# The OR Teaching Manual

*A Field Guide for the Anaesthesiologist Who Teaches Medical Students*

Bruno C. da Rosa Borges MD

# THE OR TEACHING MANUAL

*A field guide for the anaesthesiologist who teaches medical students*

## How this manual is organised

Each topic has two parts:

1. **What to teach the student** — the lesson menu, in plain language.
2. **Teaching review** — a refresher for you, with cases, pearls, and diagrams.

*Conceptual topics that lend themselves to detailed prose*

*(pain pathways, pharmacokinetics, EEG bands, cognitive bias, situational awareness)*

*are presented in full explanations rather than bullets.*

# Contents

# Introduction

Anaesthesiology is one of the most physiologically rich and immediately consequential specialties in medicine. Every induction is a small experiment in pharmacology: you give a drug, predict an effect, watch the monitors, and titrate. Every airway is an applied anatomy lesson. Every difficult case is a microcosm of teamwork, decision-making, and patient safety. For a pre-clerkship medical student, an OR rotation is therefore one of the highest-yield places they will ever stand — provided someone takes the time to point at the right things and explain why they matter.

This manual exists because we — anaesthesiologists — are often great teachers in spirit but pressed for time in practice. Cases run, the surgeon is waiting, and the student appears at the head of the bed with hopeful eyes. The intent here is to give you a pocket reference: open to a topic, see at a glance what is worth covering, and refresh the underlying concepts in a few minutes before turning to the student.

The choice of topics follows the principles laid out in modern anaesthesia education writing — that anaesthesiologists are uniquely positioned to teach applied physiology, pharmacology, decision-making under uncertainty, patient safety, and the technical instrumentation of clinical medicine. The structure is deliberately modular: any single topic can be taught in 5 minutes between the time-out and the incision, or expanded into a 30-minute teaching round at the end of a list.

> **How to use this manual on a teaching day**
>
> - Pick one topic. Don't try to cover everything in a single case.
> - Anchor the topic to something the student is watching right now (the monitor, the syringe, the airway).
> - Ask one open question before you explain. ("What do you think is happening to her blood pressure right now, and why?")
> - Close the loop: name a take-home in one sentence. Students remember sentences, not lectures.

# Part 1 · Applied physiology and pharmacology

The OR is the closed-loop pharmacology lab that students never get to see in pre-clerkship. Drugs go in, signals come out within seconds, and you adjust. Use that. The three topics that follow give a structured way to convert what you do every day into teachable physiology.

## 1.1 Cardiopulmonary physiology in real time

**What to teach the student**

- Preload, afterload, contractility — what each one means and what we do to change it.
- Oxygen delivery ($DO_2$) as the central goal of resuscitation: $DO_2 = CO \times CaO_2$.
- How positive-pressure ventilation interacts with venous return and cardiac output.
- Why an arterial line traces beat-to-beat physiology you can't see from the cuff.
- V/Q matching: dead space vs shunt and why supplemental $O_2$ helps one but not the other.

### Teaching review — refresh before you teach

The unifying concept worth giving a pre-clerkship student is oxygen delivery. Everything we do in the OR — fluid, vasopressors, inotropes, ventilation, transfusion — is in service of keeping $DO_2$ adequate for the tissues. Writing the equation on a glove or a piece of tape and walking through each variable takes ninety seconds and reframes most of clinical medicine for them.

$\mathbf{DO_2 = CO \times CaO_2} = (HR \times SV) \times (1.34 \times Hb \times SaO_2) + (0.003 \times PaO_2)$. Cardiac output is the pump term, $CaO_2$ is the carrying-capacity term. The dissolved-$O_2$ component is usually negligible at sea level — which is also why giving 100% $O_2$ to an anaemic patient does very little. Showing the student that dropping Hb from 12 to 6 halves $CaO_2$, while raising $FiO_2$ from 21% to 100% adds a rounding error to the dissolved component, is one of those lightbulb moments.

On the cardiac side, the three levers that anaesthesiologists pull constantly — preload, afterload, contractility — map directly onto the drugs in our trays. A fluid bolus moves preload. Phenylephrine moves afterload (and reflexively preload via venoconstriction). Ephedrine and adrenaline move contractility (and HR). Watching an arterial line tracing during a phenylephrine bolus, with the student narrating which lever changed, is more memorable than any lecture on Frank-Starling.

Positive-pressure ventilation deserves its own minute. PPV raises intrathoracic pressure during inspiration, which transiently reduces venous return and therefore stroke volume. This is exactly why we see arterial-line pulse-pressure variation: SV swings with the respiratory cycle. In a fluid-responsive patient (operating on the steep part of the Starling curve), the swing is exaggerated. In a fluid-replete patient, it flattens out. Pulse-pressure variation > 13% is a reasonable bedside marker of fluid responsiveness in mechanically ventilated, sinus-rhythm patients — and it's a more elegant teaching point than "give them some fluid and see what happens."

🩺 **Case: Hypotension after induction**

BP drops from 130/80 to 75/40 ninety seconds after propofol and remifentanil. Ask the student: which lever broke? In most cases it's afterload (vasodilation from propofol) plus a smaller drop in preload (venodilation) and contractility. Treatment follows the lever: phenylephrine bolus (afterload), small fluid bolus (preload), reduce or pause infusions if the patient is too deep. Walk through why ephedrine might be a better first choice in a bradycardic patient, and why noradrenaline is now the modern OR alternative to repeated phenylephrine boluses.

💡 **Clinical pearl**

If a student remembers nothing else, have them remember: "Treat the lever that broke." Hypotension is not a diagnosis — it's a sign that one of preload, afterload, contractility, rate, or rhythm has changed. Always ask which.

## 1.2 Pharmacokinetics and pharmacodynamics at the bedside

**What to teach the student**

- PK = what the body does to the drug (absorption, distribution, metabolism, excretion).
- PD = what the drug does to the body (receptor effect, dose-response).
- Volume of distribution: why lipophilic drugs (propofol, fentanyl) seem to disappear quickly after a single dose.
- Context-sensitive half-time: why the same drug behaves differently after 5 minutes vs 5 hours of infusion.
- Active metabolites: why morphine accumulates in renal failure but hydromorphone is safer.
- Receptor pharmacology: agonist, antagonist, partial agonist — with naloxone as the classroom example.

### Teaching review — refresh before you teach

Pharmacokinetics is one of the most useful topics you can teach because it explains so many things at once: why induction agents wake patients up in minutes after a bolus but take hours after a long infusion, why dosing changes in the elderly, why some drugs need a loading dose and others don't. The trick is to teach it through the drugs the student is watching go in, not through textbook compartment models.

Start with the simplest case — propofol after a single bolus. The patient is asleep within 30 seconds because the drug rapidly equilibrates with brain tissue (a small, well-perfused compartment). They wake up within 8–10 minutes not because the drug has been metabolised, but because it has redistributed to a much larger compartment of muscle and fat. This is why a single bolus is so reliable for induction: clearance is almost irrelevant. The drug isn't gone, it's just somewhere else.

Now contrast that with a 6-hour propofol infusion. The fat and muscle compartments have filled up. There is nowhere left to redistribute to. The patient now genuinely has to clear the drug, which takes much

longer. This is what "context-sensitive half-time" captures — the time required for the drug concentration in the central compartment to halve depends on the context (how long it has been infusing). The shape of this curve is dramatically different for different drugs.

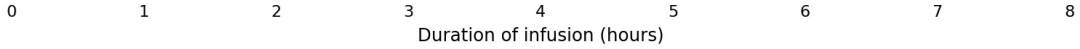


*Context-sensitive half-time. Remifentanil's curve stays flat — the reason it is the favourite agent for fast-track and for cases where fast emergence matters.*

Remifentanil is the outlier worth pointing out: it's metabolised by non-specific plasma esterases, so its CSHT is essentially flat at about 3–4 minutes regardless of infusion duration. Fentanyl and (especially) thiopental are the opposite — long infusions produce hours of residual effect.

The other PK/PD point worth a few minutes is the role of active metabolites. Morphine is cleared by the liver into morphine-3-glucuronide (no analgesic activity, mildly neuroexcitatory) and morphine-6-glucuronide (a potent μ-agonist, renally excreted). In renal failure, M6G accumulates and patients can become dangerously sedated days into therapy. Hydromorphone has a less troublesome metabolite profile, which is why it's often preferred in renal impairment. This is the kind of teaching point a pre-clerk will not get on internal medicine and will remember for the rest of their career.

> 💡 **Clinical pearl**
>
> Receptor pharmacology becomes vivid the moment you give naloxone. Antagonist with high μ-receptor affinity, short half-life (~30 minutes). The duration mismatch between naloxone and most opioids (especially methadone, fentanyl patches, sustained-release morphine) is why re-sedation is the classic teaching point. Use the same logic for flumazenil and benzodiazepines.

> 🩺 **Case: The opioid-naïve elderly patient**
>
> An 82-year-old with creatinine 180 μmol/L receives 10 mg of morphine on the ward post-op. Forty-eight hours later they are obtunded. What happened? Walk the student through reduced GFR → M6G accumulation → progressive sedation. Then ask: what should we have ordered instead? Hydromorphone, smaller doses, longer intervals, or a regional technique.

## 1.3 How anaesthetic drugs reveal core pharmacology

> **What to teach the student**
>
> - GABA-A receptor: the common pathway for propofol, benzodiazepines, volatiles, barbiturates.
> - NMDA receptor and ketamine — why ketamine is fundamentally different from everything else.
> - Opioid μ-receptor: agonist (morphine), antagonist (naloxone), partial agonist (buprenorphine).
> - Acetylcholine and muscle relaxation: nicotinic vs muscarinic receptors at a glance.
> - MAC (minimum alveolar concentration) as a clean example of potency vs efficacy.

### Teaching review — refresh before you teach

Anaesthetic agents are unusually good teaching tools for receptor pharmacology because the effects are immediate, dose-related, and reversible. Within a single case you can show a student a GABA agonist, an NMDA antagonist, an opioid receptor agonist, a muscarinic antagonist (glycopyrrolate, atropine), and an acetylcholinesterase inhibitor (neostigmine, sugammadex more selective).

GABA-A is the workhorse. Propofol, benzodiazepines, volatile anaesthetics, and barbiturates all enhance GABAergic inhibition, just at slightly different sites on the same receptor complex. This explains why their effects are additive and why combining them risks respiratory depression. It also explains why flumazenil (a competitive benzodiazepine antagonist) reverses midazolam but not propofol — they bind different sites.

Ketamine is the teaching gem because it breaks the rules. It's an NMDA antagonist (blocks glutamate excitation) rather than a GABA enhancer. Clinically that translates into preserved airway reflexes, preserved respiratory drive, sympathetic stimulation rather than depression, and dissociative analgesia. For a pre-clerkship student, the lesson is that "anaesthetic" is not one mechanism — it's an outcome that can be reached through several different pharmacologies.

MAC — *minimum alveolar concentration* — is the concentration of inhaled agent at which 50% of patients do not move in response to a standardised noxious stimulus. It's a beautifully clean teaching example. MAC values let you compare potencies of volatiles directly: sevoflurane MAC ≈ 2%, isoflurane ≈ 1.15%, desflurane ≈ 6%, nitrous oxide ≈ 104% (you cannot achieve 1 MAC with $N_2O$ alone at sea level, which is exactly why we use it as an adjunct). MAC additivity also illustrates the concept of pharmacological synergy: 0.5 MAC sevo + 70% $N_2O$ ≈ 1 MAC.

> 💡 **Clinical pearl**
>
> MAC drops with age (about 6% per decade after 40), with hypothermia, hyponatraemia, opioids, and acute alcohol intoxication. It rises with chronic alcohol use, hyperthermia, and hypernatraemia. Showing the student the depth-of-anaesthesia monitor and the end-tidal agent number, and asking them to predict whether you need more or less in the patient in front of you, integrates several concepts at once.

# Part 2 · Procedural and technical skills

Pre-clerkship students rarely get to do anything with their hands. The OR is one of the few places where letting a student attempt a venipuncture under direct supervision is appropriate, low-risk, and educationally enormous. The topics below are calibrated for the level of student you'll have: most can be demonstrated and discussed even when the student doesn't actually perform them.

## 2.1 Cardiopulmonary resuscitation and ACLS fundamentals

**What to teach the student**

- Why high-quality chest compressions matter more than anything else: rate, depth, full recoil, minimal interruptions.
- Coronary perfusion pressure as the proximate target of compressions.
- Shockable (VF, pulseless VT) vs non-shockable rhythms (PEA, asystole) — and why the algorithm differs.
- Why we give adrenaline — and the modern controversy about it.
- The importance of identifying and treating reversible causes (the Hs and Ts).
- Post-arrest care: temperature, oxygenation, blood pressure, and neuroprognostication delays.

### Teaching review — refresh before you teach

The single most useful thing to communicate to a pre-clerk is that during CPR, blood flow is whatever your hands are producing. There is no autonomous circulation. Every interruption — to check a pulse, to look at the monitor, to intubate — drops coronary and cerebral perfusion pressure to zero, and it takes 30–60 seconds of good compressions to rebuild it. This is why modern resuscitation is obsessive about minimising hands-off time.

Coronary perfusion pressure (aortic diastolic minus right atrial diastolic pressure) is the physiological target of compressions. Studies in arrested patients suggest a CPP > 15 mmHg correlates with return of spontaneous circulation. Adrenaline raises CPP by causing peripheral vasoconstriction, which is the rationale for its use; the controversy is that while ROSC rates improve, neurologically intact survival benefit is small and the dose-response is not well characterised.

The shockable-vs-non-shockable distinction is worth ten minutes. VF and pulseless VT are electrical problems with mechanical consequences — defibrillation is the definitive treatment, and time-to-shock is the most important determinant of survival. PEA and asystole are usually mechanical or metabolic problems (the heart isn't being given the conditions to work), so the treatment is to find and fix the underlying cause. The Hs and Ts (hypoxia, hypovolaemia, hypo/hyperkalaemia, hypothermia, hydrogen ion / acidosis; tension pneumothorax, tamponade, toxins, thrombosis pulmonary or coronary) are the differential a student should learn cold.

🩺 **Case: Witnessed arrest in the OR**

A patient on the table goes from sinus rhythm to coarse VF mid-laparotomy. Walk the student through what happens in the first 90 seconds: call it out, start compressions immediately (over the drape if necessary), defib pads on, charge, clear, shock at 200J biphasic, resume compressions. Emphasise: don't pause to check pulse before resuming compressions, don't intubate during compressions if the airway is already secured, don't give adrenaline before the first shock in a witnessed VF arrest. This is the kind of choreography that's invisible until someone narrates it.

⚠️ **Pitfall to flag for the student**

The most common error in pre-clerkship simulations is checking for a pulse during compressions or checking the monitor for too long between cycles. Coach the student to keep their eyes on the chest and the clock, not the monitor.

## 2.2 Airway management

**What to teach the student**

- Why airway management is the defining skill of anaesthesia and emergency medicine.
- Predictors of difficult bag-mask ventilation: bearded, edentulous, obese, snoring, advanced age.
- Predictors of difficult laryngoscopy: Mallampati class, thyromental distance, mouth opening, neck mobility.
- Bag-mask ventilation technique: head-tilt, chin-lift, two-handed seal, oral airway adjunct.
- Supraglottic devices (LMAs) as a bridge between mask and tube.
- What direct vs video laryngoscopy looks like, and what the operator is trying to achieve.
- The "can't intubate, can't oxygenate" pathway and why the front-of-neck airway exists.

### Teaching review — refresh before you teach

If you have one airway teaching session with a student, make it about bag-mask ventilation, not intubation. BMV is the skill that saves lives. Intubation is the skill that gets attention. A student who can BMV competently has bought themselves five minutes in any emergency they will ever face — and five minutes is enough to call for help.

The classic predictors of difficult BMV (mnemonic MOANS — Mask seal, Obesity/Obstruction, Age >55, No teeth, Stiff lungs) and difficult laryngoscopy (LEMON — Look externally, Evaluate 3-3-2, Mallampati, Obstruction, Neck mobility) are worth knowing, but the more useful skill is to develop a habit of asking "is this airway going to be hard, and if so, in what way?" before every induction. Showing a student your own pre-induction airway assessment, narrating each finding, is one of the highest-yield teaching moments in the OR.

Mallampati is the one most students remember from textbooks, often without understanding what they're looking at. It's a static visualisation of the posterior pharynx with the patient sitting upright, mouth

open, tongue protruded, no phonation. Class I: full uvula, soft palate, hard palate visible. Class II: tip of uvula obscured. Class III: only soft palate. Class IV: only hard palate. It's not a great solo predictor (sensitivity around 50%), but combined with thyromental distance, mouth opening, and neck extension it gives a useful pre-test probability of difficulty.

> 💡 **Clinical pearl**
>
> Three positions, three reasons: the sniffing position (neck flexion at the lower cervical spine, head extension at the atlanto-occipital joint) aligns the oral, pharyngeal, and laryngeal axes for direct laryngoscopy. The ramped position (shoulders elevated to bring the external auditory meatus level with the sternal notch) is the same idea for obese patients. The reverse-Trendelenburg adds gravity to keep diaphragmatic contents off the lungs and is your friend for pre-oxygenation in obesity.

> 🩺 **Case: The obese patient with OSA for emergency surgery**
>
> BMI 42, untreated OSA, NPO status uncertain, needs an emergency cholecystectomy. Walk the student through your plan: ramped position, denitrogenation with 100% $O_2$ for 3 minutes (or 8 vital-capacity breaths if cooperative), apnoeic oxygenation via nasal cannula at 15 L/min, rapid-sequence induction with cricoid pressure (debated but still used), video laryngoscopy as primary, prepared backup plan with LMA and front-of-neck access available. Each decision has a why.

## 2.3 IV access and common procedural skills

> **What to teach the student**
>
> - Anatomy of common access sites: dorsum of hand, forearm, antecubital fossa, external jugular.
> - Why we choose the most distal usable vein first.
> - Vein selection cues: visibility, palpability, straightness, valve location, fragility.
> - Sterile technique appropriate to the procedure (peripheral IV vs central line vs arterial line).
> - Troubleshooting: blown veins, valves, dehydration, vasoconstriction.
> - When to escalate: ultrasound-guided IV, intraosseous, central access.
> - Arterial vs venous puncture — visual, tactile, and pressure cues.

### Teaching review — refresh before you teach

IV access is the gateway procedure. Once a student can find a vein, hold a needle, and stay calm, every other percutaneous procedure becomes a variation on the theme. Spend time on the cognitive part — how to look at a forearm and decide what's most likely to work — not just the motor skill of inserting the cannula.

The classical hierarchy is: dorsum of hand → forearm → antecubital fossa → external jugular → ultrasound-guided peripheral → intraosseous → central. The reason to start distal is that if you fail proximally first, you've blown the vein for everything downstream. Antecubital fossa veins are tempting because they're easy, but they bend with elbow flexion and the patient gets occlusion alarms all night.

> 💡 **Clinical pearl**
>
> A frequently missed cue: a vein that flattens with arm elevation but refills when you drop the arm is patent. A vein that stays flat after elevation is either thrombosed or the tourniquet isn't tight enough. Teach this gesture as a quick patency check.

The rise of point-of-care ultrasound has shifted the threshold for difficult IV access dramatically. A reasonable rule for trainees: two failed blind attempts on a patient flagged as difficult, switch to ultrasound. We'll come back to ultrasound technique in section 2.4, but the IV-specific point is that ultrasound lets you cannulate veins 1–2 cm deep that you simply cannot palpate — most commonly the basilic vein in the medial upper arm.

> ⚠️ **Pitfall to flag for the student**
>
> The classic teaching error is letting the student push too hard, too fast. Slow, steady advancement until flashback, then drop the angle and advance the catheter, not the needle. Most failed IVs are technique failures of speed and angle, not anatomy.

## 2.4 Ultrasound-guided procedures: principles for the beginner

> **What to teach the student**
>
> - What ultrasound is: high-frequency sound waves, reflected by tissue interfaces, processed into images.
> - Probe selection: linear (high frequency, shallow, high resolution) vs curvilinear (low frequency, deep, lower resolution).
> - Knobology: gain (overall brightness), depth, frequency.
> - Identifying artery vs vein: round vs oval, thick vs thin wall, non-compressible vs compressible, pulsatile vs not, Doppler.
> - In-plane vs out-of-plane needle approach — what's visible, and the trade-offs.
> - The cardinal rule: keep the needle tip in view at all times.

### Teaching review — refresh before you teach

Ultrasound has become the universal procedural skill in modern anaesthesia and is rapidly becoming so in emergency medicine, internal medicine, and critical care. A pre-clerk who learns the basics from you is years ahead of their peers. The conceptual framework matters more than scanning skill at this stage — they will not be doing a brachial plexus block, but they should understand what they're seeing on a screen by the time they hit clerkship.

*Left: distinguishing artery from vein on cross-section. Right: in-plane vs out-of-plane needle approach. Both are essential mental models for any ultrasound-guided procedure.*

The artery-vs-vein distinction is the universal first lesson. Veins are oval, thin-walled, and collapse with gentle probe pressure. Arteries are round, thick-walled, non-compressible, and pulsate. If you're still unsure, colour Doppler shows the alternating direction of arterial flow. This is the same skill whether you're putting in an IV, an arterial line, or a central line — and getting it wrong has consequences ranging from inconvenient (cannulating a vein when you wanted an artery for a line) to catastrophic (carotid puncture during attempted internal jugular access).

The in-plane vs out-of-plane choice depends on the goal. In-plane (long-axis) shows the entire needle as a bright line — the gold standard for nerve blocks where you must know exactly where the tip is going relative to a target. Out-of-plane (short-axis) is faster for vascular access because the target vessel is shown in cross-section, but you only see the needle as a single bright dot. The trap is mistaking the needle shaft for the tip — which is why slow, deliberate advancement with constant tip-walking (tilting the probe to keep the dot in view as you advance) matters.

> ⚠️ **Pitfall to flag for the student**
>
> Students often want to look at their hands instead of the screen. Coach the rule: "once the needle is in the skin, your eyes are on the screen, not your hands." The probe is doing the seeing for you.

## 2.5 How clinical monitors work — and how they lie

> **What to teach the student**
>
> - Pulse oximetry: the physics, the assumptions, and the ways it fails.
> - Capnography waveforms: phases and patterns of normal and abnormal traces.
> - Non-invasive vs invasive blood pressure measurement.
> - ECG lead placement and what each lead 'sees'.
> - Why we look at trends, not single numbers.
> - Why every monitor reading needs to be interpreted in clinical context.

### Teaching review — refresh before you teach

#### Pulse oximetry

Pulse oximetry is the canonical example of an indispensable monitor that lies confidently. It works on the Beer-Lambert principle: oxyhaemoglobin and deoxyhaemoglobin absorb red (660 nm) and infrared (940 nm) light differently. The pulsatile component of the signal is assumed to come from arterial blood, so the ratio of absorptions is calibrated against direct measurements of arterial saturation in healthy volunteers. That last clause is where most failures live.

Pulse oximetry will mislead you (or read incorrectly) in: carbon monoxide poisoning (carboxyhaemoglobin absorbs almost identically to oxyhaemoglobin at 660 nm — the $SpO_2$ reads 100% while the patient is hypoxic); methaemoglobinaemia (the $SpO_2$ trends toward 85% regardless of true saturation); severe

anaemia (the absolute number can be reassuring while $DO_2$ is critically low); poor perfusion (no pulsatile signal, no reading, or a misleading low one); nail polish, especially blue and black; deeply pigmented skin under poor signal conditions; and motion artefact.

> 💡 **Clinical pearl**
>
> Teaching question: "Your patient has $SpO_2$ 100% on room air, looks blue, and was just rescued from a house fire. What's going on?" Carbon monoxide poisoning. The pulse oximeter is reading carboxyhaemoglobin as oxyhaemoglobin. You need a co-oximeter or arterial blood gas to see the truth. This single example rewires how a student thinks about all monitoring.

## Capnography

Capnography is arguably the most important monitor in anaesthesia — and the one with the steepest hidden curriculum. Every $CO_2$ waveform tells a story about ventilation, perfusion, equipment, and metabolism. A student who learns to read capnograms is years ahead.

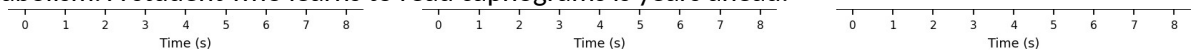


*Capnography waveform patterns. Normal trace at top left; common pathologies elsewhere. The shape often diagnoses the problem before the numbers do.*

A normal capnogram has four phases. Phase 0 (inspiration, baseline at zero), Phase I (anatomical dead-space exhalation, still near zero — gas from conducting airways with no $CO_2$), Phase II (rapid upstroke as alveolar gas reaches the sensor), Phase III (alveolar plateau with slight upslope; the value at end-expiration is the $EtCO_2$). The shape, baseline, height, and slope each carry diagnostic information.

- Sudden flat trace → arrest, disconnection, oesophageal intubation, complete obstruction.
- Slow gradual decline → falling cardiac output, hyperventilation, hypothermia.
- Sudden drop with preserved waveform → pulmonary embolus, sudden severe hypotension.
- Elevated baseline ($CO_2$ > 0 in inspiration) → rebreathing, exhausted $CO_2$ absorber.
- "Shark fin" upslope (no plateau) → bronchospasm, severe COPD, kinked tube.
- Curare cleft (notch in plateau) → returning neuromuscular activity during partial paralysis.

> 💡 **Clinical pearl**
>
> $EtCO_2$ during CPR is one of the strongest prognostic markers we have. Sustained $EtCO_2$ < 10 mmHg after 20 minutes of high-quality compressions is associated with very low survival. A sudden rise in $EtCO_2$ during compressions often signals ROSC before any other monitor changes.

> 🩺 **Case: Sudden $EtCO_2$ drop in mid-laparoscopy**
>
> $EtCO_2$ drops from 36 to 22 mmHg over 60 seconds. Arterial saturation begins to fall. Differential: gas embolus from $CO_2$ insufflation (consider, but rare with proper technique), pulmonary embolus, sudden hypotension, circuit disconnection, accidental extubation, cardiac arrest. Walk through how the rest of the picture (BP, $ETCO_2$ shape, breath sounds, surgical view) refines the differential. This is exactly the kind of integrated

diagnostic moment that pre-clerks rarely see explicitly named.

### Blood pressure: NIBP vs arterial line

The non-invasive blood pressure cuff is convenient and surprisingly inaccurate at the extremes. It tends to overestimate hypotension and underestimate hypertension, and it samples discretely (every 3–5 minutes typically). The arterial line gives beat-to-beat measurement, a waveform that contains additional information (slope, dicrotic notch, pulse-pressure variation), and access for blood sampling. The teaching point is that monitoring choice should match clinical risk: a healthy young patient for a 30-minute procedure doesn't need an arterial line; a septic patient on noradrenaline does.

## 2.6 The anaesthesia machine and oxygen delivery devices

### What to teach the student

- The basic gas flow path: pipeline → flowmeters → vaporiser → common gas outlet → circuit → patient.
- Vaporiser function: how a constant agent concentration is produced from a liquid agent.
- The circle circuit: inspiratory and expiratory limbs, soda lime absorber, fresh gas inflow.
- Fail-safes: oxygen pressure failure alarm, hypoxic guard, scavenging.
- Oxygen delivery devices in escalating order: nasal prongs, simple mask, non-rebreather, Venturi, high-flow nasal oxygen, NIV, intubation.
- $FiO_2$ estimates and their limitations.

### Teaching review — refresh before you teach

Most students go through medical school without ever understanding what the box behind the patient's head actually does. Five minutes with the back of a machine open is the most memorable physics lesson in the curriculum. If you can't open the back, even just walking through the diagram below works.

*Simplified gas flow path of a modern anaesthesia machine. Each stage is a teaching opportunity in physics, safety engineering, and clinical practice.*

The vaporiser is the magic. Liquid sevoflurane (or iso, or des) sits in a chamber. A controlled fraction of the fresh gas flow is diverted through the chamber, picks up agent vapour to saturation, and rejoins the bypass stream. The ratio is precisely controlled so that the dial setting (e.g. 2%) corresponds to the actual percentage of agent in the output gas. Modern vaporisers are temperature- and flow-compensated; desflurane requires a heated, pressurised vaporiser because of its low boiling point (~23°C).

The circle circuit is an elegant solution to gas economy. Exhaled gas passes through an absorber (soda lime), which removes $CO_2$; the cleaned gas then mixes with fresh gas and is rebreathed. This conserves volatile agent, conserves heat and humidity, and reduces operating-room pollution. The trade-off is that

the absorber can become exhausted (colour change indicator) and must be monitored. An exhausted absorber is one cause of an elevated capnography baseline.

> 💡 **Clinical pearl**
>
> The oxygen pressure failure alarm and the hypoxic guard (which prevents you from delivering an $FiO_2$ < 25% by mechanical interlock between $O_2$ and $N_2O$ flowmeters) are the kind of layered, redundant safety engineering that defines the specialty. They exist because at some point in history, patients died from exactly the failures these systems prevent. The Swiss cheese model in section 4 makes this explicit.

### Oxygen delivery devices and $FiO_2$

$FiO_2$ delivered by face devices is highly variable because patients entrain room air with each breath. Standard estimates: nasal cannula at 1–6 L/min ≈ 24–40% (every L/min adds about 4%); simple face mask at 6–10 L/min ≈ 35–55%; non-rebreather mask at 10–15 L/min ≈ 60–80% (rarely 100% even with a good seal); Venturi mask delivers a controlled $FiO_2$ from 24–60% by entraining room air at fixed ratios; high-flow nasal oxygen delivers warmed, humidified gas at flows up to 60 L/min and reliably achieves $FiO_2$ close to 100%, plus small amounts of PEEP. The teaching point is that the device chosen should match both the $FiO_2$ target and the patient's tolerance — a confused, agitated patient won't tolerate a tight non-rebreather mask, but might tolerate HFNO.

# Part 3 · Clinical management

These topics cover what we actually do during a case: control pain, manage fluids, support failing physiology, and use simulation to practise rare emergencies. They are the territory where pre-clerkship physiology meets clinical decision-making.

## 3.1 Pain management — pathways, perception, and multimodal analgesia

**What to teach the student**

- The four steps of pain processing: transduction, transmission, modulation, perception.
- Pain types: nociceptive (somatic and visceral), neuropathic, nociplastic — and why they respond to different drugs.
- Multimodal analgesia: hitting multiple targets simultaneously to spare opioids.
- Equianalgesic dosing and routes (oral vs IV, immediate vs sustained release).
- Regional anaesthesia at a conceptual level — what a peripheral nerve block actually is.
- Chronic pain ≠ acute pain prolonged: the role of central sensitisation.
- Opioid stewardship in 2026: the harms, the indications, the alternatives.

### Teaching review — the four steps of pain, in detail

Pain is one of the most ubiquitous symptoms in medicine and one of the most poorly taught. Most medical schools cover the anatomy of the spinothalamic tract in week three of neuroscience and never return to it. The result is graduates who treat pain by reflex ("add some morphine") rather than by mechanism. The framework that follows — the four-step model of pain processing — is the single most useful conceptual tool an anaesthesiologist can give a student. It also happens to map directly onto the classes of analgesics we use.

*The four-step model of pain processing. Every analgesic strategy targets one or more of these steps.*

#### Step 1 — Transduction

Transduction is the conversion of a noxious stimulus (heat, cold, mechanical pressure, chemical irritation, tissue injury) into electrical activity in a primary afferent neuron. The receptors that do this are called nociceptors. They are free nerve endings of two main fibre types: Aδ fibres, which are myelinated, fast-conducting (5–30 m/s), and responsible for the sharp, localised "first pain" you feel when you stub your toe; and C fibres, which are unmyelinated, slow-conducting (0.5–2 m/s), and responsible for the dull, aching, poorly localised "second pain" that follows.

Tissue injury releases an inflammatory soup that lowers nociceptor thresholds: prostaglandins, bradykinin, histamine, serotonin, substance P, hydrogen ions, ATP, cytokines. This soup explains primary hyperalgesia (the area around an injury becomes more sensitive to pain) and is the target of the most

widely used class of analgesics in medicine — the NSAIDs. NSAIDs inhibit cyclo-oxygenase enzymes and reduce prostaglandin synthesis, which reduces nociceptor sensitisation at the source. Acetaminophen acts centrally rather than peripherally, but is conventionally grouped with the antipyretic analgesics. Topical capsaicin works by depleting substance P from C fibres.

> 💡 **Clinical pearl**
>
> When you teach NSAID pharmacology, frame it as "turning down the volume on the nociceptor" rather than as a generic anti-inflammatory. This connects the drug class to the physiology in a way students remember.

### Step 2 — Transmission

Once the nociceptor fires, the signal is transmitted along the primary afferent into the dorsal horn of the spinal cord. There it synapses (most commonly in laminae I, II, and V) onto second-order neurons that decussate within a segment or two and ascend in the contralateral spinothalamic tract to the thalamus. From the thalamus, third-order neurons project to the somatosensory cortex (where pain is localised), the anterior cingulate (where it acquires its emotional weight), and the insula (where it gains its visceral, "this is bad for me" quality).

Local anaesthetics act on transmission by blocking voltage-gated sodium channels, preventing action potential propagation. This is what every nerve block, epidural, spinal, and field infiltration is doing. The reason a local anaesthetic block is so effective is that it interrupts pain at the wire — no signal gets through, no matter how loud the nociceptor is shouting. The reason it eventually wears off is that the drug diffuses away from the nerve and is cleared by the circulation.

> 💡 **Clinical pearl**
>
> Differential blockade is the elegant teaching point. Smaller, more myelinated fibres are blocked first (sympathetic preganglionic B fibres), then sensory C and Aδ, then motor Aα. This is why a patient with a fresh epidural will feel warm and vasodilated before they lose pinprick, and lose pinprick before they lose motor function. Showing this progression in real time in an awake regional case is unforgettable for a student.

### Step 3 — Modulation

Modulation is what happens at the spinal cord and brainstem to amplify or dampen the pain signal before it reaches consciousness. This is the step that explains why two patients with identical injuries can have radically different pain experiences, and why pain is genuinely modulable rather than purely a stimulus-driven phenomenon.

The dorsal horn has its own micro-circuitry. Inhibitory interneurons release GABA and glycine. Descending pathways from the periaqueductal grey and rostral ventromedial medulla release noradrenaline and serotonin onto dorsal horn neurons, dampening pain transmission. Endogenous opioids — endorphins, enkephalins, dynorphins — act on μ, δ, and κ receptors to inhibit pain at multiple levels. This is the system exploited by every exogenous opioid we give.

The classical "gate control theory" (Melzack and Wall, 1965) describes how non-painful input — touch, vibration, pressure — activates Aβ fibres that, via interneurons in the dorsal horn, inhibit pain transmission from C fibres. This is the physiology behind rubbing your shin after a knock, the analgesic effect of TENS, and parts of acupuncture. Pre-clerks find this delightful because it explains an everyday human experience in mechanistic terms.

Pharmacologically, modulation is targeted by opioids (μ-receptor agonism in the dorsal horn and brainstem), $\alpha_2$-adrenergic agonists (clonidine, dexmedetomidine — mimicking descending noradrenergic inhibition), tricyclic antidepressants and SNRIs (boosting descending serotonergic and noradrenergic tone), and gabapentinoids (modulating presynaptic calcium channels at the dorsal horn, reducing neurotransmitter release). This is also where ketamine has effect — NMDA antagonism in the dorsal horn prevents the central sensitisation that drives wind-up.

### Step 4 — Perception

Perception is the conscious experience of pain — the part the patient actually reports. It is constructed in the cortex from the ascending nociceptive signal modified by attention, expectation, mood, memory, social context, and meaning. This is not soft science: functional imaging reliably shows that the same nociceptive stimulus produces different cortical activation patterns depending on whether the subject expects pain, is distracted, is anxious, is depressed, or has been told the stimulus will be analgesic (placebo).

The clinical implication is that perception is a legitimate analgesic target. Anxiolytics, distraction (music, conversation, virtual reality), preoperative education, cognitive behavioural therapy, hypnosis, and the simple act of explaining what is happening all reduce pain perception. The teaching point for a pre-clerk is that "managing pain" is not just about finding the right drug — it includes everything that shapes the patient's experience.

> 💡 **Clinical pearl**
>
> The placebo effect is real, large (often 30–40% of the analgesic effect of moderate-dose opioids), and ethically usable. The way you talk to a patient about their pain is part of their treatment. This is one of those points pre-clerks rarely hear stated explicitly.

### Putting it together: multimodal analgesia

Multimodal analgesia is the application of the four-step model: hit several steps at once with smaller doses of each, instead of pushing one class to its toxicity limit. A typical post-laparotomy multimodal regimen might combine: paracetamol (mixed central and peripheral), an NSAID if not contraindicated (transduction), regional or neuraxial block (transmission), low-dose ketamine and/or gabapentinoid (modulation, anti-sensitisation), opioid for breakthrough (modulation), and good preoperative communication and post-operative reassurance (perception). Every component has a different mechanism, so the analgesic effects add or even synergise, while the side effects of any single class are kept low.

🩺 **Case: Post-thoracotomy pain**

A 62-year-old has a thoracotomy for lobectomy. What's your multimodal plan? Walk through: thoracic epidural or paravertebral block (transmission), regular paracetamol and NSAID if renal function permits (transduction), ketamine infusion at 0.1–0.2 mg/kg/h intraoperatively and continued at low dose (modulation, prevention of chronic post-thoracotomy syndrome), low-dose patient-controlled opioid for breakthrough, early mobilisation. Each ingredient targets a different point on the pathway. The reason this matters: thoracotomy has one of the highest rates of chronic post-surgical pain (~30–50%), and aggressive multimodal management at the time of surgery is one of the few interventions that reduces it.

### Acute, chronic, and nociplastic pain

An important framing for any student going into clinical years: chronic pain is not just acute pain that has lasted a long time. It is a different physiological state. After persistent nociceptive input, the dorsal horn undergoes structural and functional changes — increased receptor density, NMDA receptor activation, loss of inhibitory interneurons — collectively called central sensitisation. The result is wind-up: the same input now produces an exaggerated response. Allodynia (pain from non-painful stimuli) and hyperalgesia (exaggerated pain from mildly painful stimuli) emerge. The treatment shifts: opioids work less well, anti-neuropathic agents (gabapentinoids, TCAs, SNRIs, ketamine), interventional procedures, and a biopsychosocial approach become more important.

The newer term nociplastic pain (introduced by IASP) describes pain arising from altered nociception without clear evidence of tissue damage or somatosensory lesion — fibromyalgia, irritable bowel syndrome, some chronic low back pain, complex regional pain syndrome. It is not psychological, not malingering, and not nociceptive in the conventional sense. The pathophysiology is central sensitisation as the primary phenomenon. Treating these patients with NSAIDs and opioids alone is futile and often harmful.

⚠️ **Pitfall to flag for the student**

The single biggest teaching point for a pre-clerk going into clinical years: "functional" pain is not fake pain. The nervous system is genuinely processing pain, just from internal rather than external sources. The therapeutic implications follow directly from the four-step model.

## 3.2 Fluid balance and resuscitation

**What to teach the student**

- Body fluid compartments: total body water, intracellular vs extracellular, intravascular vs interstitial.
- Why crystalloids redistribute and only ~25% stays intravascular.
- Balanced crystalloids (Ringer's, Plasma-Lyte) vs 0.9% saline — and the chloride question.
- Static markers (CVP, urine output) vs dynamic markers (pulse pressure variation, passive leg raise).
- Goal-directed therapy: how much, how fast, when to stop.

- Fluid overload is a complication, not a side effect.

### Teaching review — refresh before you teach

Fluid management has shifted dramatically over the last fifteen years. The era of liberal crystalloid resuscitation ("4 litres in the first hour for sepsis") is over. Modern practice is goal-directed: give a defined bolus, assess response, decide whether more is warranted. The teaching point is that "more fluid" is a treatment with side effects (pulmonary oedema, gut oedema with delayed return of function, dilutional coagulopathy, hyperchloraemic metabolic acidosis with saline) and the indication needs to be specific.

Body water is approximately 60% of body weight in adults, distributed two-thirds intracellular and one-third extracellular. Of the extracellular fluid, only about a quarter is intravascular — the rest is interstitial. When you give a litre of crystalloid, only ~250 mL stays in the vasculature; the rest equilibrates within minutes to the interstitium. This is why colloids were once popular (they stay intravascular longer) and why the question "how much should I give?" has no clean answer — it depends on whether you're trying to fill the tank or hydrate the tissues.

> 💡 **Clinical pearl**
>
> Normal saline is neither normal nor saline. It contains 154 mmol/L of sodium and 154 mmol/L of chloride — well above the chloride concentration of plasma (~100 mmol/L). Large-volume saline resuscitation produces a hyperchloraemic metabolic acidosis and may impair renal blood flow. Balanced crystalloids (Ringer's lactate, Plasma-Lyte) more closely approximate plasma electrolyte composition and are now generally preferred for resuscitation, particularly in sepsis (the SMART and PLUS trials).

Assessing volume status remains imperfect. Static markers (CVP, urine output, mucous membranes, capillary refill) are insensitive and non-specific. Dynamic markers — pulse-pressure variation in mechanically ventilated patients, the response to a 250 mL bolus, the passive leg raise as a reversible auto-bolus — are more useful. Point-of-care ultrasound (IVC variability, lung B-lines for pulmonary congestion, cardiac function) is rapidly becoming the standard of care.

## 3.3 Critical care fundamentals

### What to teach the student

- The four classes of shock and their mechanism.
- Vasopressor selection: noradrenaline first-line for most distributive shock.
- Basic ventilator settings: tidal volume, respiratory rate, PEEP, $FiO_2$, mode.
- Lung-protective ventilation: 6–8 mL/kg ideal body weight, plateau pressure < 30 $cmH_2O$.
- Sedation in ICU: shorter is better, daily interruption, awareness of delirium.
- The deteriorating ward patient: NEWS score, MET/RRT activation, early escalation.

### Teaching review — refresh before you teach

*Four categories of shock, defined by which part of the oxygen-delivery chain has failed. Treatment follows mechanism.*

Shock is a clinical state in which oxygen delivery is inadequate for tissue demand. The four traditional categories — hypovolaemic, cardiogenic, distributive, obstructive — each break a different link in the $DO_2$ chain you taught in section 1.1. Recognising which is which is the first task in resuscitation; treatment follows mechanism. Distributive shock from sepsis (low SVR) needs vasopressors and source control, not litres of fluid into already-leaky vessels. Cardiogenic shock (failing pump) needs inotropes and may need mechanical support, not aggressive fluid. Obstructive shock (PE, tamponade, tension pneumothorax) needs the obstruction relieved, not pharmacology.

Mechanical ventilation has its own conceptual layer worth a brief tour. Volume-controlled modes deliver a set tidal volume; the resulting pressure depends on lung compliance. Pressure-controlled modes deliver a set pressure; the resulting tidal volume depends on compliance. PEEP (positive end-expiratory pressure) keeps alveoli open at end-expiration, improves V/Q matching and oxygenation, but reduces venous return and may worsen haemodynamics in a hypovolaemic patient. Lung-protective ventilation — 6–8 mL/kg ideal body weight, plateau pressure under 30 $cmH_2O$, modest PEEP — was originally proven to reduce mortality in ARDS and is now standard for most ventilated patients.

> 💡 **Clinical pearl**
>
> The single most important thing for a pre-clerk to take away from a critical-care discussion: the ventilator and the vasopressor are temporary supports while you fix the underlying problem. They are not treatments in themselves. Good ICU care is about identifying and addressing the cause.

## 3.4 Introduction to simulation-based learning

> **What to teach the student**
>
> - Why simulation is now central to medical education: rare events, deliberate practice, psychological safety.
> - What "psychological safety" means in a sim — and why it matters.
> - How to receive feedback without becoming defensive.
> - Debriefing models: PEARLS, advocacy-inquiry, the "plus-delta" approach.
> - The difference between technical and non-technical skills (the latter often matter more).

### Teaching review — refresh before you teach

Simulation has gone from a curiosity twenty years ago to a core part of anaesthesia, emergency, and critical-care training. The rationale is straightforward: rare emergencies are too rare to learn by experience alone, and deliberate practice in a safe environment produces durable skill change. The harder lesson, especially for pre-clerks, is that simulation is most valuable for the non-technical skills — leadership,

communication, situational awareness, decision-making under pressure — which are exactly the skills that get the least attention in the rest of the curriculum.

Psychological safety is a precondition for productive simulation. The participants need to know that mistakes will be examined non-judgementally, that the goal is learning rather than evaluation, and that what is said in the debrief stays in the debrief. This is unfamiliar territory for medical students who have spent two years being graded constantly. Naming it explicitly at the start of any simulation experience — "this is a learning environment, not a test" — changes how they engage.

> 💡 **Clinical pearl**
>
> The advocacy-inquiry approach to debriefing is worth teaching explicitly. Instead of "that was wrong" (judgement) or "why did you do that?" (often heard as accusatory), use "I noticed you went straight to intubation — I was thinking BMV first; I'm curious about your reasoning." The pattern is: state your observation, share your interpretation, ask for theirs. It surfaces mental models without triggering defence.

# Part 4 · Decision-making and patient safety

If applied physiology is the science of anaesthesia, decision-making and patient safety are its art. These topics rarely appear in the pre-clerkship curriculum and are exactly the territory where anaesthesiologists have the most to offer. Each section is shorter on bullet points and longer on conceptual prose, because the topics genuinely benefit from explanation.

## 4.1 High-stakes decision-making under uncertainty

**What to teach the student**

- Type 1 thinking (fast, automatic, pattern-recognition) vs Type 2 (slow, deliberate, analytical).
- Recognition-primed decision-making in experts.
- How time pressure changes the cognitive task.
- When to commit to a plan and when to keep gathering information.
- The role of explicit verbal commitment ("I think this is X, here's why") in surfacing reasoning.

### Teaching review — refresh before you teach

Most clinical decisions, even high-stakes ones, are not made by working through a textbook differential. Experts mostly use Type 1 thinking: pattern-recognition based on accumulated experience. The book by Daniel Kahneman (Thinking, Fast and Slow) and the work of Gary Klein on naturalistic decision-making are good background reading. The clinical implication is that pattern-recognition is fast and usually right, but it fails predictably in unfamiliar situations and under stress.

Recognition-primed decision-making — described by Klein in firefighters, ER physicians, and military commanders — proceeds something like: the situation matches a familiar pattern; an action plan is recalled; the plan is mentally simulated to check it works; if so, it is executed; if not, it is modified. This is fundamentally different from the formal decision analysis taught in evidence-based medicine, but it is what experts actually do, and teaching pre-clerks to value this mode (rather than dismiss it as "intuition") is important.

Where Type 1 fails, Type 2 has to take over. The skill of senior clinicians is recognising when their pattern-match is failing and explicitly switching modes. Cues: the case feels "off" without a clear reason; the patient is not responding as expected to standard treatment; the differential is unusually broad; you are getting tired or distracted. In these moments, the right move is to slow down, write out a differential, and go through it deliberately.

**💡 Clinical pearl**

A teachable habit: before committing to a plan, articulate it out loud — "I think this is X because of Y, and the plan is Z" — and watch the team's reaction. The act of verbalising forces some Type 2 review, and the team's reaction (nods, hesitation, raised eyebrows) is information. This is also the foundation of cognitive aids and

checklists — they force the surfacing of reasoning that would otherwise stay implicit.

## 4.2 Cognitive heuristics and bias

**What to teach the student**

- Heuristics are useful shortcuts that fail predictably under specific conditions.
- Anchoring: latching onto the first piece of information.
- Confirmation bias: seeking and weighting evidence that confirms an initial hypothesis.
- Availability bias: recent or vivid cases come to mind first.
- Premature closure: settling on a diagnosis before considering alternatives.
- Fixation error: persisting with a wrong plan because it has been mentally committed to.

### Teaching review — refresh before you teach

The literature on diagnostic error is fundamentally a literature on cognitive bias. Around 10–15% of diagnoses in routine practice are wrong, and a substantial proportion of those errors are not from missing information but from misweighted information. Teaching pre-clerks to name the biases they are about to make is one of the highest-yield interventions in medical education.

Anchoring is the tendency to rely too heavily on the first piece of information encountered. The handover says "this is a chest pain, probably MI" and the receiver structures their workup around that hypothesis even when the clinical picture diverges. The remedy is the explicit pause: "if I were seeing this patient cold, what would I think?"

Confirmation bias is the tendency to seek information that supports an existing hypothesis and discount information that contradicts it. The remedy is the active search for disconfirming evidence: "what would I see if I'm wrong?" In differential diagnosis, this means explicitly asking "what doesn't fit?" rather than building the case for the leading hypothesis.

Availability bias is the tendency for recent, vivid, or memorable cases to come to mind first. The trainee who saw a missed PE last week starts diagnosing PE in everyone with a vague chest complaint. The remedy is awareness — naming the recent case and explicitly broadening the differential.

Premature closure is the most common diagnostic error in retrospective analyses. A reasonable hypothesis is formed early, the case is mentally closed, and subsequent information is fitted to the diagnosis rather than allowed to revise it. Fixation error is the procedural cousin: the team persists with a plan that isn't working because they have committed to it. Both are addressed by deliberate "stop and reassess" moments — "are we sure this is what's happening?"

**🩺 Case: The classic CICO fixation**

A textbook teaching case: an elective surgical patient becomes a cannot-intubate, cannot-oxygenate

emergency. The team makes multiple attempts at laryngoscopy, then at LMA placement, while saturations fall. Hands stay on the airway; nobody calls for surgical airway. Twenty minutes into the crisis, the patient arrests. The cognitive error is fixation: each new attempt feels like progress, the team is committed to the airway-from-above approach, and the unfamiliar action (front-of-neck access) is repeatedly deferred. The teaching point is the explicit time-out: "if I haven't oxygenated in three attempts, the next move is plan D." Cognitive aids and decision algorithms exist precisely to interrupt fixation.

## 4.3 Situational awareness and clinical workflow

### What to teach the student

- Endsley's three levels of situational awareness: perception, comprehension, projection.
- Task saturation and how it degrades situational awareness.
- The "10-second pause" as a deliberate awareness-restoration tool.
- Why team situational awareness matters more than individual awareness in complex environments.

### Teaching review — refresh before you teach

Situational awareness, as defined by Mica Endsley in the human-factors literature, is the perception of elements in the environment, the comprehension of their meaning, and the projection of their status into the near future. Most adverse events involve breakdowns at one of these levels. The case file usually shows that information was available, but not perceived; or perceived but not understood; or understood but not extrapolated to the patient's near-future state.

Task saturation is the predictable enemy of situational awareness. As cognitive load rises, the field of attention narrows; we attend to the task in front of us and stop monitoring the periphery. In the OR, this is when the anaesthesiologist absorbed in a difficult cannulation misses the falling end-tidal $CO_2$. The countermeasure is structural: when task-saturated, delegate one of your simultaneous tasks. The single most important thing a senior trainee can learn is to call for help before they need it.

### 💡 Clinical pearl

The 10-second pause: in any prolonged or evolving case, deliberately stop, look at all the monitors, and ask yourself "what is this patient doing?" — not "what am I doing?" Senior anaesthesiologists do this several times per case, often without realising it. Teaching it explicitly accelerates trainee development.

## 4.4 Team dynamics and crisis resource management

### What to teach the student

- Closed-loop communication: state, repeat, confirm.
- Role clarity: who is the team leader, and who is doing what.
- Calling for help early — the Goldilocks problem of escalation.

- Flat hierarchy in emergencies: anyone with relevant information speaks up.
- Graded assertiveness: CUS — Concerned, Uncomfortable, Safety issue.

### Teaching review — refresh before you teach

Crisis resource management began in aviation in the 1970s after a series of crashes traced back not to mechanical failures but to crew communication breakdowns. It was imported into anaesthesia by David Gaba and colleagues in the 1990s and is now woven through anaesthesia training internationally. The core insight is simple: in complex high-stakes environments, technical skill is necessary but not sufficient. Teams fail at communication, role allocation, leadership, and assertiveness more often than at procedure.

Closed-loop communication is the basic unit. The leader directs ("give 1 mg adrenaline"), the receiver repeats back ("1 mg adrenaline"), and confirms when done ("adrenaline 1 mg given, time 14:23"). Sounds trivial. In the absence of explicit closed loops, things drop. Teaching this to a pre-clerk through a sim and showing the difference in error rates is one of the most striking demonstrations in medical education.

Graded assertiveness — using a structured way to escalate concern — addresses one of the deepest problems in medical hierarchy: the junior who notices something off but doesn't speak up. The CUS framework gives them a graded vocabulary: "I'm concerned about…" → "I'm uncomfortable with this plan…" → "This is a safety issue, we need to stop." The advantage of teaching the words explicitly is that the junior can deploy them without having to invent the script under stress.

> 💡 **Clinical pearl**
>
> Two-challenge rule: if you raise a concern and it isn't acknowledged, raise it again, more explicitly. If still not acknowledged, escalate up the hierarchy or invoke a stop. This is borrowed directly from aviation and is now standard in many surgical safety programmes.

## 4.5 Patient safety and the Swiss cheese model

### What to teach the student

- Latent vs active errors: most adverse events involve both.
- The Swiss cheese model of accident causation.
- Just culture: addressing system failures without scapegoating individuals.
- M&M as a learning tool, not a blame ritual.
- Checklists: why something so simple has such large effects.

### Teaching review — refresh before you teach

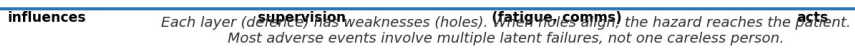


*Reason's Swiss cheese model. Each defensive layer has weaknesses (holes); when they align, hazards reach the patient. Most adverse events involve multiple latent failures.*

James Reason's Swiss cheese model is the foundational framework for thinking about patient safety. The picture is intuitive: each layer of defence (training, equipment design, supervision, checklists, monitoring) has weaknesses (holes). On any given day, hazards bounce off the layers and don't reach the patient. Occasionally, the holes align — a tired clinician, a poorly designed drug ampoule, an unfamiliar room layout, a missed double-check — and the hazard reaches the patient. The model's deep point is that adverse events almost never have a single cause. Asking "who messed up?" is the wrong question. Asking "how did the layers align?" is the right one.

The latent vs active error distinction follows. Active errors are the proximal cause — the wrong drug given, the wrong site marked. Latent errors are the upstream conditions that made the active error possible — labels that look alike, fatigue from a 16-hour shift, a culture in which the junior didn't feel able to question the consultant. A just culture is one that examines latent errors openly and only sanctions individuals for genuine reckless behaviour, not for honest mistakes within a flawed system.

The WHO Surgical Safety Checklist (2008) is the most-studied example of a low-cost, high-impact safety intervention. Implementation reduced surgical mortality by ~40% in the original study. Its mechanism is not complicated — it forces explicit verbal confirmation of patient identity, procedure, allergies, antibiotic prophylaxis, anticipated blood loss, and equipment availability — but it works precisely because it converts implicit assumptions into explicit checks. The teaching point: simple structural interventions can outperform complex clinical ones.

> ⚠️ **Pitfall to flag for the student**
>
> M&M (morbidity and mortality) conferences vary enormously in quality. Done well, they are non-blaming, system-focused, learning-oriented, and shared across the team. Done badly, they become public shaming sessions that drive errors underground and worsen safety. Modelling a healthy M&M culture for trainees is one of the most lasting things you can do.

## 4.6 Ethics in acute and perioperative care

> **What to teach the student**
>
> - The four pillars: autonomy, beneficence, non-maleficence, justice.
> - Capacity: a decision-specific assessment, not a global label.
> - Consent in the emergency setting: implied, presumed, and the role of substitute decision-makers.
> - DNR/DNAR in the OR: suspended vs honoured — the perioperative nuance.
> - Goals of care: not the same as DNR; broader, harder, more important.
> - Futility: a contested concept best avoided as a unilateral declaration.

### Teaching review — refresh before you teach

Ethics in anaesthesia tends to surface in time-pressured situations: a Jehovah's Witness presents in haemorrhagic shock; a frail elderly patient with established DNR comes for hip fixation; a young patient

with a head injury cannot consent; a family demands "everything" for a patient with no realistic chance of recovery. These are not classroom dilemmas — they are Tuesday afternoon.

Capacity is the most useful ethical concept to teach a pre-clerk because it is frequently misused. Capacity is decision-specific (a patient may have capacity to refuse one treatment but not another), context-specific, and time-specific (capacity can fluctuate). The standard test is: does the patient understand the relevant information; can they retain it; can they weigh it; can they communicate a choice? Disagreement with the medical team is not, on its own, evidence of incapacity.

Perioperative DNR is one of those quiet topics that catches students by surprise. A patient with a community DNR who comes for surgery has typically agreed to anaesthetic interventions that look very much like resuscitation — intubation, mechanical ventilation, vasopressors, even direct cardiac massage in extremis. The ethical and practical question is: is the DNR suspended for the perioperative period, honoured strictly, or modified to a procedure-specific agreement? The answer is reached through explicit conversation with the patient (or their substitute decision-maker) before induction, ideally documented. Doing this conversation well is a master skill.

> **🩺 Case: The Jehovah's Witness with massive haemorrhage**
>
> A 34-year-old practising Jehovah's Witness presents to your trauma bay with a ruptured ectopic pregnancy. She is hypotensive, Hb 60, refusing blood. She has full capacity, has documented her wishes, and has a husband who supports the refusal. What's your plan? Walk the student through: respecting the autonomous refusal of blood, while pursuing every alternative (cell salvage, fluid resuscitation, tranexamic acid, fibrinogen concentrate, surgical haemostasis, factor concentrates if accepted). The key teaching point is that autonomy is binding even when the consequences are severe — and that an ethical framework gives you a structure for an emotionally extreme situation.

# Part 5 · Professional identity formation

The OR is one of the most concentrated environments in medicine for observing how a team actually functions and how individual clinicians develop their sense of professional self. These last two topics often get squeezed in between the more technical material, but they are arguably the most important things a pre-clerk takes away from a rotation.

## 5.1 Interprofessional collaboration

**What to teach the student**

- Who is in the OR and what they actually do: anaesthesia, surgery, nursing, anaesthesia assistants, perfusionists, technicians.
- Why role clarity changes by phase: induction, surgery, emergence, recovery.
- Handovers: SBAR, I-PASS, and why structure matters.
- Conflict in the OR: when it happens, why, and how teams repair after it.

### Teaching review — refresh before you teach

*Use for any handoff, escalation, or consult call. Removes ambiguity.*

*SBAR: a structured handover format. Use for any transfer of clinical responsibility — between OR and recovery, between shifts, or between the floor and the ICU.*

Pre-clerks rarely see how interprofessional teams actually work; they have been told the parts but not how they fit. The OR is a teaching gift here. Within ten minutes of being in the room they can see who decides what, how communication flows, where the friction points are, and how repair happens after a tense moment. Naming the roles explicitly as you teach helps: "the anaesthesia assistant just opened the difficult-airway trolley because of the way I said 'have a Bougie ready' — that's how we communicate concern without alarming the patient."

Handovers are where most patient-safety incidents originate. Information dropped at handover is information that gets reinvented or lost downstream. SBAR (Situation, Background, Assessment, Recommendation) is the most widely adopted structure; I-PASS (Illness severity, Patient summary, Action items, Situation awareness, Synthesis) is more comprehensive and now standard in many academic centres. The teaching point is that the structure does work the speaker would otherwise have to do under pressure: it forces them to assemble the information into a usable form.

## 5.2 Formulating and prioritising management plans

**What to teach the student**

- ABCDE: treat first what kills first.
- Parallel vs serial thinking — and when each is appropriate.

- The reassessment loop: every intervention is a new data point.
- Plan articulation: speaking the plan out loud as a team-coordination tool.

## Teaching review — refresh before you teach

*ABCDE primary survey. The order is fixed because of physiological priority: airway compromise kills before breathing failure, breathing failure before circulatory collapse, and so on.*

ABCDE is one of the most successfully exported teaching frameworks in medicine. Its genius is not the content — anyone could list those five — but the discipline of order. Airway is addressed before breathing, and breathing before circulation, because of a strict physiological hierarchy: an obstructed airway kills in three minutes, inadequate breathing in six to ten, circulatory collapse over a longer window. The framework forces the user to address what kills first, even when other things look more dramatic.

Parallel vs serial thinking is a more advanced concept worth at least naming. In a single-operator situation, you must think serially: do A, then B, then C. In a team, work can be parallelised: while you secure the airway, someone else gets venous access; while you reassess, someone else draws up the next drug. Senior anaesthesiologists become very good at delegating in parallel, which is partly why a good team can do in two minutes what a solo operator does in ten.

The reassessment loop closes the framework. Every intervention is a new data point: did the BP rise after the bolus, did the saturation improve after the airway adjustment, did the patient settle after the dose? If yes, continue; if no, revise. The most common error in teaching cases is acting without reassessing — "give more fluid" without checking whether the first fluid did anything.

> 💡 **Clinical pearl**
>
> Speaking the plan out loud is a coordination tool, not just a teaching habit. "Plan: I'm going to bag-mask, you draw up another 50 of propofol, we'll attempt direct laryngoscopy in 30 seconds, and if I'm not in by attempt two we move to LMA." This sounds excessive in print; in real time it aligns the team in five seconds and reduces the cognitive load on every member. Pre-clerks rarely see this articulated and find it transformative when they do.

# Part 6 · Brain-wave monitoring and depth of anaesthesia

Of all the monitors in the modern OR, the depth-of-anaesthesia monitor is the one most likely to leave a pre-clerk wondering "what does that number actually mean?" It is also one of the most beautiful applications of physiology — the EEG, generated by the synchronised activity of cortical neurons, changes in characteristic ways as a patient transitions from awake to anaesthetised, and those changes can be quantified, displayed, and acted on. This section is intentionally detailed because the underlying physiology is rich and rarely taught well.

## 6.1 What the EEG is showing you

### What to teach the student

- The EEG records the summed electrical activity of millions of cortical neurons.
- It's analysed in frequency bands — gamma, beta, alpha, theta, delta — each with characteristic clinical meaning.
- Higher frequency = more cortical activity (typically); lower frequency = less.
- Anaesthesia produces a stereotyped progression of EEG changes that can be quantified.
- BIS, Entropy, and other monitors are processed-EEG numbers derived from these changes.

### Teaching review — refresh before you teach

The electroencephalogram measures the postsynaptic potentials of pyramidal neurons in the cerebral cortex, summed across millions of cells, and recorded as voltage fluctuations on the scalp. The raw signal is small (microvolts) and highly oscillatory. To make sense of it clinically, we decompose it into frequency bands using Fourier analysis. Each band corresponds to a different physiological state, and the relative power in each band shifts predictably under anaesthesia.

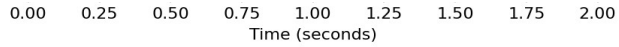


*The five EEG frequency bands. Each is generated by a different neural rhythm and corresponds to a different state of consciousness. Anaesthesia shifts the dominant frequency progressively from gamma/beta toward delta.*

### The five bands, in clinical context

**Gamma (>30 Hz) —** the highest-frequency activity, generated by tightly synchronised inhibitory networks. Gamma is the rhythm of active cognition: focused attention, sensory binding, working memory. Anaesthesia rapidly suppresses gamma, which is why patients lose the ability to integrate information almost immediately on induction.

**Beta (13–30 Hz) —** the EEG of an awake, alert, problem-solving brain. Anxious patients often show prominent beta. Beta also appears paradoxically during the early phases of induction with GABAergic agents (propofol, benzodiazepines) — the so-called "paradoxical excitation," a transient burst of fast

frontal activity before slowing takes over. This is worth pointing out because students will see BIS values briefly rise on induction with propofol and may misinterpret it.

**Alpha (8–12 Hz) —** the rhythm of relaxed wakefulness, classically seen with eyes closed. Under propofol anaesthesia at surgical depths, a striking change occurs: alpha activity becomes coherent across the frontal cortex ("frontal alpha") at high amplitude. This is one of the EEG hallmarks of adequate hypnotic depth with propofol and is what BIS algorithms are partly detecting.

**Theta (4–7 Hz) —** the rhythm of drowsiness, light sleep, deep meditation, and surgical anaesthesia. Theta activity emerges as patients deepen from light sedation toward surgical levels. In children, theta is a more prominent component of normal awake EEG than in adults — one of the reasons paediatric BIS values are less reliable as direct indicators of anaesthetic depth.

**Delta (0.5–4 Hz) —** the slowest rhythm, characteristic of deep sleep, deep general anaesthesia, and pathological states (coma, severe encephalopathy). Delta dominance during anaesthesia indicates a deeply suppressed cortex. As anaesthesia deepens further, the EEG transitions through delta-dominance into burst suppression (alternating periods of activity and electrical silence) and ultimately to a flat trace — too deep for almost any clinical purpose.

## 6.2 Processed EEG monitors: BIS and friends

**What to teach the student**

- BIS (bispectral index) is a number from 0 to 100, derived from a proprietary algorithm.
- 0 = no electrical activity (flat EEG, deep coma); 100 = fully awake.
- Surgical anaesthesia target: typically 40–60.
- BIS does not directly measure consciousness — it measures EEG features that correlate with hypnotic depth.
- It does not work equally well for all anaesthetic agents (poor for ketamine, dexmedetomidine, $N_2O$).
- It is one input among many — clinical signs still matter.

### Teaching review — refresh before you teach

*The BIS scale. Numbers reflect a processed combination of frequency content, burst suppression ratio, and inter-frequency synchrony. The 40–60 target band corresponds to surgical anaesthesia for most GABAergic agents.*

BIS was the first commercially successful processed-EEG depth monitor and remains the most widely used. The algorithm is proprietary, but the inputs are well-described: relative power in the high-frequency bands (suggesting wakefulness), the burst-suppression ratio (suggesting deep anaesthesia), and a measure of inter-frequency synchrony (the bispectral component, which is sensitive to the coherent frontal alpha that emerges at surgical depths). The output is a single number between 0 and 100 that corresponds, in studies on healthy adults receiving GABAergic anaesthesia, to a graded continuum from awake to flat EEG.

It is essential to teach what BIS does not do. It does not measure consciousness directly — there is no "consciousness meter." It measures EEG features that correlate with hypnotic depth in healthy adults receiving the agents on which it was calibrated (mainly propofol and volatiles). For ketamine, BIS values can be paradoxically high because ketamine produces high-frequency activity even at deep dissociative depths. For dexmedetomidine, BIS underestimates depth because the EEG resembles natural sleep. For nitrous oxide, BIS is essentially insensitive.

> 💡 **Clinical pearl**
>
> The B-Aware and B-Unaware trials, taken together, suggest BIS-guided anaesthesia reduces the (already very low) incidence of intraoperative awareness in high-risk patients (e.g. cardiac surgery, trauma, paralysed and TIVA without volatile end-tidal monitoring). It does not clearly reduce awareness across all patient populations, and it is not a substitute for clinical vigilance and end-tidal agent monitoring when those are available.

Newer monitors — Entropy, SedLine, Narcotrend — use different algorithms but operate on the same principle: extract features from the EEG that change predictably under anaesthesia, and present them as a number or pattern the clinician can act on. SedLine in particular displays the raw EEG and the density spectral array (DSA), which is a colour-coded display of frequency power over time. The DSA lets the clinician see the actual EEG patterns (frontal alpha emergence, burst suppression onset) rather than relying on a derived index. For teaching, the DSA is the most pedagogically useful display because it makes the underlying physiology visible.

## 6.3 Clinical applications and limitations

**What to teach the student**

- Avoiding intraoperative awareness in high-risk patients.
- Avoiding excessively deep anaesthesia, which is independently associated with worse outcomes in elderly and frail patients.
- Detecting burst suppression (typically at BIS < 40), which is increasingly seen as a marker of anaesthetic over-dose.
- Why elderly patients need lower drug doses to reach the same BIS as young patients.
- Why BIS is unreliable in children under 1, in patients on ketamine or dexmedetomidine, and during cardiopulmonary bypass.

### Teaching review — refresh before you teach

The most important clinical use of depth monitoring in 2026 is avoiding excessively deep anaesthesia, particularly in elderly and frail patients. Multiple observational studies and a growing body of randomised data link prolonged time in burst suppression to postoperative delirium, prolonged hospital stay, and possibly long-term cognitive decline. The mechanism is debated — whether burst suppression causes harm or marks a vulnerable brain — but the practical implication is clear: keep older patients out of burst suppression unless there is a specific indication.

Age-related dose-response shifts are dramatic. An 80-year-old reaches BIS 45 on roughly half the propofol that a 25-year-old needs. The MAC of volatile agents drops about 6% per decade after age 40. A pre-clerk who has only seen healthy young patients anaesthetised may not appreciate how different the dose-response curve is in geriatric patients, and BIS-guided dosing makes this visible in a way the syringe pump cannot.

> ⚠️ **Pitfall to flag for the student**
>
> BIS values can be artefactually altered by EMG (the patient lightening and tensing facial muscles), pacemaker artefact, electrocautery, warming blankets, and certain surgical instruments. A sudden BIS change without a clinical correlate should prompt a check of the raw EEG before acting — never blindly increase the propofol because the BIS rose.

> 🩺 **Case: The frail elderly patient drifting into burst suppression**
>
> A 78-year-old, BMI 22, having a hemicolectomy. You're running 0.7% sevoflurane and an opioid infusion. BIS reads 28. You look at the DSA: the periods of activity are getting shorter, periods of silence longer — burst suppression. What do you do? Reduce the volatile, reassess the opioid, and consider whether you've been targeting BIS 40–60 in someone whose brain is already saying "this is plenty." Walking through this with a student and showing them the burst-suppression pattern is one of the most physiologically vivid moments in clinical teaching.

# Appendix A · One-line take-home for each topic

If you teach a topic in a hurry and want a closing line for the student, use these:

- Cardiopulmonary physiology: "Treat the lever that broke — preload, afterload, contractility, rate, or rhythm."
- Pharmacokinetics: "How long a drug lasts depends on how long you've been giving it."
- Pharmacology: "Anaesthesia is an outcome reached through several mechanisms — GABA, NMDA, opioid, others."
- CPR: "Blood flow during arrest is whatever your hands are producing — minimise interruptions."
- Airway: "Bag-mask ventilation buys you five minutes; learn it before you learn to intubate."
- IV access: "Start distal; the cognitive part is harder than the motor part."
- Ultrasound: "Once the needle is in the skin, your eyes are on the screen, not your hands."
- Monitors: "Every monitor reading is an interpretation, not a fact — know how each one lies."
- Anaesthesia machine: "The vaporiser is magic; everything else is plumbing and safety engineering."
- Pain: "Hit several steps in the pathway at once. That's what multimodal means."
- Fluids: "More fluid is a treatment with side effects. Have an indication."
- Critical care: "The ventilator and the vasopressor are temporary — fix the underlying cause."
- Decision-making: "Recognise when your pattern-match is failing and slow down."
- Cognitive bias: "Always ask: what would I see if I were wrong?"
- Situational awareness: "Take a 10-second pause and ask what the patient is doing — not what you're doing."
- CRM: "Closed-loop communication, graded assertiveness, call for help early."
- Patient safety: "Most adverse events have many causes. The right question is 'how did the holes align?'"
- Ethics: "Capacity is decision-specific, not a global label."
- Teamwork: "Speaking the plan out loud aligns everyone in five seconds."
- ABCDE: "Treat first what kills first; reassess after every intervention."
- Brain-wave monitoring: "BIS doesn't measure consciousness — it measures EEG features that correlate with depth."

# Appendix B · A teaching session in 10 minutes

When a registrar or rotating student appears at the head of the bed and you have a single case to work with, here's a teaching skeleton that fits any topic:

- **Minute 1 — Anchor:** Point at something they can see right now (the monitor, the syringe, the airway). Ask one open question.
- **Minutes 2–4 — Concept:** Give the framework in plain language. One diagram, one sentence per layer.
- **Minutes 5–7 — Application:** Tie it to what's happening in this case. "This is why we just gave X."
- **Minute 8 — Pitfall or pearl:** Name one common error or surprising fact related to the topic.
- **Minute 9 — Reflection:** Ask what they would do differently if they were running the case.
- **Minute 10 — Take-home:** One sentence. Make sure they can repeat it back.

# Closing

Teaching in the OR is one of the underappreciated privileges of our specialty. The students you teach today are the colleagues you will rely on a decade from now — in your hospital, on your retrieval calls, on the other end of a phone at 3 a.m. when something is going wrong. The time you spend explaining why we just did what we did is rarely counted in any productivity metric, but it compounds. This manual is here to make that time a little easier.

Add to it. Cross out what doesn't work. Write in the margins. The best teaching manuals are the ones that look used.

*— end —*