaging shows a perforated appendix and the patient has successful surgical management.

The shift

A common diagnosis should become safer with time, not merely be accepted at the beginning.

Field lesson

Don't ignore intuition, and don't blindly trust it either. Test it against physiology, behaviour and trajectory.

Chapter 20 — The older trauma patient who “just fell”

The case

An 82-year-old woman presents after a fall at home. She has left hip pain, cannot weight-bear, and has a shortened, externally rotated leg. Heart rate 108, blood pressure 106/64, respiratory rate 22, oxygen saturation 96% on air, temperature 36.1°C. She takes an anticoagulant and was on the floor for several hours. The handover reads: “Likely fractured neck of femur.”

The trap

The visible injury becomes the whole case. Because the diagnosis looks obvious, the clinician waits for imaging and routine results before acting on the wider risk.

Running RADIAL

She isn’t peri-arrest, but tachycardia, borderline blood pressure, a raised respiratory rate, anticoagulation and a prolonged lie together signal physiological vulnerability. Airway, breathing, circulation, conscious level, glucose where indicated, temperature and the full injury pattern all need assessing — exposing enough to examine the back, flanks, chest wall and pelvis while maintaining dignity. Worst-case thinking should include occult bleeding, head injury, additional fractures, hypothermia, complications of a long lie, and a medical precipitant for the fall itself.

What doesn’t fit “just a hip fracture”: an unclear mechanism, amnesia, confusion, disproportionate tachycardia, bruising away from the obvious injury, chest wall tenderness, abdominal or flank findings, or a physiological pattern pain alone doesn’t explain. Early analgesia and appropriate monitoring, investigation and escalation should follow local trauma, anticoagulation and frailty pathways — not every result needs to be back before the patient is made safer. On reassessment: has pain improved, has the heart rate settled, is blood pressure stable, has confusion appeared, and has examination or investigation revealed a second injury or medical cause? This patient had a hidden haematoma to their posterior chest wall, hiding a radiological flail chest requiring admission and nerve blocks to control pain and prevent deterioration.

The shift

The obvious injury is not permission to stop examining the patient.

Field lesson

In older trauma, modest abnormalities and posterior injuries can carry substantial risk. The first diagnosis may be visible; the dangerous one may be behind the patient.

Chapter 21 — The patient between teams

The case

A patient is referred to one specialty, reviewed, and passed to another. Both assessments are defensible on their own, but neither team makes the final plan. Hours later the patient remains in hospital — physically safe, but without active direction, clear ownership, or a defined reassessment point.

The trap

Everyone has contributed information. No one owns the decision. Responsibility diffuses because each team is waiting for another to supply definitive certainty.

Running RADIAL

Clinical stability still needs confirming — waiting in hospital doesn't automatically make a patient safe — and immediate physiological needs and monitoring must stay appropriate throughout. The worst-case question here is what harm delay itself could cause: deterioration, missed treatment, prolonged fasting, medication omission, psychological distress, loss of function, or an unsafe discharge process, depending on the case.

What doesn't fit: the patient is described as “awaiting a decision,” but no decision-maker is named; the plan says, “specialty review,” but the specialty hasn't accepted the task; the patient has been seen, but no one can say what happens next. The action here is to name an owner, define the decision required, set a time, and escalate across professional or organisational boundaries if the ambiguity persists. Reassessment asks whether the decision has now occurred, whether the patient's condition or needs have changed during the wait, and whether the current plan is still appropriate.

The shift

The case is not safe just because it has been seen.

Field lesson

A handover without ownership is not a safe endpoint.

Chapter 22 — The task that disappeared at handover

The case

During a senior sweep of a clinical area, a patient is found waiting for review of completed blood tests and a final discussion. The task had been held by a senior clinician but was never completed, written down or handed over. The patient has waited several hours for a decision that takes only minutes once it's rediscovered.

The trap

The clinical reasoning here was never the hard part — task completion failed under workload. It's worth naming plainly: patient safety isn't only threatened by the hardest cases. Ordinary unfinished work can cause avoidable delay, frustration and lost trust just as easily.

Running RADIAL

The first step is confirming the patient has stayed clinically stable during the wait and addressing any new physiological concern before assuming this is still a simple results review. Worst-case thinking asks what result, symptom or change would make the delay significant, and what should be reviewed before the patient can safely leave. What doesn't fit: the patient is "ready for discharge," but the required senior review hasn't happened; the task is assumed to exist but isn't on any visible list; the plan is complete in the clinician's head but incomplete in the system.

The action is to complete the decision, communicate with the patient, and make any follow-up explicit — and for future tasks, to use a visible list, named ownership and closed-loop handover. Before leaving an area or a shift, a brief sweep for patients seen but not decided, results not closed, discharge blockers and plans that exist only in memory serves as the reassessment step here.

The shift

Even senior clinicians lose tasks. Safe systems assume memory will sometimes fail.

Field lesson

Productivity isn't only picking up the next patient. It's reducing hidden risk before handover.

Chapter 23 — The abnormal result the computer did not emphasise

The case

A clinician reviews a set of laboratory results. The electronic system uses colour, arrows or bold text to highlight values outside the reference range, and attention is naturally drawn to the most visually salient abnormalities. One clinically dangerous result is less prominent — only modestly outside range, presented in an unfamiliar format, or significant because of the patient's context rather than the laboratory flag.

The trap

The interface decides what attracts attention. The clinician reviews what is highlighted rather than what could change the plan.

Running RADIAL

The starting point is the patient, not the screen: are these results consistent with the current physiology and appearance? From there, ask which result relates to an immediate airway, breathing, circulation, disability or exposure risk, and — before scanning the screen at all — what dangerous result you are specifically looking for in this patient. That might be a severe electrolyte disturbance, significant anaemia, hypoglycaemia, acidosis, renal dysfunction, myocardial injury, or another context-dependent finding.

What doesn't fit is any result that doesn't match the current diagnosis, the treatment response or the expected trajectory, and any "normal" result that is inappropriate for this patient or changing rapidly. The response should follow clinical significance rather than the colour of the interface, with urgent abnormalities escalated and critical results communicated and acknowledged. Reassessment means returning to the patient after acting, confirming the result has been addressed, checking the response, and closing the task.

Three questions for results review: What result would change the plan? What result would be dangerous to miss? What result does not fit the current diagnosis?

Field lesson

Digital salience is not clinical significance.

PART V - Using RADIAL in real clinical work

Chapter 24 — The ten-second and one-minute versions

A framework only matters if it can survive a busy department. RADIAL shouldn't need a worksheet at the bedside — its most important form is six questions held in sequence.

The ten-second version

- Stable, unstable or compensating?
- What kills first?
- What cannot be missed?
- What does not fit?
- What needs action now?
- What has changed?

That is enough to redirect attention.

The one-minute version

With slightly more time, RADIAL structures a verbal summary: “This patient is currently compensating rather than clearly stable. Breathing and circulation are the active physiological issues. The leading diagnosis is pneumonia, but the immediate risks are sepsis, respiratory failure and cardiac pathology. Confusion and borderline blood pressure do not fit a simple chest infection. We have started monitoring and treatment, requested targeted investigations, and need senior reassessment after the blood gas and repeat observations.”

That summary is useful precisely because it carries more than a label — the current physiological state, the active priorities, the likely diagnosis, the dangerous alternatives, the contradictions, the actions already taken, and the next point of reassessment, all in one breath.

RADIAL in the notes

Documentation doesn't need to grow longer to be more useful — it needs to make the risk model and the next steps clear. For example: “Physiologically vulnerable older fall: HR 108, BP 106/64, RR 22, anticoagulated, prolonged lie. Obvious left hip injury but risk not limited to fracture. ABCDE completed. Concerns include occult bleeding, head injury and complications of prolonged immobility. Analgesia, monitoring, relevant imaging and blood tests arranged. Senior review requested. Reassess after analgesia, repeat observations and initial results.” A note like this tells the next reader what matters and what remains open — more useful than a longer, unstructured list ever is.

Chapter 25 — Escalation as a decision threshold

Escalation is often taught purely as communication, but before it's communication, it's a decision — recognising that the current level of support, expertise, authority, monitoring or treatment is no longer sufficient for the visible risk.

That threshold can be hard to reach for two reasons. The clinician may fear appearing uncertain, inexperienced or alarmist, and the senior receiving the call may be difficult to find, inconsistent, or dismissive. Both are patient-safety problems in their own right.

Make the threshold explicit

Escalate when physiology is abnormal or worsening; when a high-consequence diagnosis remains plausible; when a contradiction the current plan can't explain has appeared; when treatment isn't working; when the decision sits beyond your competence or authority; when an operational barrier is delaying something necessary; when a patient's needs exceed the current clinical area; or when workload or fatigue has reduced your own ability to manage safely alone.

The escalation statement

A concise escalation can follow a simple structure — current state: “The patient is compensating, with a rising respiratory rate and borderline blood pressure.” Risk: “I am concerned about early sepsis and respiratory failure.” Contradiction: “The confusion is not explained by a simple COPD exacerbation.” Request: “I need you to review the patient now and help decide whether they require resuscitation-area care.” The request matters most of all — a description without a requested action can leave responsibility just as ambiguous as before.

Receiving escalation safely

Senior clinicians carry two tasks when concern is raised: assessing the patient and the decision and protecting the team's future willingness to raise concerns at all. A junior clinician or a newly trained triage nurse may escalate frequently while they're still calibrating risk, and that is expected. Belittle them for it, and the immediate interaction may go quieter — but the department becomes less safe.

The GMC expects medical leaders to build an environment where errors and concerns can be discussed safely and acted on promptly [19], and research on speaking-up climates shows clinicians will withhold safety concerns when they fear negative interpersonal consequences or believe nothing will change [15,16]. Approachability, in other words, isn't a soft leadership extra — it's part of the escalation system itself.

Visible risk



**Can I manage this safely at my
current level?**



No / uncertain



Escalate clearly



- ☐ What is happening?
- ☐ What am I worried about?
- ☐ What has changed or does not fit?
- ☐ What do I need?

Chapter 26 — Handover, ownership and closed loops

Handover isn't complete once information has been spoken. A safe handover transfers the patient's current state, the working diagnosis and its uncertainty, the active risks, the actions already completed, the tasks still outstanding, the contingency plans, ownership of the next decision, and confirmation that the receiver has understood.

The I-PASS multicentre intervention combined structured handover content with training, observation and a wider implementation bundle, and was associated with reductions in medical errors and preventable adverse events — without increasing handover time [13]. The lesson isn't that one mnemonic solves handover on its own; it's that structured communication, contingency planning, receiver synthesis and implementation discipline all matter together.

A RADIAL handover

- Recognise instability: how ill is the patient now — stable, unstable or compensating?
- ABCDE prioritisation: which physiological problem is active?
- Define worst-case risk: what deterioration or diagnosis must the receiving clinician protect against?
- Identify contradictions: what doesn't fit the leading diagnosis or current plan?
- Act and escalate: what has been done, what remains, and what needs senior or specialty involvement?
- Loop and reassess: when should the patient be reviewed, what result must be checked, and what change should trigger action?

Before the handover ends, the receiver should be able to answer three things: what am I responsible for, by when, and what would make the plan change? If those answers aren't clear, the handover isn't finished yet.

Chapter 27 — Senior visibility and verification

Senior clinical decision-making is often described through knowledge and procedural skill, and those matter enormously, particularly for high-impact, low-frequency events. But much of senior practice is less dramatic than that: being visible enough to be found, hearing concerns without discouraging the next one, deciding which patients need personal review, helping others organise uncertainty, resolving ownership between teams, protecting task completion and handover, and holding a working picture of the whole department.

Visibility is operational

If ten clinicians each spend several minutes hunting for a senior decision-maker, the department loses clinical time and decisions slow down — and, just as importantly, junior clinicians learn that help is hard to reach.

When to verify personally

Senior clinicians are constantly choosing between trusting a colleague's assessment and reviewing the patient themselves. Personal verification tends to matter more when the patient is unstable or compensating; when the proposed disposition carries high consequence; when the history, examination and results don't cohere; when the presenting clinician is uncertain or inexperienced; when the case has already passed through several hands; or when the patient isn't responding as expected. The aim isn't to repeat every assessment. It's to place senior attention where it changes risk.

Senior verification without disempowerment

A personal review shouldn't erase the junior clinician's own reasoning. Asking them what their leading diagnosis is, what worries them most, what doesn't fit, what they've already done, and what they need from you turns the review into a decision and a teaching moment at once.

Situational awareness

Emergency departments demand awareness at the level of the patient, the team, and the department as a whole. Research in emergency medicine links team situational awareness with clinical performance and points to the importance of a shared mental model in dynamic environments [17,18]. For the senior clinician, that means looking beyond the next patient to the pattern of the department: who is deteriorating, who is waiting for a senior decision, which abnormal results are still open, which patients have been seen but not decided, where ownership is unclear, and what needs to be handed over explicitly rather than assumed. This is clinical leadership functioning as risk detection.

Chapter 28 — Results review, discharge and the senior sweep

The most visible work in an acute department is new assessment. The least visible risk usually sits in the unfinished work behind it — a patient seen but not decided, a result returned but not interpreted, a discharge letter left incomplete, a specialty discussion held with no clear ownership afterwards, a senior review requested that now exists only in someone's memory, a patient who improved but whose observations were never repeated.

The senior sweep

Before handover, or before the end of a shift, it's worth running the department through five questions: who is waiting for a senior decision, which abnormal results or investigations remain unclosed, who has been seen without a clear plan, which discharge, transfer or referral is blocked by an unresolved decision, and what needs to be handed over explicitly rather than assumed? This isn't administrative tidying — it is a direct reduction of hidden risk.

Before discharge

A brief final loop is worth running every time: is the patient physiologically safe now? Has the dangerous plausible diagnosis been sufficiently addressed? What didn't fit, and has it been explained? Has the patient improved as expected? Are the outstanding results and follow-up owned by someone? And does the patient understand what change should bring them back urgently? Discharge is not only a conclusion about the current diagnosis — it is a prediction about the patients near future outside the clinical environment.

Results do not close themselves

A requested test creates a future obligation: someone has to know that it's back, what it means in context, whether it changes the plan, whether the patient needs reassessing, and whether the loop is closed. Technology can display data. It cannot guarantee ownership.

Senior sweep

Before handover, check:

- 
- 1 **Who is waiting for a senior decision?**
 - 2 **Which results remain unclosed?**
 - 3 **Who has been seen but not decided?**
 - 4 **Which discharge or referral is blocked?**
 - 5 **What must be handed over explicitly?**
- 

Chapter 29 — RADIAL as team language

RADIAL isn't only for a clinician thinking silently on their own. In a deteriorating patient, different team members are often already performing different parts of the loop — a healthcare assistant notices the patient looks different, a nurse recognises abnormal physiology and begins assessing, a paramedic supplies trajectory and context, a junior doctor defines the immediate risks, a senior clinician makes the treatment, escalation and disposition decisions, and the team together reassesses the response. No single profession owns clinical thinking.

Shared questions create a shared model

Teams work more safely when everyone understands not just the tasks but the current risk model behind them. Useful prompts to share out loud include:

- Are we calling this patient stable, unstable or compensating?
- What is the active ABCDE problem?
- What are we most worried about missing?
- What doesn't fit our current explanation?
- What needs to happen in the next ten minutes?
- And when are we reassessing?

Questions like these turn parallel activity into coordinated care.

Psychological safety and clinical safety

A team cannot use RADIAL well if junior or non-medical staff are afraid to say that something doesn't fit. Psychological safety isn't the absence of standards or challenge — it is the ability to take the interpersonal risk of asking a question, admitting uncertainty, or raising a concern without being humiliated for it [14]. A demanding clinical environment can still be psychologically safe; the standard can stay high while the response to escalation stays respectful.

Senior clinicians shape this through ordinary behaviour — responding to concerns promptly, thanking staff for escalating before debating the conclusion, explaining decisions consistently, separating the person from the problem, admitting when new information changes the plan, and making it normal to ask what doesn't fit. The aim isn't to make every concern correct. It is to make the important ones speakable.

PART VI - Building a personal clinical thinking system

Chapter 30 — Your danger phrases and corrective questions

A framework earns its keep when it interrupts your own predictable pattern. Your danger phrase is the sentence you're most likely to say just before your strength tips into a risk.

Overloaded Processor

Strength: thorough, detail-aware, risk-sensitive. Danger phrase: "There are too many possibilities." Corrective question: "Which three risks change what I do now?" Practice habit: externalise the working model and choose the next safe action.

Premature Closer

Strength: decisive, efficient, action oriented. Danger phrase: "This is probably just..." Corrective question: "What would make this explanation incomplete or unsafe?" Practice habit: name one dangerous alternative and one contradiction before closing the case.

Hesitator

Strength: cautious, reflective, safety conscious. Danger phrase: "I will just wait for..." Corrective question: "Is there enough visible risk to act now?" Practice habit: define a decision threshold, act on the reversible safety steps, and name the cost of delay.

Intuitive Sprinter

Strength: rapid pattern recognition and confidence under pressure. Danger phrase: "I know what this is." Corrective question: "What would prove my instinct wrong?" Practice habit: name the impression, test it against physiology, and plan the second look.

Your personal interruption phrase

Choose one line you can actually use on shift — not everything is equally urgent; pause before it becomes fixed; visible risk justifies action; trust it, then test it; the case is not safe because it has been seen; a correct diagnosis can still be incomplete; stable is a moment, safety is a trajectory; run the loop. The best phrase isn't the most elegant one. It's the one that actually changes your behaviour.

Chapter 31 — Deliberate practice with real cases

Clinical reasoning develops through experience, but experience alone isn't enough. Practice becomes valuable when the clinician gets feedback, compares the initial model with what later emerged, and identifies which cue should change future decisions.

The three-case exercise

Choose three recent cases — one that went well, one that felt uncomfortable, and one where the diagnosis or plan changed along the way. For each, write down the initial frame (what was the first label or explanation?), the physiological state (stable, unstable or compensating?), the worst-case risk (what dangerous possibility or trajectory mattered most?), the contradiction (what didn't fit the first explanation?), the action threshold (what prompted action or escalation, and was it early enough?), the reassessment (what changed after treatment, time, results or further history?), the pressure pattern you fell into — overload, early closure, hesitation, or an intuitive sprint — and, finally, a future cue: what specific thing will you notice earlier next time?

From reflection to behaviour

A weak reflection ends with “I will be more careful.” A stronger one identifies a repeatable behaviour: “In future breathless patients with hyperglycaemia, I will still ask whether the history and examination support a second cardiac process.” “Before discharging after senior review, I will confirm that every requested result has a named owner.” “When a junior says something does not fit, I will ask for the contradiction before offering my own conclusion.” “At the end of the shift, I will run a senior sweep for patients seen but not decided.” Specific cues like these are what create usable learning.

Teaching with RADIAL

Vague feedback — “you need to think more broadly” — rarely changes anything. It lands better as: “Your first diagnosis was reasonable. The missing step was defining the worst-case risk.” “You recognised the abnormality, but you didn't set a reassessment point.” “The plan was safe initially. It became unsafe when the patient failed to improve and the loop was never restarted.” Feedback like this gives the learner a method, not just a verdict.

Chapter 32 — Confidence without false certainty

Confidence in acute care is often misunderstood. It isn't the belief that you're always right, and it isn't speaking loudly, acting quickly, or simply appearing certain. Real clinical confidence is quieter than that.

It is the ability to act while still thinking — to say, “This is my working diagnosis, and these are the risks I am still checking.” It is the ability to escalate without embarrassment, to ask a junior colleague what you may have missed, and to change your mind when the patient changes. It is the ability to move before certainty while still respecting the uncertainty that remains, and the discipline to finish the work — to see the result, complete the handover, return to the patient, and close the loop.

RADIAL is built to support that kind of confidence. Not false certainty. Safe movement.

The final loop

When everything feels important, run the loop. When the first diagnosis feels obvious, run the loop. When you're waiting for certainty, run the loop. When instinct is moving quickly, run the loop. When the nurse says the patient is different, run the loop. When the result doesn't match the patient, run the loop. When the patient has been seen but not decided, run the loop. After treatment, after investigation, before transfer, and before discharge — run the loop.

RADIAL quick reference

R — Recognise instability

Question: Stable, unstable or compensating? Look for abnormal physiology, altered behaviour, subtle deterioration, limited reserve, trajectory, and team or family concern. Risk: false reassurance from a snapshot, score or near-normal number.

A — ABCDE prioritisation

Question: What kills first? Look for immediate airway, breathing, circulation, disability and exposure threats. Risk: treating the diagnostic label before securing physiology.

D — Define worst-case risk

Question: What cannot be missed? Look for a plausible diagnosis, complication or trajectory that would change immediate management. Risk: confusing the most likely diagnosis with a sufficiently safe one.

I — Identify contradictions

Question: What does not fit? Look for unexplained findings, discordant physiology, an unexpected trajectory, ambiguous ownership, or incomplete task closure. Risk: filtering inconvenient information through the first explanation.

A — Act and escalate

Question: What needs action now? Look for the next step that reduces immediate risk, the escalation threshold, and the named owner. Risk: waiting for certainty, or assuming someone else will decide.

L — Loop and reassess

Question: What has changed? Look for the response to treatment, evolving observations, new results, revised history and new concern. Risk: assuming the original diagnosis and plan remain correct.

One-line memory aid: Safe now? First threat? Worst case? What does not fit? What needs doing? What has changed?

The four patterns at a glance

Overloaded Processor

STRENGTH

Thoroughness

RISK

Everything feels equally important

CORRECTIVE QUESTION

Which three risks change what I do now?

Premature Closer

STRENGTH

Decisiveness

RISK

Diagnosis fixes too early

CORRECTIVE QUESTION

What would make this unsafe?

Hesitator

STRENGTH

Caution

RISK

Waits for certainty

CORRECTIVE QUESTION

Is there enough visible risk to act?

Intuitive Sprinter

STRENGTH

Rapid pattern recognition

RISK

Instinct goes untested

CORRECTIVE QUESTION

What would prove me wrong?

Evidence base and limitations

RADIAL is an author-developed framework. It has not been prospectively validated as a diagnostic tool, a clinical prediction rule, or a patient-safety intervention, and no claim is made that using the acronym will independently reduce diagnostic error or improve outcomes.

Its components are, however, aligned with several established bodies of evidence and professional practice: diagnostic safety as a collaborative, longitudinal process rather than a single act; [1] the complementary roles of pattern recognition and analytical reasoning; [2] selective deliberate reflection in complex, difficult or bias-prone cases; [3–5] cognitive load and the value of structuring complex information; [6] mixed evidence for generic cognitive-forcing training, which supports a cautious approach to any claims made here; [7,8] the potential effect of fatigue and contextual distractors on clinical reasoning; [9,10] ABCDE assessment, early help and regular reassessment in deteriorating patients; [11] standardised physiological assessment through NEWS2; [12] structured handover and contingency planning; [13] psychological safety, speaking up and leadership response to concerns; [14–16,19] and situational awareness and shared mental models in emergency teams; [17,18]

None of this supports the simple conclusion that every diagnostic error results from cognitive bias, that slow reasoning is always safer, or that a checklist can replace expertise. Knowledge, experience, systems, communication, workload, technology and access to timely investigation all shape diagnosis and care together.

RADIAL should be used as a prompt for structured clinical thinking, a shared language for supervision and escalation, a framework for reflection and teaching, and a reminder to build reassessment and ownership into every plan. It should not be used to override clinical judgement, delay treatment, replace disease-specific guidance, or justify indiscriminate investigation.

References

1. National Academies of Sciences, Engineering, and Medicine. Improving Diagnosis in Health Care. Washington, DC: National Academies Press; 2015.
2. Norman GR, Eva KW. Diagnostic error and clinical reasoning. *Medical Education*. 2010;44(1):94-100. doi:10.1111/j.1365-2923.2009.03507.x.
3. Mamede S, Schmidt HG, Penaforte JC. Effects of reflective practice on the accuracy of medical diagnoses. *Medical Education*. 2008;42(5):468-475. doi:10.1111/j.1365-2923.2008.03030.x.
4. Mamede S, van Gog T, van den Berge K, et al. Effect of availability bias and reflective reasoning on diagnostic accuracy among internal medicine residents. *JAMA*. 2010;304(11):1198-1203. doi:10.1001/jama.2010.1276.
5. Mamede S, Hautz WE, Berendonk C, et al. Think twice: effects on diagnostic accuracy of returning to the case to reflect upon the initial diagnosis. *Academic Medicine*. 2020;95(8):1223-1229. doi:10.1097/ACM.0000000000003153.
6. van Merriënboer JJG, Sweller J. Cognitive load theory in health professional education: design principles and strategies. *Medical Education*. 2010;44(1):85-93. doi:10.1111/j.1365-2923.2009.03498.x.
7. Sherbino J, Yip S, Dore KL, Siu E, Norman GR. The effectiveness of cognitive forcing strategies to decrease diagnostic error: an exploratory study. *Teaching and Learning in Medicine*. 2011;23(1):78-84. doi:10.1080/10401334.2011.536897.
8. Lambe KA, O'Reilly G, Kelly BD, Curristan S. A cognitive forcing tool to mitigate cognitive bias: a randomised control trial. *BMC Medical Education*. 2019;19:12. doi:10.1186/s12909-018-1444-3.
9. Collen J, Durning S, Berk J, et al. Exploring sleep duration and clinical reasoning process in resident physicians: a thematic analysis. *Journal of Clinical Sleep Medicine*. 2024;20(8):1279-1289. doi:10.5664/jcsm.11134.
10. Monteiro SD, Sherbino J, Sibbald M, Norman GR. Disrupting diagnostic reasoning: do interruptions, instructions, and experience affect the diagnostic accuracy and response time of residents and emergency physicians? *Academic Medicine*. 2015;90(4):511-517. doi:10.1097/ACM.0000000000000614.
11. Resuscitation Council UK. The ABCDE Approach. Updated July 2024.
12. Royal College of Physicians. National Early Warning Score (NEWS) 2: Standardising the Assessment of Acute Illness Severity in the NHS. Updated report of a working party. London: RCP; 2017.
13. Starmer AJ, Spector ND, Srivastava R, et al. Changes in medical errors after implementation of a handoff program. *New England Journal of Medicine*. 2014;371:1803-1812. doi:10.1056/NEJMs1405556.

14. Edmondson AC. Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*. 1999;44(2):350-383. doi:10.2307/2666999.
15. Schwappach DLB, Gehring K. Speak up-related climate and its association with healthcare workers' speaking up and withholding voice behaviours: a cross-sectional survey in Switzerland. *BMJ Quality & Safety*. 2018;27(10):836-843. doi:10.1136/bmjqs-2017-007388.
16. Martinez W, Etchegaray JM, Thomas EJ, et al. Barriers to speaking up about patient safety concerns. *Journal of Patient Safety*. 2021;17(8):e1437-e1443.
17. Lowe DJ, Ireland AJ, Ross A, Ker J. Exploring situational awareness in emergency medicine: developing a shared mental model to enhance training and assessment. *Postgraduate Medical Journal*. 2016;92(1093):653-658. doi:10.1136/postgradmedj-2015-133772.
18. Gardner AK, Kosemund M, Martinez J. A simulation-based approach to measuring team situational awareness in emergency medicine: a multicenter, observational study. *Academic Emergency Medicine*. 2017;24(10):1239-1247.
19. General Medical Council. Raising and Acting on Concerns about Patient Safety. Updated December 2024.

Next step — from framework to practice

A playbook can introduce a framework. Repeated practice is what develops the judgement to use it under pressure.

Clinical Thinking Under Pressure: The RADIAL Course will take the framework into realistic, evolving clinical cases. You will practise recognising deterioration, identifying what does not fit, defining the risk that cannot be missed, choosing when to act or escalate, and reassessing when the patient or the information changes.

The course will include:

- evolving case-based lessons
- practical decision checkpoints
- escalation and handover exercises
- guided reflection on your own pressure pattern
- tools that can be used in clinical practice, supervision and teaching

The objective is not to memorise an acronym. It is to become more deliberate at the moments when uncertainty, workload and consequence meet.

Continue with RADIAL

Join the early-access list for course updates, new clinical cases and launch information:

[Oakmere Clinical Services — RADIAL Course Early Access](#)

Run the loop.

About the author



Dr James Wallace is a Consultant in Emergency Medicine, clinical leader, author and medical educator. His work spans acute care, clinical governance, emergency preparedness, remote and prehospital medicine, supervision and professional education.

He is the founder of **Oakmere Clinical Services** and the author of *Diagnosis Leader*. His work focuses on helping clinicians recognise risk, make safer decisions under pressure, communicate uncertainty clearly and lead effectively in complex clinical environments.

He developed RADIAL from patterns he repeatedly encountered in frontline practice: capable clinicians who know their medicine and care deeply about their patients, but still meet the very human limits of workload, uncertainty, interruption and incomplete information.

His central teaching principle is simple:

Good clinical thinking is not the absence of uncertainty. It is the ability to recognise risk, act safely and reassess deliberately while uncertainty remains.

Further books, clinical resources and RADIAL updates: www.oakclinservices.com