BASIRA

CANCER METABOLISM
RESEARCH & TRAINING

بصيرة

TRAINING CURRICULUM

CHAPTER THREE

The Science

What cancer is, and how a cell powers and builds itself

The biology, from the beginning. No prior molecular biology is assumed — only that you are willing to hold several new ideas at once. By the end you should be able to look at a tissue section, a gene name and a metabolic pathway and know what each is telling you.

Founder & Programme Lead	Aroob Gomosani
Prepared & delivered by	Haneen Marghalani
Version	2.0 · second cohort
Date	2026

Gomosani, A., & Marghalani, H. (2026). Basira Curriculum. Zenodo.
https://doi.org/10.5281/zenodo.21791751

3·0

Before you begin

This is the longest chapter, and the only one that is mostly facts rather than judgement. That is deliberate: you cannot generate a good scientific question in a field you do not know. Chapter 1 gave you the machinery of thinking. This chapter gives you something to think about.

A word on scope. Cancer is not one disease, and this chapter does not treat it as one. But teaching every cancer shallowly teaches nothing, so we do something better: we follow one cancer — colorectal — properly, from a normal crypt to an invasive tumour with its own evolving sub-populations, and then use other cancers as deliberate contrasts. Learn one well, then generalise.

WHERE THIS SITS

	Chapter	What it answers
	1 · Thinking	How does science reason, and how do I avoid fooling myself?
	2 · Practice	How do we behave — integrity, collaboration, communication.
▶️	3 · The Science	What is cancer, and how does a cell power itself?
	4 · Data & Tools	What data exists, and how do we handle it without breaking it?
	5 · The Toolkit	The statistics and code, as a reference you return to.
	6 · Build Your Own	Now go find a question of your own and run it.

WHAT YOU SHOULD BE ABLE TO DO BY THE END

- Trace the path from DNA to protein, and say what an expression measurement and a dependency measurement each capture.
- Distinguish an oncogene from a tumour suppressor by mechanism, mutation type and example — and explain why it matters.
- Describe cancer as evolution within a tissue, and say what follows from that for treatment.
- Look at a stained tissue section and identify epithelium, stroma, immune cells and a tumour.
- Explain why a tumour sample is a mixture of cell types, and what that does to any measurement made on it.
- Name the core metabolic pathways, where each happens, and why the mitochondrion matters so much.
- State the Warburg effect and, more importantly, say what it does not mean.

- Read a piece of cancer-metabolism literature and identify where an untested assumption is hiding.

3.1
The cell, and what it has to do

Every cell in your body carries the same instruction set, and yet a neuron and a colon lining cell could not be more different. That paradox has a simple resolution: cells do not use all of their instructions at once. They read different parts. Almost everything in this chapter follows from that one fact.

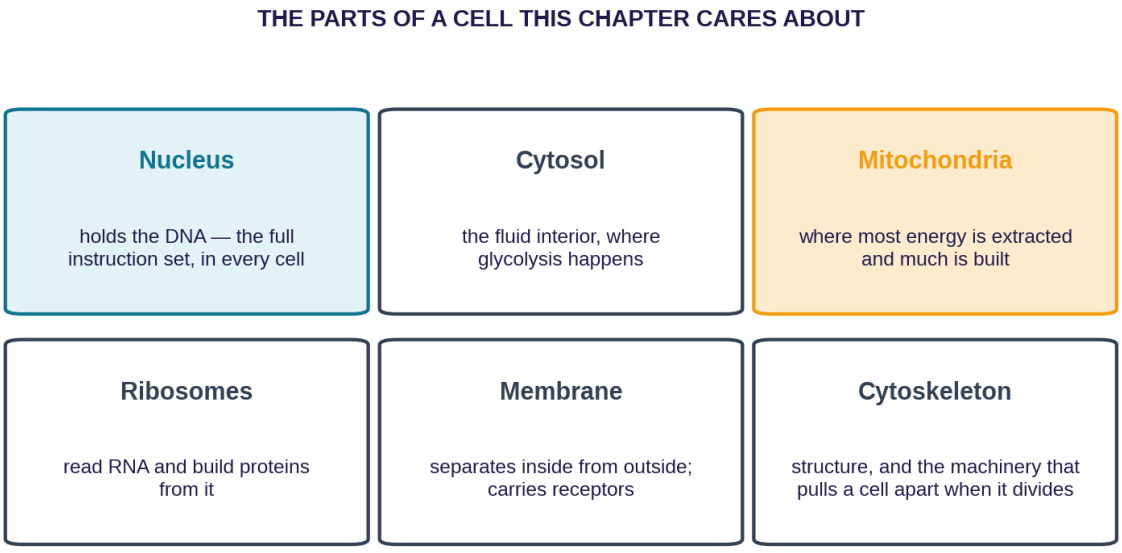

Figure 3.1 The parts that matter for this chapter. There are many more; these are the ones that recur.

3.2
The instructions: DNA, RNA, protein

A gene is a stretch of DNA. To use it, the cell first transcribes it into messenger RNA — a working copy of that one gene — and then translates that RNA into a protein, the molecular machine that actually does something: catalysing a reaction, forming a structure, carrying a signal. This flow is called the central dogma, and knowing it tells you exactly what our two kinds of measurement are measuring.

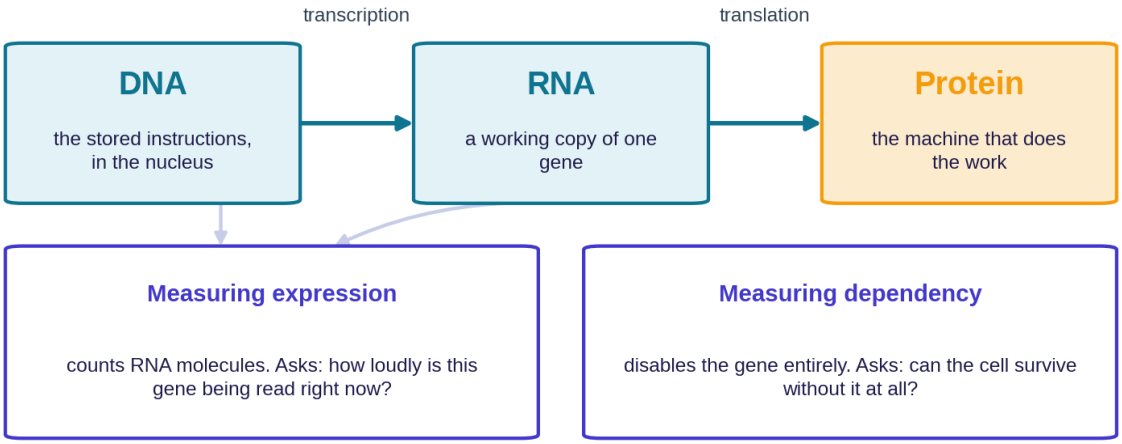

Figure 3.2 The flow, and the two very different questions you can ask about any gene along it.

◆ KEY IDEA — **Two questions that sound similar and are not**

Expression asks how loudly a gene is being read. Dependency asks what happens if the cell cannot read it at all.

A gene can be read very loudly and be entirely dispensable. A gene can be read quietly and be absolutely essential. Keeping these apart is one of the most useful habits you will build here.

3.3 How the instructions get damaged

Cancer begins with changes to the DNA itself. The ones that concern us are somatic — acquired during a person's life, in the tissue, rather than inherited from a parent. They come in several forms, and the form matters.

FIVE WAYS THE INSTRUCTIONS GET DAMAGED

Missense	one letter changes, so one amino acid changes	sometimes harmless, sometimes switches the protein permanently on
Nonsense	one letter change inserts a premature STOP	the protein is truncated and usually broken
Frameshift	letters added or removed, not in multiples of three	everything downstream is garbled — almost always destroys the protein
Splice-site	the joins between coding segments are disrupted	the RNA is assembled wrongly; the protein usually fails
Copy number	the whole gene is amplified to many copies, or deleted	turns the dose sharply up or sharply down

The unifying question is direction: does the change BREAK the protein, or switch it permanently ON?

Figure 3.3 Not all mutations are equal. What matters is what each does to the protein at the far end.

3.4
Accelerators and brakes

Cancer-driving genes fall into two categories that are, in a real sense, mirror images of one another. Think of a cell's growth as a car. Two quite different failures can make it race out of control: jamming the accelerator down, or cutting the brakes.

	Oncogene	Tumour suppressor
	the accelerator, jammed down	the brakes, cut
Normal job	drives growth when growth is appropriate	restrains growth, repairs damage, triggers cell death
What goes wrong	GAIN of function — switched permanently on	LOSS of function — broken or removed
How	recurrent activating “hotspot” mutations, or amplification	nonsense, frameshift or splice-disrupting mutations, or deletion
Copies affected	one is enough	usually both — the “two-hit” rule
Examples	KRAS · EGFR · BRAF · PIK3CA	TP53 · APC · PTEN · RB1

Figure 3.4 The distinction that governs how a driver gene is detected in data. Getting it wrong makes an analysis fail silently.

The asymmetry in the last row is worth pausing on. A single jammed accelerator is enough to speed a car, so one mutated copy of an oncogene suffices. But a car has redundant braking, so both copies of a tumour suppressor usually have to be disabled before the brake fails — the classic two-hit requirement.

> ▲ **CAUTION — Why this becomes a line of code later**
>
> Because the two classes are altered by different kinds of mutation, you cannot detect them the same way. An activating hotspot mutation in an oncogene does not appear in a catalogue of damaging loss-of-function mutations — because it is not a loss of function.
>
> Screen an oncogene against the loss-of-function list and you find zero mutants, with no error and no warning, and conclude wrongly that there is nothing there. The biology decides the lookup table.

> ? **CHECKPOINT — Checkpoint 1**
>
> - Why can two cells with identical DNA be completely different from each other? (§3.1)
> - A gene is transcribed at very high levels in a tumour. Does that tell you the tumour needs it? Explain. (§3.2)
> - Which mutation types tend to break a protein, and which tend to switch one permanently on? (§3.3)
> - TP53 is a tumour suppressor and KRAS is an oncogene. For each, say what kind of mutation you would look for, and what happens if you look for the wrong kind. (§3.4)

3.5 What cancer actually is

At its core, cancer is uncontrolled cell proliferation arising from damage to the genome. A cell acquires mutations that release it from the normal restraints on growth. Its descendants inherit those mutations and acquire further ones, and among them, whichever variants grow fastest or survive best come to dominate.

That is evolution by natural selection, playing out inside a tissue, over years. Mutation supplies the variation; the tumour's environment does the selecting. It is also why cancer is so hard to defeat — you are not fighting a fixed target but a population under selection, and treatment is itself a selective pressure.

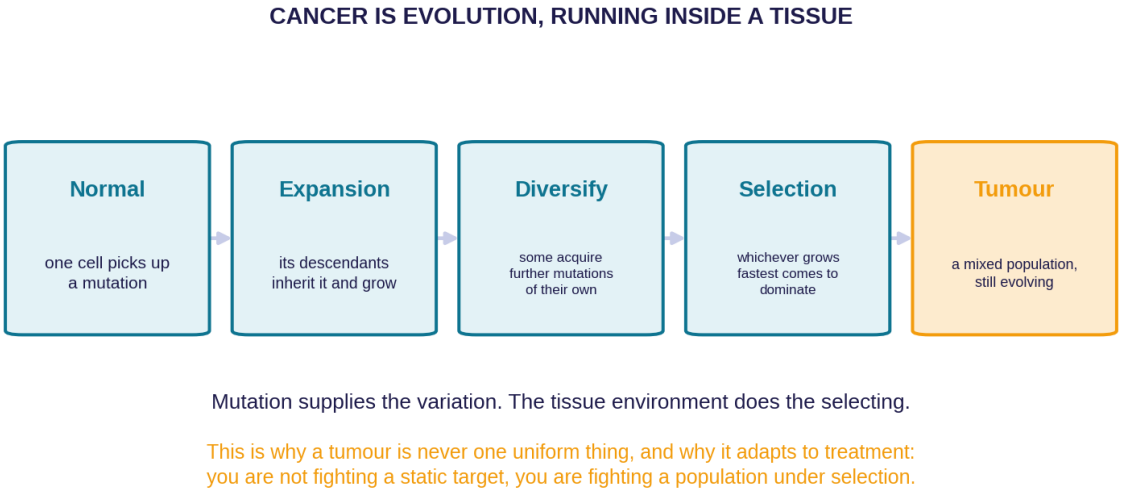

Figure 3.5 The process, in five steps. Note that the endpoint is a mixed population, not a uniform one.

This is not a theoretical picture. It can be seen directly in a resected specimen.

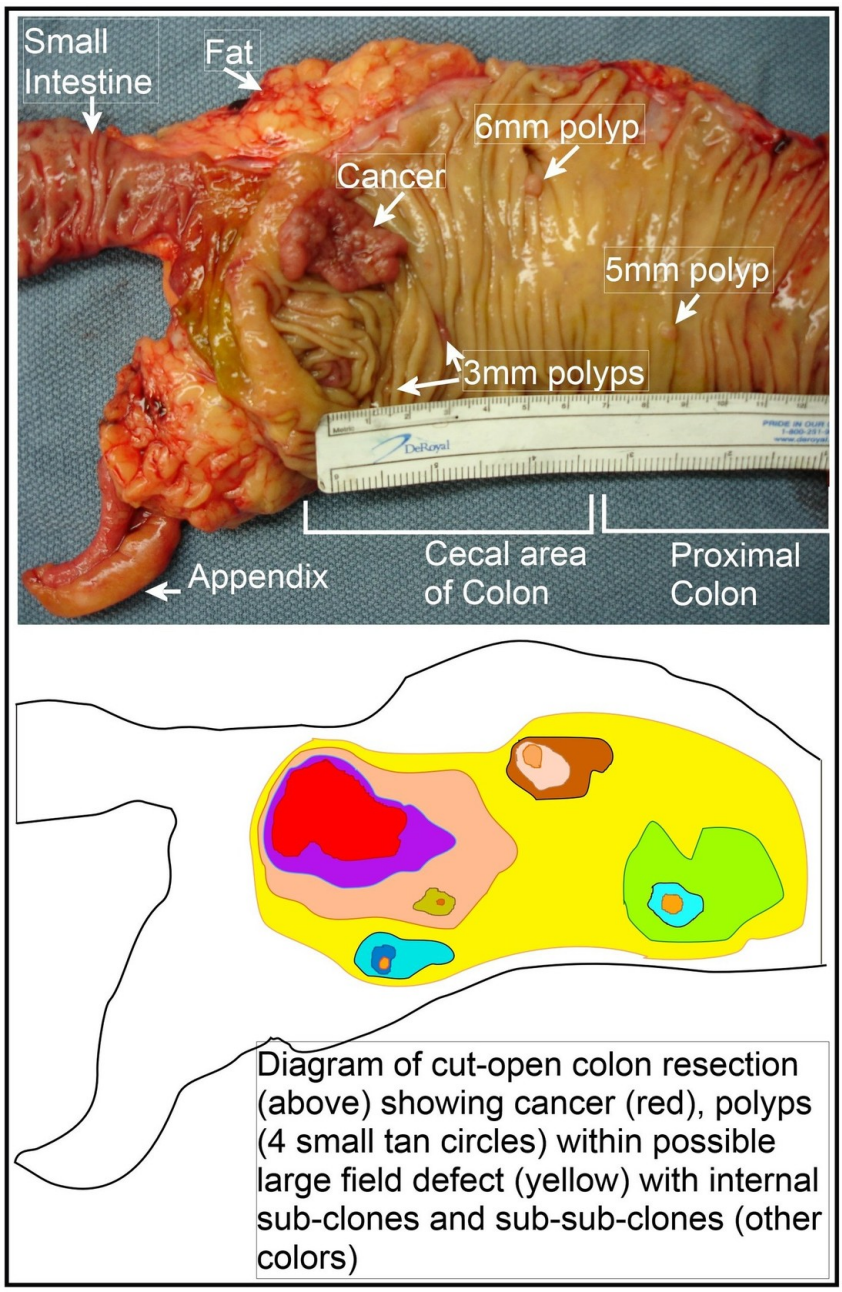

Figure 3.6 A resected colon segment: an invasive cancer alongside four separate polyps, with the schematic below mapping a field defect and the sub-clones and sub-sub-clones within it. One patient, one region of tissue, many related but distinct populations.

Source: Wikimedia Commons — see file page for author and licence.

■ DEFINITION — **Driver and passenger mutations**

A tumour genome accumulates many mutations, but only a few — the drivers — actually confer a growth advantage. The rest are passengers, carried along because they happened to be in a cell that succeeded for other reasons.

Telling them apart is one of the central tasks of cancer genetics, and one of the easiest places to fool yourself: a mutation being frequent is not the same as a mutation mattering.

3.6

Cancer is many diseases

“Cancer” is a category, not a disease. A leukaemia and a glioblastoma share a logic — uncontrolled proliferation driven by genomic damage — but very little else. They arise in different tissues, are driven by different genes, behave differently and are treated differently.

DIFFERENT CANCERS, DIFFERENT DRIVERS — AND ONE THAT IS EVERYWHERE

		and also…
Colorectal	APC · KRAS · SMAD4	TP53
Lung adenocarcinoma	KRAS · EGFR · STK11	TP53
Breast	PIK3CA · GATA3 · ERBB2	TP53
Pancreatic	KRAS · CDKN2A · SMAD4	TP53
Melanoma	BRAF · NRAS · CDKN2A	TP53
Glioblastoma	EGFR · PTEN · IDH1	TP53

Most drivers are tissue-specific: the same organ, the same handful of genes. A few, like TP53, are broken almost everywhere.

Figure 3.7 Each tissue has its own characteristic drivers. A few genes, TP53 above all, are broken almost everywhere.

The variety runs deeper than the organ. Within a single cancer of a single organ, the tissue can look strikingly different from patient to patient.

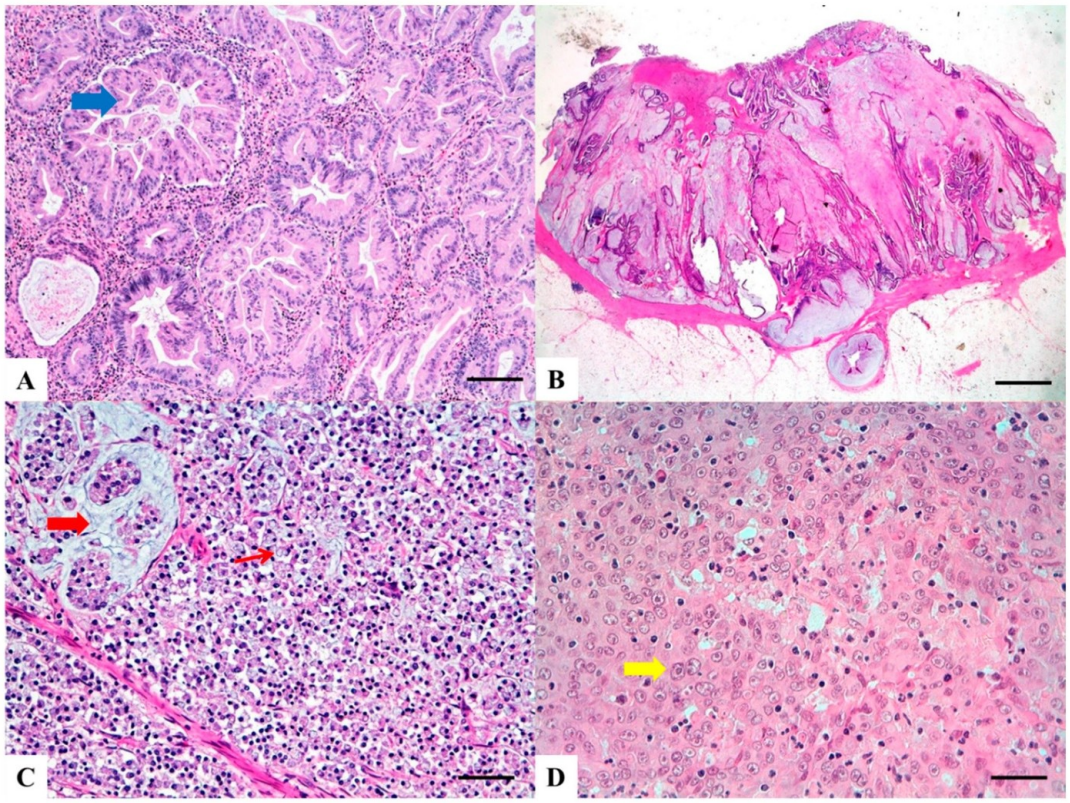

Figure 3.8 Four histological subtypes of colorectal carcinoma: serrated, mucinous, signet-ring and medullary. All four are colorectal cancer. All four came from the same organ.

Source: Wikimedia Commons — see file page for author and licence.

◆ **KEY IDEA — The practical consequence for your own work**

A finding that holds across several cancer types is more credible than one confined to a single type — and a finding confined to one type may still be important, provided you say so.

The failure to avoid is quietly generalising: testing in one tumour type and writing as though you had discovered something about cancer.

3.7 The hallmarks

However a tumour arises, it tends to acquire a common set of capabilities. These are usually called the hallmarks of cancer, and they are a useful checklist for asking what a given tumour has managed to do.

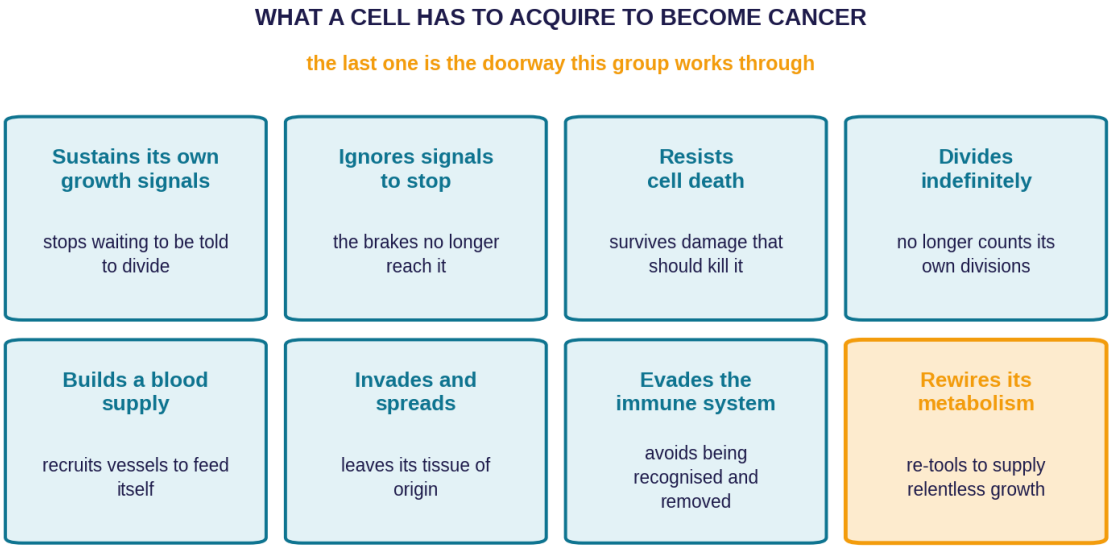

Figure 3.9 Eight capabilities a cell must acquire. The last is where this group's work begins.

The reprogramming of metabolism is not a side effect of cancer. It is a recognised, defining capability — which is why the second half of this chapter is about how a cell powers and builds itself.

3.8
Reading a tissue section

Much of what we know about a tumour begins with someone looking at a thin slice of it under a microscope. You will not be asked to diagnose anything, but you should be able to look at a section and see what is there. Start with normal.

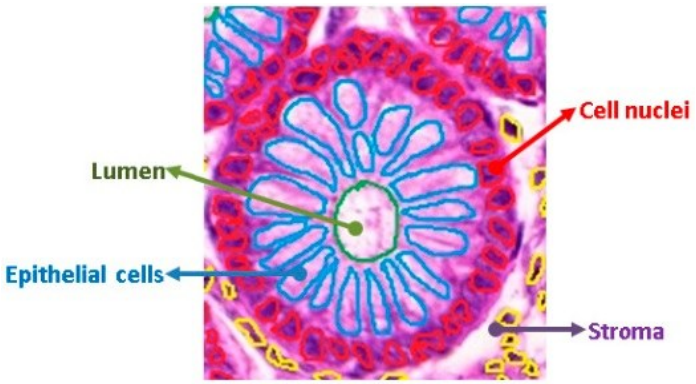

Figure 3.10 Normal large-intestinal crypts. The epithelial cells form an orderly ring around a central lumen, their nuclei neatly arranged at the outer edge, with stroma between the crypts.
Source: Wikimedia Commons — see file page for author and licence.

Note the order. The cells are one layer thick, uniformly sized, with their nuclei lined up in a row. The architecture is repetitive and predictable. Now compare a cancer arising in the same tissue.

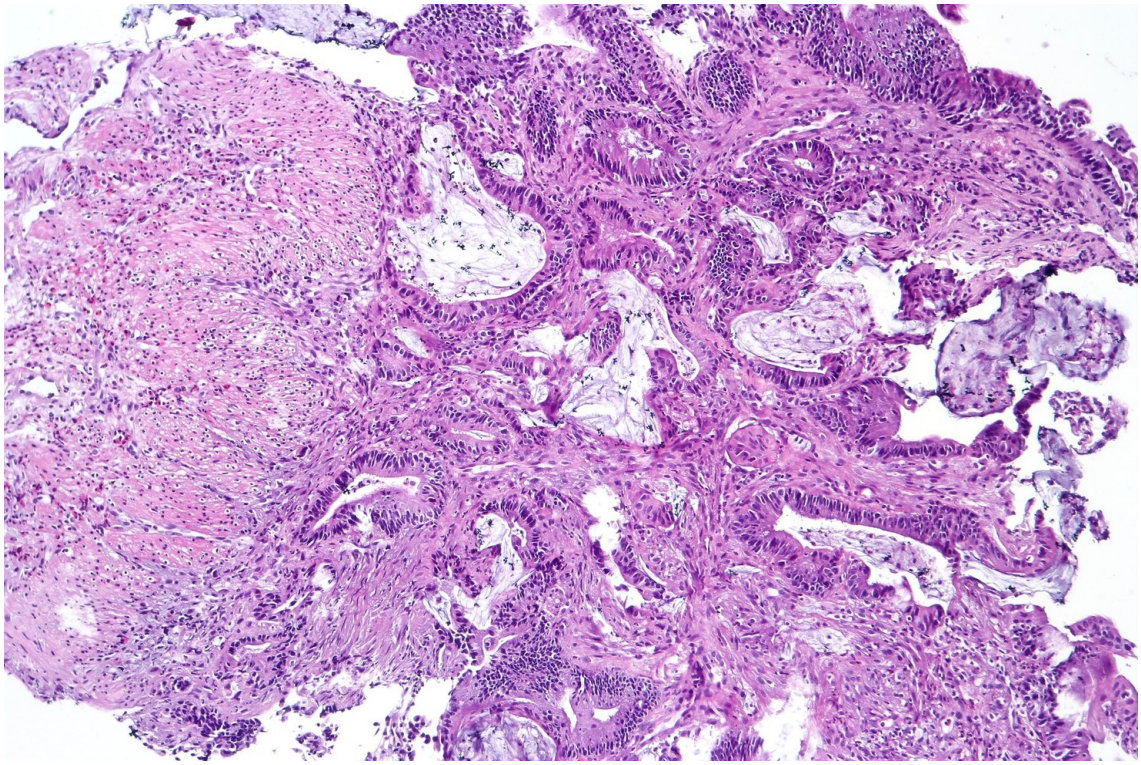

Figure 3.11 Colorectal adenocarcinoma. The orderly crypt architecture is gone — glands are irregular, crowded and fused, nuclei are enlarged, darker and no longer aligned, and the whole structure invades the surrounding tissue rather than respecting its boundary.

Source: Wikimedia Commons — see file page for author and licence.

You do not need training to see the difference. That loss of order is what a pathologist is reading, and it is the visible consequence of everything in the first half of this chapter.

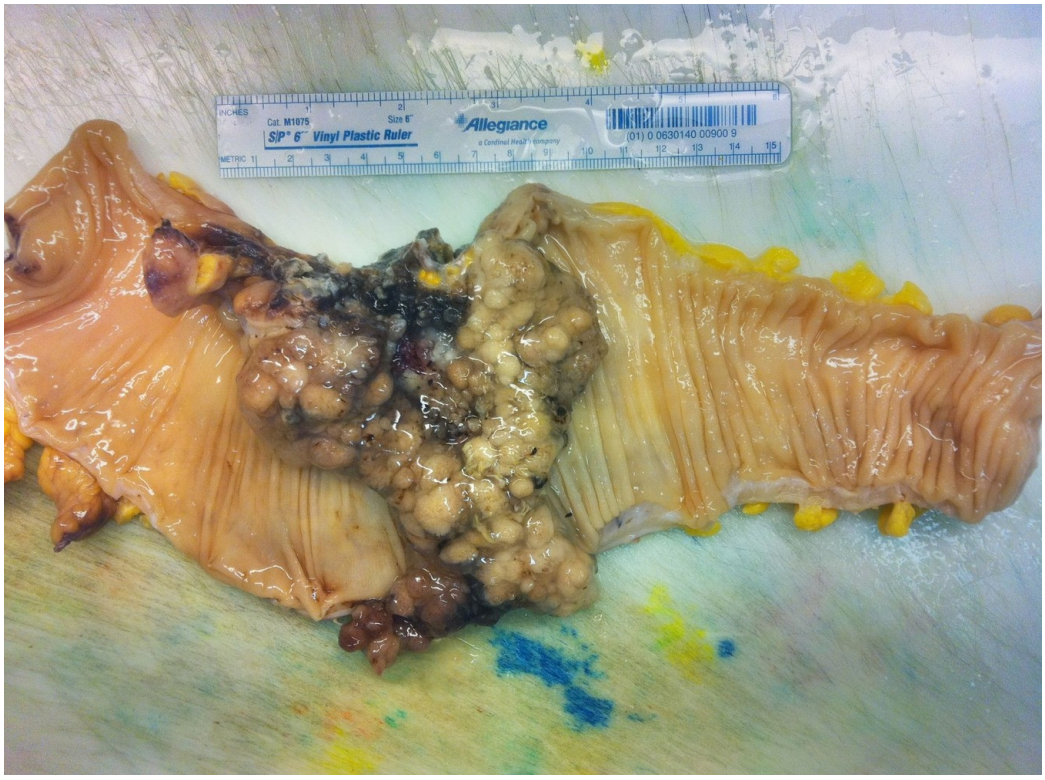

Figure 3.12 The same disease at a scale you can see without a microscope: a large colonic cancer arising within an adenoma, in an opened resection specimen. The benign polyp and the invasive cancer are continuous with one another.

Source: Wikimedia Commons — see file page for author and licence.

3.9
A tumour is a tissue, not a bag of cancer cells

Here is a fact that is easy to overlook and that governs a great deal of how data must be analysed. A tumour is not a pure population of malignant cells. It is a whole tissue — an ecosystem.

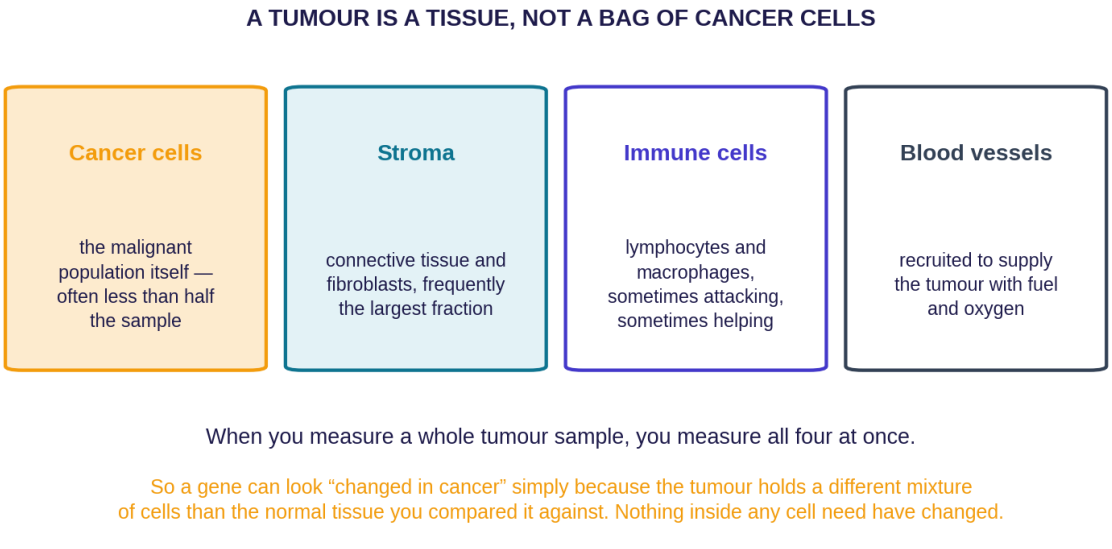

Figure 3.13 Four populations, measured together whenever you measure a tumour sample.

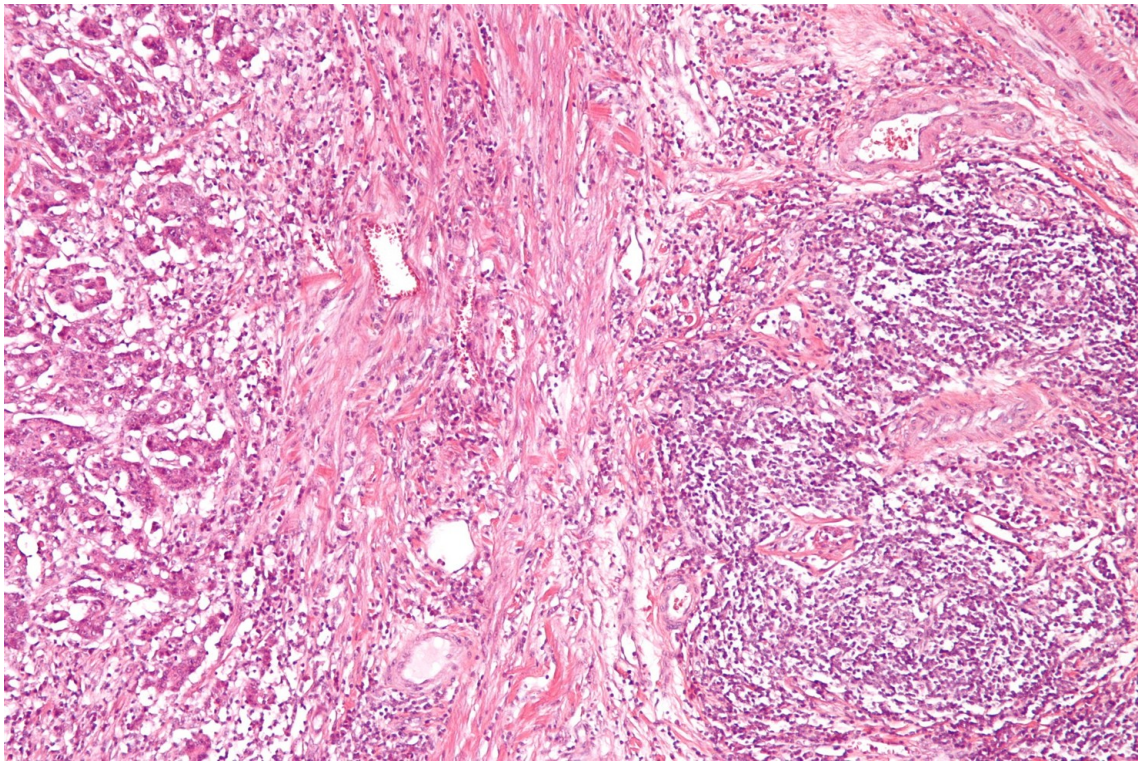

Figure 3.14 The ecosystem, visible in one field. Tumour on the left, a dense band of immune cells on the right, and stroma running between them. A measurement made on this whole section averages all three.

Source: Wikimedia Commons — see file page for author and licence.

▲ **CAUTION — The composition problem, stated once so it is unmistakable**

Tumour tissue and normal tissue do not merely differ in what their cells are doing. They differ in which cells are present at all.

So a gene can appear “changed in cancer” purely because the tumour contains more epithelium and less fat, with nothing having changed inside any individual cell. Any honest analysis has to separate a real change within cells from a change in the proportions of cells — and Chapter 5 spends much of its length on exactly that.

? **CHECKPOINT — Checkpoint 2**

- Explain in your own words why cancer is described as evolution, and name one practical consequence for treatment. (§3.5)
- What is the difference between a driver and a passenger mutation, and why is frequency alone a poor guide? (§3.5)
- Looking at Figures 3.10 and 3.11, list three specific differences you can see between normal and malignant tissue. (§3.8)
- A gene shows higher expression in tumour samples than normal samples. Give two entirely different explanations, only one of which is about cancer biology. (§3.9)

3.10

Metabolism: how a cell powers and builds itself

Metabolism is the sum of the chemical reactions by which a cell obtains energy and materials, and uses them to stay alive, function, and — crucially for cancer — divide. It has two complementary halves: catabolism breaks large molecules down to release energy, and anabolism spends that energy building the large molecules a cell needs. The energy released is captured in a universal currency, ATP.

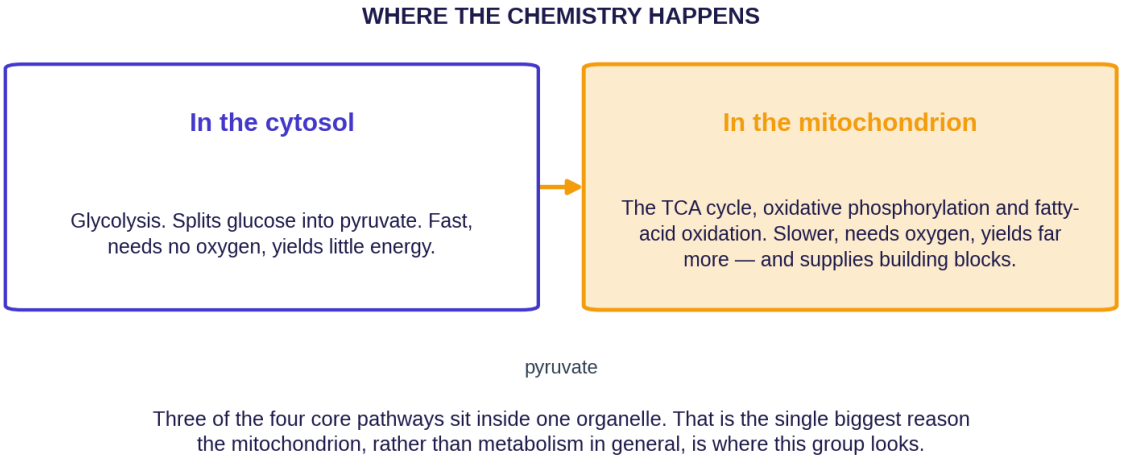

Figure 3.15 Location matters, and it is why this group's attention is where it is.

3.11
The core pathways

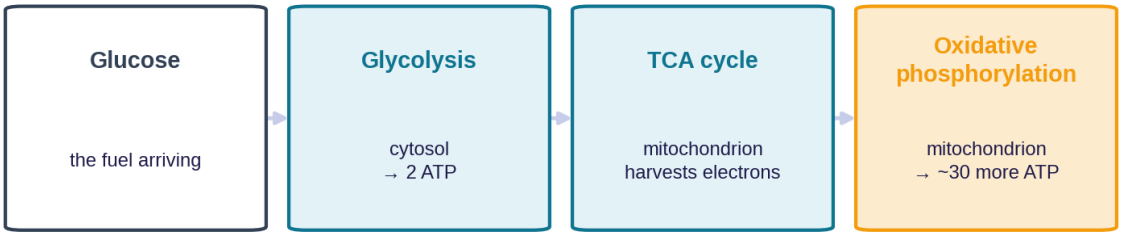

The mitochondrion produces roughly fifteen times more ATP than glycolysis alone.

But energy is only half the story. The TCA cycle is also where a growing cell draws the raw material for new nucleotides, amino acids and lipids — which is why a dividing cell needs it so badly.

Figure 3.16 Glucose in, ATP out — but the ATP is only half of what the cell is extracting.

3.12

The mitochondrion

The mitochondrion is popularly called the cell's powerhouse, and it is — but that phrase badly undersells it. Beyond generating most of the cell's ATP, mitochondria are hubs for biosynthesis, for managing the cell's redox balance, for signalling, and for controlling programmed cell death. They are woven into whether a cell lives, grows or dies, which is precisely why cancer exploits them and why therapy might target them.

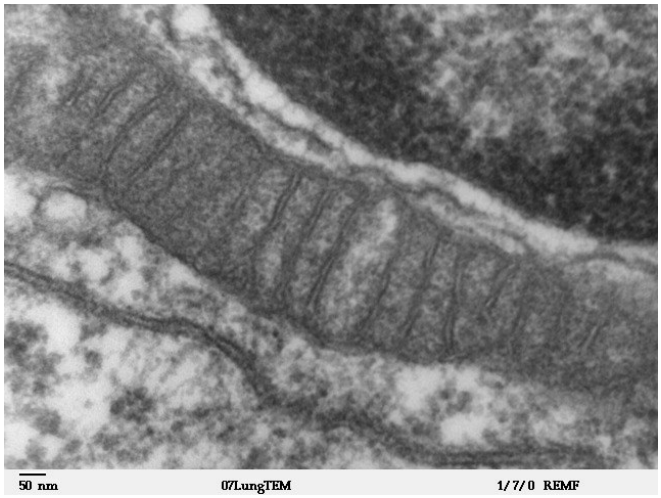

Figure 3.17 A mitochondrion in mammalian lung tissue, imaged by transmission electron microscopy. The double membrane and the folded cristae are directly visible. Scale bar 50 nm.

Dartmouth College Electron Microscope Facility, via Wikimedia Commons — public domain.

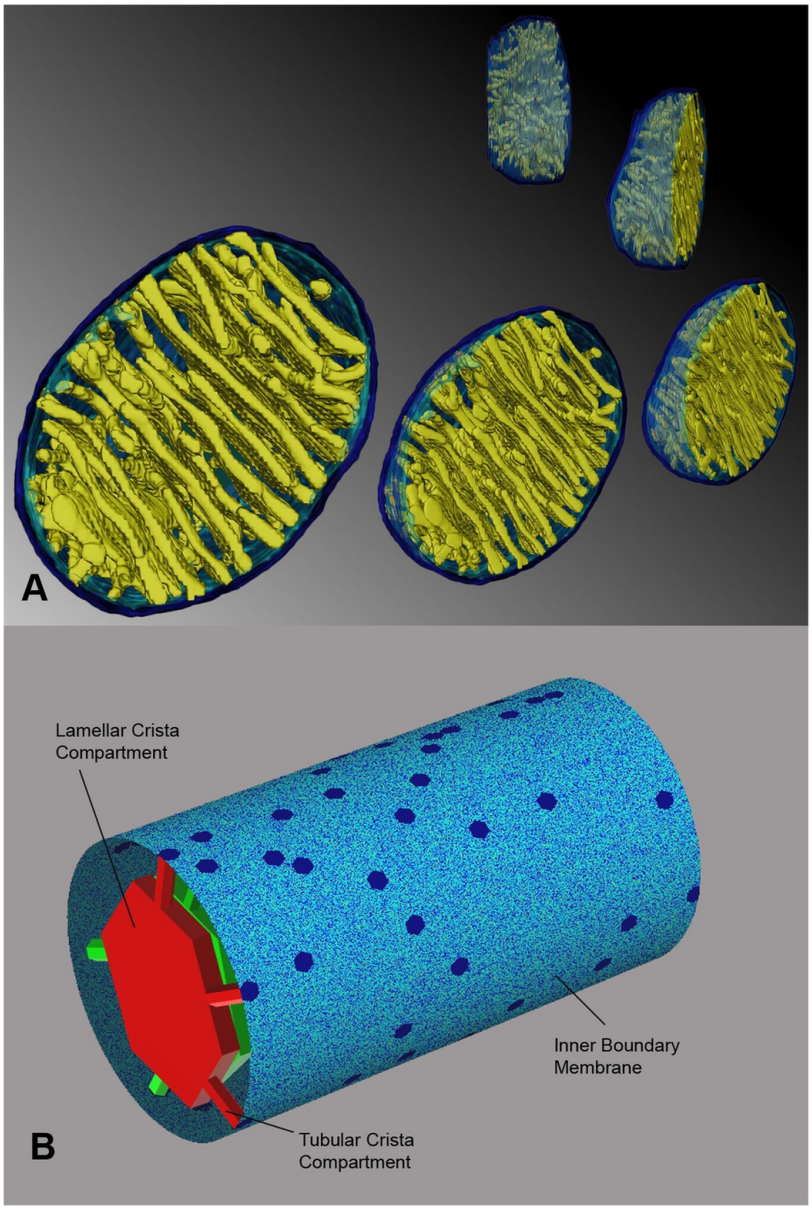

Figure 3.18 The same structure reconstructed in three dimensions from tomography, with cristae in yellow and the boundary membranes in blue. The lower panel names the compartments.

Source: Wikimedia Commons, CC BY 3.0 — see file page for author.

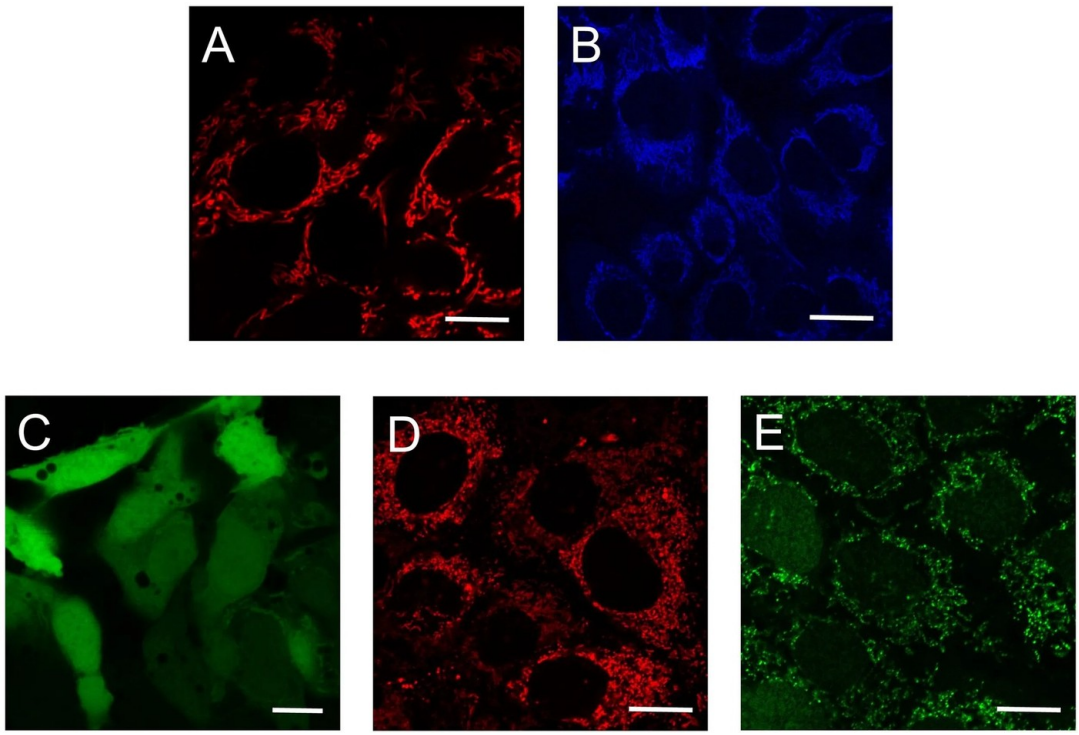

Figure 3.19 And in living cells: mitochondria are not isolated beans but a dynamic, branching network that divides, fuses and moves.

Frontiers in Oncology 13:1152553 (2023), CC BY 4.0.

A genetic curiosity makes mitochondria especially interesting to study, and it matters for how we search them.

THE MITOCHONDRION RUNS ON TWO GENOMES

13 proteins

encoded by the mitochondrion's own small genome, a relic of its ancient bacterial ancestry

~1,100 proteins

encoded in the nucleus, built in the cytosol, then imported into the mitochondrion

Catalogued in a resource called MitoCarta.

Almost everything that builds and runs a mitochondrion is written in the nuclear genome — which is why a nuclear mutation can wreck mitochondrial function.

Figure 3.20 Thirteen proteins from one genome, roughly eleven hundred from the other.

3.13
Why a dividing cell is different

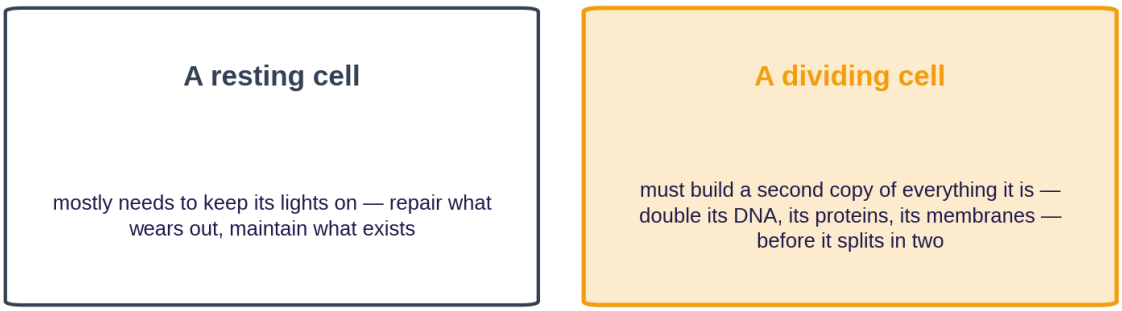

Figure 3.21 The reframe that makes cancer metabolism make sense.

This distinction — between metabolism for energy and metabolism for building — is central to understanding cancer, because cancer is above all a disease of inappropriate cell division.

3.14
The Warburg effect

The best-known feature of cancer metabolism is also the most misunderstood. In the 1920s Otto Warburg observed that many tumours consume glucose voraciously and secrete lactate even when oxygen is plentiful — apparently choosing an inefficient route when an efficient one is available.

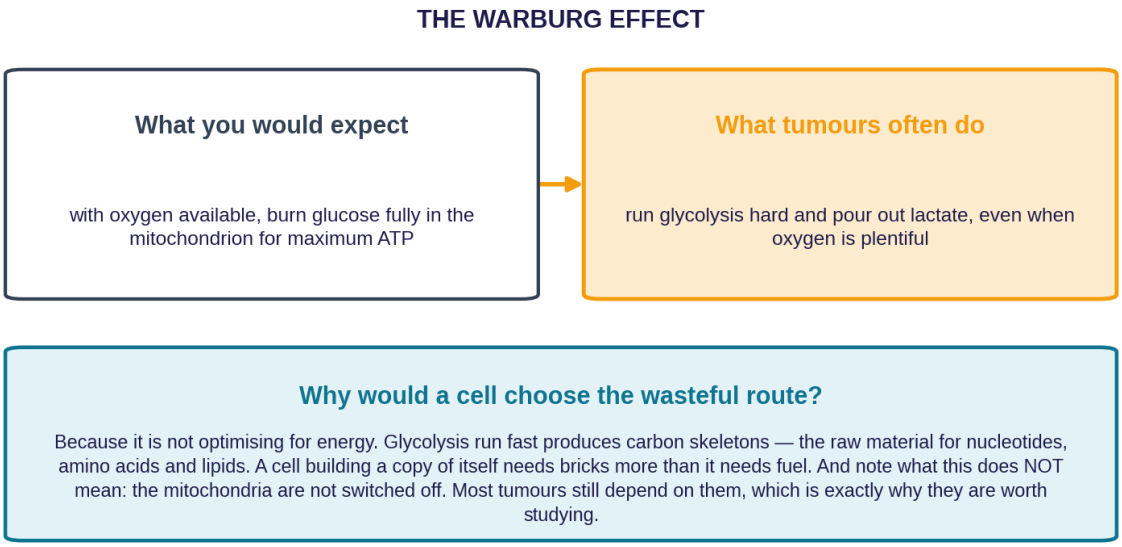

Figure 3.22 The observation, and the resolution of its apparent paradox.

▲ CAUTION — **What the Warburg effect does not mean**

It does not mean cancer cells have switched off their mitochondria. This is a genuinely common misreading, and it has misled real research programmes.

Mitochondrial metabolism remains essential in most tumours, supplying the building blocks of proliferation. That is exactly why mitochondrial genes are a promising place to look for dependencies a cancer cell cannot do without.

? CHECKPOINT — Checkpoint 3

- Name the core metabolic pathways and say where in the cell each one happens. (§3.11)
- Give three jobs the mitochondrion does besides making ATP. (§3.12)
- Roughly 1,100 mitochondrial proteins are encoded in the nucleus and 13 in the mitochondrion itself. Why does that matter when hunting for cancer dependencies? (§3.12)
- Explain the Warburg effect to someone who has just told you “cancer cells don't use their mitochondria.” (§3.14)

3.15
Finding a question of your own

You now know enough biology to start looking for a question. But a field this size is not a blank page — it is crowded, and most of the obvious questions have been asked. The skill is in seeing where the genuine gaps are.

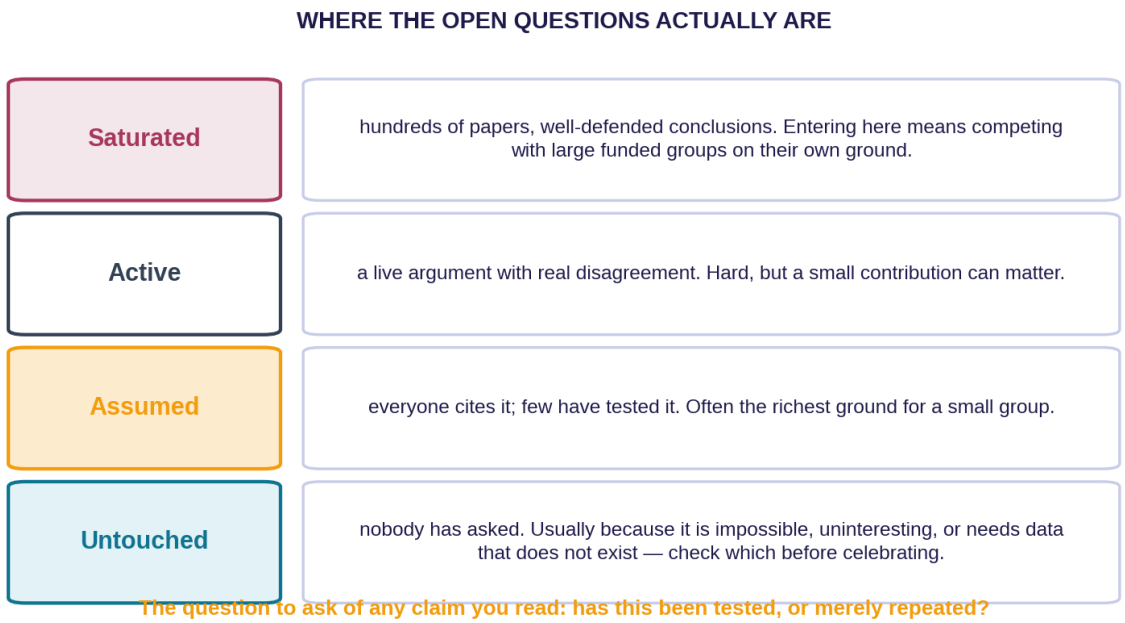

Figure 3.23 Four kinds of territory. Only two of them are worth a small group's time.

The third zone is the one to learn to see. A claim that everybody cites and nobody has tested is the most productive thing a small group can find, because testing it requires no new data — only the discipline to ask whether the emperor is dressed. When you read, keep a running note of sentences that are asserted rather than demonstrated.

3.15.1 Four ideas this group explored and abandoned

Killing your own ideas is a skill, and it is not taught often enough. Here are four directions this group pursued in cancer metabolism, and why each was set aside. Each was abandoned for a different reason, and the reasons are the lesson.

- **Metabolic prognostic signatures.** Abandoned on saturation. The literature is enormous, the methods are formulaic, and adding another signature would have contributed nothing. A field being large is a reason to look elsewhere, not a reason to join in.
- **Collateral lethality.** Abandoned on saturation, and on the realisation that the well-known examples were well known precisely because they were the easy ones.
- **Single-cell metabolic heterogeneity.** Abandoned on feasibility. The question was good; the data required to answer it properly was not available to us. Recognising this early saved months.
- **Single-gene metabolic dependency in one tumour type.** Abandoned on the result. A sustained analysis returned controlled nulls with positive controls firing correctly — the method worked and the effect was not there. This is the most important kind of abandonment, and the hardest.

> ◆ **KEY IDEA — What survived, and why**
>
> The direction that survived was not the most exciting one. It was the one where a specific, testable gap remained after everything else had been eliminated — an assumption widely relied upon and not directly checked.
>
> Notice that this is strong inference from Chapter 1, applied to a whole research programme rather than a single result. You enumerate the possible directions, and you kill them until the survivors are worth your time.

3.16 End-of-chapter questions

Write a paragraph on each. Several have no single right answer.

- Explain to a friend with no biology background why cancer is not one disease. Use two specific examples.
- A paper reports that gene X is more highly expressed in tumours than in normal tissue and concludes it is a promising drug target. List every objection you can raise.
- Look again at Figure 3.6. What does the presence of four separate polyps alongside one cancer suggest about how the disease began in that patient?

- Why might a drug that kills a dividing cell harm a patient, and what property would a better target need to have?
- Pick any cancer type not covered here. Find its three most common driver genes and say, for each, whether it is an oncogene or a tumour suppressor and how you can tell.
- Find one sentence in a review article that is asserted rather than demonstrated. Write down what evidence would be needed to test it.
- Of the four abandoned directions in §3.15.1, which do you think was hardest to give up, and why?

3.17
Go deeper

Every link is live. [Watch] is video, [Read] is written, [Source] is the original paper.

IF YOU ARE STARTING FROM ZERO

[Watch] **Crash Course Biology — full series** — DNA, transcription, translation, cell division and respiration. Cover §3.1–3.3 and §3.10–3.11 painlessly.

[Read] **Khan Academy — Biology library** — Free and worked. The cellular respiration and central dogma sections are the relevant ones.

[Read] **NCI — What Is Cancer?** — The clearest short account there is, written for the public by people who know.

THE ORIGINALS

[Source] **Hanahan & Weinberg (2011) — Hallmarks of Cancer: The Next Generation** — The organising framework for §3.7, and the paper that put metabolic reprogramming on the list.

[Source] **Hanahan (2022) — Hallmarks of Cancer: New Dimensions** — The updated view. Read it after the 2011 paper, not instead of it.

[Source] **Vander Heiden, Cantley & Thompson (2009) — Understanding the Warburg Effect** — Why a dividing cell's metabolism differs. The source for §3.13 and §3.14.

[Source] **Vogelstein et al. (2013) — Cancer Genome Landscapes** — Drivers versus passengers, oncogenes versus suppressors, across many cancers. The source for §3.4 and §3.6.

[Source] **Pavlova & Thompson (2016) — The Emerging Hallmarks of Cancer Metabolism** — A structured tour of how cancer rewires metabolism.

[Source] **Rath et al. (2021) — MitoCarta3.0** — The inventory of nuclear-encoded mitochondrial proteins behind §3.12.

3.18
Glossary

Anabolism / Catabolism The building-up and breaking-down halves of metabolism. Catabolism releases energy; anabolism spends it.

ATP The cell's universal energy currency, produced by catabolism and spent by everything else.

Central dogma The flow of genetic information: DNA is transcribed to RNA, which is translated to protein.

Clonal evolution The process by which mutation and selection produce a genetically mixed, adapting tumour population.

Cristae The folds of the mitochondrial inner membrane, where oxidative phosphorylation takes place.

Driver / Passenger mutation A driver confers a growth advantage and helps cause the cancer. A passenger is incidental.

Gain / Loss of function A mutation that makes a protein overactive or permanently on, versus one that breaks or removes it.

Gene expression The degree to which a gene is transcribed into RNA in a given cell; in effect, a cell's operating programme.

Glycolysis The cytosolic pathway splitting glucose into pyruvate. Fast, oxygen-independent, low yield.

Hallmarks of cancer The set of capabilities tumours tend to acquire, including reprogrammed metabolism.

Hotspot mutation A specific, recurrent activating point mutation; the characteristic way oncogenes are switched on.

MitoCarta A curated inventory of the roughly 1,100 nuclear-encoded mitochondrial proteins.

Mitochondrion The organelle producing most cellular ATP, and a hub for biosynthesis, redox balance and cell death.

Oncogene A growth-promoting gene that drives cancer when a gain-of-function change switches it on. The accelerator.

Oxidative phosphorylation The mitochondrial process generating the bulk of cellular ATP using oxygen.

Somatic mutation A mutation acquired during life, in the tissue, rather than inherited.

Stroma The connective-tissue component of a tissue or tumour, distinct from the epithelium.

TCA cycle The mitochondrial cycle that oxidises fuel and supplies precursors for biosynthesis. Also the Krebs cycle.

Tumour microenvironment The full tissue of a tumour — cancer cells plus stroma, immune cells and vasculature.

Tumour suppressor A growth-restraining gene that permits cancer when a loss-of-function change disables it. The brakes.

Warburg effect The reliance of many tumours on glycolysis with lactate secretion even when oxygen is available.