

貳拾肆 · Last Chapter

無盡 · Chapter 24 — the manual. How to find the gaps in your own art, generate candidates, and refuse to believe them until a body says so.

Chapter 24 — the manual. Every other chapter in this book taught you a thing to train. This one teaches you how to make more chapters — find the gap, draft the lesson, refuse to believe it until a body says so — and then hands the method over. It is called the Last Chapter because it is the last one I write alone.

Part I · The queue that grows

Start with a number, not an idea.

Long before any of this was automated, my own backlog carried a bucket labelled drills that need my body to validate. Entries went in and did not come out. Not because they were hard — because there was only one body, it was frequently sore, and the list grew faster than the body could clear it.

That is the finding, and it arrived before the machine did: **generating candidate knowledge was never the constraint**. Reading, cross-referencing, spotting a gap, drafting a drill — these are cheap and they were always cheap. What is expensive is the part where somebody stands up and finds out.

Worked, on mine. The bucket held four drills, each named and specified, each waiting on a body:

DRILL	FROM
Feed the Line — drop-and-tilt to slip, then add the kick	Hen / L25b
Drop-a-Zone — hand smothered → elbow → low kick	Ten-Chi-Jin / L26
Door Both Ways — FKD move ↔ Taido technique on a shared sotai	Axis Graft / L25
Sho-to-Sho alternating-kick order — still explicitly "loose"	L22b

Logged 4 July. Written in an afternoon. As of this chapter, **52 days later**, not one of the four has been through a body. Nothing was blocking them. No equipment was missing, no research was outstanding. There was simply one body, and it had other things to do.

Four drills, fifty-two days, zero validated. That ratio is the entire argument of this chapter, and it was already true before anything was automated.

Do this first. Before you build anything, count your own backlog. How many things do you believe about your practice that you have never actually tested? If that number is small, you are not reading enough. If it is large, you have just found your real bottleneck — and it is not the one you assumed.

Part II · Build the corpus

The generator is only ever as good as what it can see. Three rules, learned by getting them wrong.

Concretely, a corpus is: source texts chunked into passages, embedded, and stored in a vector database you can query by meaning rather than by exact string. Mine holds **40 collections** — training manuals, art-specific texts, sports-science material — and every query in this chapter was run against it in seconds. You do not need forty. You need three systems that disagree, which is the minimum for absence to be informative.

Primary sources, whole. Not summaries. If a system has a founding text, that is what goes in. A summary has already thrown away the thing you are looking for — the odd detail the summariser judged unimportant is precisely where the gaps live.

Breadth beats depth, at first. One art deeply indexed cannot show you a gap, because a gap is a comparison. You need at least three systems before absence means anything.

Mind the language of your queries. This one cost me. Searching my corpus for `xingyi` returned hits; searching it for 形意 returned nothing at all. Not because the material was missing in one form and present in the other — because the corpus is entirely English-language and romanised. Had I searched only the characters, I would have concluded the corpus was empty. Had I searched only the romanisation, I would have missed that `qi` appears 351 times while 氣 appears zero, which tells you exactly what kind of library you have built.

The trap: a null result from a search is ambiguous. It means either the idea is absent or your query is in the wrong language. Always test the second before believing the first.

Part III · Point the gap analysis at your own art

This move has a name and it is already published: **KEGA — Knowledge Extension via Gap Analysis**. Its claim is that a coherent system implies more than it states, and that the implied content can be derived from the shape of what is missing. It has been run on Goodfellow's Deep Learning and on the Guiguzi; this chapter runs it on a martial art.

KEGA sorts work into three modes, and **knowing which mode you are in decides what your output is allowed to claim**:

MODE	WHAT IT PRODUCES — AND WHAT A MARTIAL ART IS
Science	Coherent formal systems. Output is verifiable — you can prove or test it on paper. A biomechanical claim about levers sits here.
Art	Output is plausible and constrained, but not uniquely determined. Most of a martial art lives here. This is why the ledger exists: Art Mode output cannot be proven, only tried.
Experiment Generation	The middle layer — derive the constrained set of tests that would confirm or falsify a claim. This is the mode you actually want. Do not generate lessons; generate the drills that would settle them.

An honest amendment to my own method. KEGA's central trade is that it converts a discovery problem into a verification problem, on the grounds that verification is cheaper. In formal domains that is true — checking a derivation costs less than finding one. **In an embodied art it inverts.** Verification needs a body, a warm-up, recovery time, and it returns $n=1$. Discovery got cheaper and verification did not, so the whole cost of research moved onto the slowest component. That is not a flaw in KEGA. It is what KEGA looks like when you point it at something made of tissue.

Here is the move, worked end to end on a real gap, including the part where it fails.

Step 1 — seed it with a question, not a vocabulary. This is the part I got wrong first, and the correction is the most useful thing in this chapter.

My first version asked you to describe the subject before searching for it: list the members of the system, then list the technical terms a text teaching each one would contain. For Chinese internal boxing that means naming 三拳 — Taiji, Bagua, Xingyi — and then, for Xingyi, naming the five fists. **Which requires you to already know the thing you are asking about.** The seed had the answer baked into it, so it could only ever find gaps in subjects I could already describe.

Seed with a question instead:

What is the 80/20 of Wing Chun?
What are the core fundamentals of Xingyi?
How is unsoku actually performed?

A question needs no vocabulary list, because **the question is the probe**. Semantic search takes questions natively — that is what it is for — and the enumeration you were going to write by hand comes back out of the sources instead, which means it is grounded rather than remembered.

Step 1b — the guard is already inside the question. Semantic search has one property that will fool you: **it can never return nothing**. Ask it for a subject your library does not contain and it

returns the nearest neighbours, at confident-looking distances, with no signal that anything is wrong. Asked "what are the core fundamentals of Xingyi" against a corpus holding no Xingyi at all, mine returned eight passages — seven on Tai Chi and Qigong, and one that did say "Xingyiquan", inside a list of style names in a Qigong book. Keyword matching, for all its faults, had exactly one virtue this lacks: it could say nothing here.

The fix costs nothing, because the subject is sitting in the question already. Pull it out and require every returned passage to actually mention it. Then count:

QUESTION	PASSAGES MENTIONING THE SUBJECT
80/20 of Wing Chun	8 / 8
80/20 of Capoeira	8 / 8
80/20 of Brazilian Jiu Jitsu	7 / 8
fundamentals of Xingyi	1 / 8

That last row is worth sitting with, because it is the honest version of a result I first wrote down as a clean zero. One passage in eight did name Xingyi — in a list of style names. **Named is not documented.** A ratio near zero is not a rounding error to argue away; it is the corpus telling you it can supply the word and not the content.

That ratio is the whole test. Retrieval for recall, the subject name for precision, and the fraction between them tells you whether your library can answer the question at all. It needs no threshold tuning — and I say that having spent an evening tuning thresholds before noticing the ratio was sitting there the entire time.

The question shape decides how much of the answer is yours. "How is unsoku performed" is answerable from sources — books describe technique. "What is the 80/20" is not: no book ranks its own contents, so any 80/20 you get back is inference wearing citations. Both are legitimate questions. Only one of them produces a claim the corpus actually made, and your Limit section has to say which you asked.

Step 2 — ask what your own canon says about each member. Mine said:

ART	CHAPTERS	VERDICT
Taiji	Ch11, Ch15, Ch19	mined
Bagua	Ch13	mined
Xingyi	none	absent

Zero files. Not thin — absent. Two legs of a three-legged system, written about extensively, and the third never mentioned once. That is a gap analysis producing a result in about ten seconds, and it is the kind of thing that is invisible from inside because you never notice the chapter you did not write.

Step 3 — and this is the step people skip — ask whether your corpus can actually fill it. Mine could not:

SEARCHED	HITS
pi / beng / zuan / pao / heng quan	0
splitting / crushing / drilling / pounding / crossing fist	0
five element fist	0
santi	3

The entire technical core of Xingyi — its five fists — was not in my library, in either language. So the honest output of that gap analysis was not a chapter. It was a **shopping list**.

The result that matters: gap analysis is bounded by its corpus. It finds holes in what you have read, which is not the same as holes in the art. A generator that forgets this will happily invent the missing leg rather than admit it cannot see it — and the invention will be fluent, plausible, and worthless.

Part IV · Draft the lesson

A gap is not a lesson. It is a vacancy. This is the step where you fill it, and it is the step this book has so far only ever done without ever explaining.

The skeleton. Every chapter here has the same bones. Steal them:

SECTION	JOB
The claim	One sentence. What is true that a reader does not already know? If you cannot write it in one sentence you have a topic, not a lesson.
The evidence	Where it came from. Name sources. If the only source is your own feeling, say that — it is still a lesson, it is just a weaker one.
The mechanism	Why it is true. A lesson that says "do this" and not "because this" cannot be argued with, and a claim that cannot be argued with cannot be corrected.
The drill	What the reader physically does, with a number attached. Reps, angles, a count that should move week over week.
The limit	What this does not establish. Non-optional.

Worked, gap to draft. Take a gap where the corpus actually can help — the opposite of the Xingyi case. My footage of Taïdo's unsoku happo teaches the eight steps feet-first, then adds the upper body in a second lesson. The gap: nothing in my canon said what the upper body is actually constrained by. Working the skeleton:

SECTION	FILLED
---------	--------

Claim	In unsoku the guard's elbow line is fixed and its height is free.
Evidence	The instructor is explicit: high guard or low guard, he does not mind — what matters is that the elbows stay on the inside, and that the elbows change before the feet do.
Mechanism	The elbow line is what closes the centre. Guard height only changes which target you cover; elbow position changes whether you have a structure at all. So the art constrains the load-bearing part and leaves the cosmetic part to the practitioner.
Drill	Step the eight-step pattern with hands only, no feet. 3×20 changes. Count how many times the elbows drift outside the line. Score should fall week over week.
Limit	One instructor, one school's interpretation. Not tested against other Taïdo lineages, and not tested on my body under pressure.

That is a complete lesson and it took ten minutes, because the gap analysis had already done the hard part — deciding what was worth writing about. Note what it is not: it is not validated. It goes into the ledger at **derived** and stays there until somebody drills it.

Prefer Experiment Generation mode. Notice that the strongest row in that table is the Drill, not the Claim — because a drill with a count is falsifiable and a claim is not. If you find yourself unable to write the Drill row, you have generated an opinion. Go back to the gap and ask what test it implies instead of what lesson it implies.


The rule that keeps this honest: write the Limit row before you write the Claim row. If you draft the claim first you will unconsciously shape the limit to protect it. Written first, the limit tells you how big a claim you have actually earned.

Part V · The honesty ledger

If you build only one part of this system, build this one. A generator without a ledger is not a research tool, it is a landfill with good grammar.

Every artifact carries exactly one status:

STATUS	MEANS	WHO MAY SET IT
derived	Generated from a gap. A hypothesis.	the machine
corroborated	An independent source agrees. Still untested on a body.	the machine
body-validated	Someone did it, or an instrument measured it.	a body, or an instrument. Never the machine.



The honesty ledger

Gold: **derived** → **corroborated**. The machine may cross that arrow on its own.

Crimson: **body-validated**. Only a body or an instrument writes it, and the arrow points back — validation is something done to a body. The wall is structural, not a policy.

Worked, on a real claim. Follow this book's own central finding through the three states and note where it stops:

STATE	WHAT EARNED IT
derived	Ch21. One clip, pose-estimated: the pelvis peaked 200 ms after the thigh. A measurement of one repetition on one day.
corroborated	Ch23. Two unrelated witnesses land on the same joint — a Korean stylist naming the pelvis as what karate's upright torso locks out, and Taido's founder rebuilding the art around movement in three dimensions. Neither could see the clip.
body-validated	Not reached. The direction is corroborated. The thresholds — chamber under 90°, retraction ratio ≤ 1.0 — were authored from coaching literature and have never been tested against a corpus of clips or against another body.

That is what a correctly-kept ledger looks like: the flagship finding of the book sits at **corroborated**, not at the top. If your ledger has everything in the third column, it is not a ledger, it is a trophy cabinet.

Two rules make it hold. **No citation, no candidate** — if the generator cannot name where an idea came from, it does not produce the idea at all. And **the status is printed on the artifact**, not buried in metadata, so a reader always knows what they are holding.

One more, which is the rule that makes this chapter honest rather than triumphant: **generation pauses when the unvalidated queue gets long**. If you have ten untested drafts waiting, you do not need an eleventh. You need to go outside. The system is built to enforce the thing this book has argued for twenty-three chapters, rather than merely mention it.

Part VI • Validate on your own body

The ledger's third column is the whole point, and it is the only part that cannot be automated, delegated or accelerated.

Use what Chapter 21 established: film it, measure geometry rather than impressions, and let the tool report its own blind spots. Use what Chapter 23 established: a drill with a number attached beats a drill with an adjective attached. Count something. Watch the count move.

And carry the caveat forward intact. A single clip supports a claim about direction — this kick was thrown from the knee, that sequence was inverted. It does not support a threshold. Quoting an angle to one decimal place tells you the estimator was confident; it tells you nothing about whether that number is right for anyone else's body, including your own on a different day.

Worked, on one kick. The whole loop, as it actually ran:

STEP	WHAT HAPPENED
1 · Film	Side view, support foot unoccluded. One rep, thrown normally.
2 · Measure	Chamber 118.8° · extension 164° · lean 42.9° · COM drift 0.001 m · retraction 1.62 · pelvis 200 ms after thigh.
3 · Compare	Three failures against criterion: loose chamber, slow retraction, inverted sequence. One excellent result: COM drift.
4 · Pick one	Sequencing. Not because it was worst, but because the other two are what a knee-thrown kick produces — fix the cause, not the symptoms.
5 · Retest	Two weeks. Same camera position, same distance, same clothing.

Step 4 is the one that takes discipline. Three failures is three temptations to fix everything at once, and a session that changes three variables tells you nothing about which one moved the number.

The failure mode to watch for in yourself: a derived lesson you never tested is not knowledge. It is a to-do list wearing the costume of a finding. The pleasure of generating is real and it is a trap — it feels like progress and produces none.

Part VII · Two masters, opposite calls

There is a serious argument against everything in this chapter, and it deserves to be made properly rather than knocked down.

While researching Chapter 23 I watched a Korean-style master analyse Taïdo. Partway through he mentions that he once worked out a method for kicking high reliably — significant enough that he changed his own technique around it. Then he declines to explain it. His reasoning, paraphrased: every new piece of know-how is followed by everyone catching up, and once that happens your school has no advantage left. So it stays inside the school.

That is not paranoia. It is a coherent strategy, it is how most transmission has worked for most of history, and it is why so many arts have a layer that is only taught to people who stay twenty years.

I am doing the opposite, and here is the argument, stated so you can weigh it against his.

A method kept in one school is validated by one school. Secrecy protects an advantage and starves the evidence. The Chapter 23 finding only became credible because three unrelated witnesses converged on it — a pose estimator, a founder who rebuilt karate after learning to steer a torpedo through nets, and a competitor who beat faster opponents on a broken foot. None of them could see the other two. Had any of them kept their method in-house, that convergence never happens and the finding stays a hunch.

He is protecting an edge. I would rather have the evidence. Both of us are trading something away, and you should know which trade you are making.

Part VIII · 無盡

So: which limit was the illusion?

Not the body. The body is extremely real. It gets tired, it gets injured, it needs a day off, and it can only honestly test a few things a week. Every hour spent generating drills does nothing to change that, and this book has spent three chapters insisting on exactly that point.

The illusion was the **ceiling**. It always feels as though there is a fixed amount to learn and that one gets steadily closer to the end of it. There is not. Every gap closed opens two more; every graft between two systems creates a seam that has never been examined. The material is unbounded, and it was unbounded the entire time you were worrying about running out.

And the second illusion — the quieter one — is that **the rate is fixed**. The rate is one body per lifetime only if you insist on being the only one measuring. Hand the method over and the denominator changes. Validation stops being what one person can survive and becomes what a practice can accumulate.

That is why this chapter is a manual and not a demonstration. A demonstration would have been me showing you a machine. A manual is the machine.

Volume I closed with a gate left open — build, learn, or find the holes. That was an invitation. This is the same gesture with the tools attached.

"Limits, like fears, are often just an illusion."
— Michael Jordan

This is Chapter 24.

"There are 24 hours in a day, so why are you only using 8?"
— Mamba Out

Called the Last Chapter because it is the last one written alone. Whether it stays the last one is now, properly, not up to me.