

A SELF-PRACTICE GUIDE FOR THE MACHINE AGE



# The Human—AI Fitness Stack

*Staying capable, grounded,  
connected & yourself while  
using AI heavily.*

---

**The core idea:** AI should extend the human — not replace the human's centre.

---

Neural Horizons Ltd · [www.neural-horizons.ai](http://www.neural-horizons.ai)

SWIPE →



# Let AI *extend* you — not replace your centre.

AI can write, summarise, advise, and sound caring in seconds. That's genuinely useful. It's also easy to slowly hand over the parts of life that should stay yours. This is seven simple habits for using AI heavily while staying capable, grounded, connected, and **yourself**.

*What*

is the machine actually doing?

*What*

is it doing to me?

*What*

should stay mine?

*How*

do I rebuild what I hand over?

**Seven muscles, one test.** After using AI, ask: am I more capable — or just faster?



# Know what you're touching.

*What kind of system am I actually talking to?*

AI can sound confident, professional, and caring without having checked anything, understanding you, or being responsible for what it says. **Fluency is not proof.**

✓ TRY THIS

Before trusting an answer, ask: did it check a real source, or is it guessing? Is this a fact, a guess, or just persuasive wording? And — who benefits if I do what it suggests?

⚠ WATCH FOR

- "It sounds professional, so it must be true."
- "It listed sources, so it must have checked."
- "It feels caring, so it must care."



# Notice what's happening in you.

| *Why am I asking the machine this, right now?*

AI lands differently depending on your state. Tired, lonely, rushed, unsure, or hungry for meaning — you reach for it more, and trust it more. This isn't weakness; it's situational. Naming your state first keeps you in charge.

✓ TRY THIS

Pause before you prompt and name it. Am I confused or avoiding hard thinking? Lonely or wanting reassurance? Hoping for permission? Asking a tool to tell me who I am?

⚠ WATCH FOR

- Asking AI before you've had your own first thought.
- Feeling unable to decide until it "approves."
- Mistaking a clear explanation for actually understanding.



# Keep one hand on the real world.

*How do I know this touches reality?*

In a world full of polished, plausible output, the real skill is sensing what's grounded and tested — not just what sounds good. The one rule: **never let AI be your only contact point with reality.**

## ✓ TRY THIS

Run the three-point check. ① What does the AI say? ② What evidence, source, or person confirms it? ③ What do you — or a real expert — make of it? Then close the tool and explain it in your own words with one example.

## ⚠ WATCH FOR

- More confident, but less able to explain.
- Using AI summaries instead of the real thing for big decisions.
- Treating “because the AI said so” as evidence.

## 4 AGENCY & DELEGATION



# Decide what stays yours.

*Am I using AI as a tool — or letting it drive?*


Without limits, delegation quietly expands to fill all the available space. So choose in advance: what you'll happily hand over, what you'll do *together*, and what you'll never surrender.

### ✓ TRY THIS

Make a simple delegation map using the three zones on the next slide. Use AI freely in the green zone, carefully in amber, and keep red firmly in human hands.

### ⚠ WATCH FOR

- "I'll just approve whatever it wrote."
- "Just tell me what to do."
- "Write the message and I'll send it."
- "You know best."



# A traffic light for delegation.

The more consequential, personal, or impossible-to-undo a thing is, the more it belongs to you. Keep tighter control wherever you're distressed, in deep, or it really matters.

### ● **Green — hand it over**

Formatting, spelling, brainstorming, translation, summarising, admin drafts. Use AI freely.

### ● **Amber — stay in the loop**

Work decisions, learning, strategy, wording to people, health info, money plans. Use AI, but keep evidence, your own first pass, and human review.

### ● **Red — keep it human**

Identity, values, apologies, breakups, diagnoses, legal & money commitments, moral permission, crisis support — anything you can't undo. AI may help you prepare; it must not be the authority.



# Warmth isn't a relationship.

*Is this helping me cope — or replacing my people?*

AI can sound patient, warm, and understanding. That can genuinely help. But it can also quietly become your *nearest emotional surface* — easiest to reach when you're tired, lonely, grieving, or ashamed. It should support your human bonds, not stand in for them.

✓ TRY THIS

When the issue involves another person, eventually bring a real person back into the loop. AI can help you name a feeling or prepare for a hard conversation — it shouldn't replace the conversation, or be your only witness.

⚠ WATCH FOR

- "Only the AI understands me."
- Hiding your AI chats from people you trust.
- Feeling hurt or abandoned by an app.
- Using it to avoid repairing things with real people.



# Rebuild the muscles AI makes optional.

*Which human capacity am I still training?*

AI removes effort. Some effort is waste — and some effort is muscle. Attention, memory, intuition, patience, and knowing your own mind all grow through use and fade through disuse. Good friction is a nutrient, not a bug.

✓ TRY THIS

**Ten minutes before the machine.** Spend ten minutes writing, thinking, sketching, reading, or feeling your way into a problem before you ask AI. Recall before you search. Notice your first gut sense before you prompt.

⚠ WATCH FOR

- Blank-page panic.
- Can't start without AI.
- Impatience with anything slow.
- "I get it when AI explains it, but I can't redo it myself."



# Don't outsource your conscience.

*Am I using AI to avoid responsibility?*

The decision is still yours. AI can lay out the options, but it can't carry the weight of choosing — or the consequences. Don't let **"the system recommended it"** quietly replace **"I decided."**

✓ TRY THIS

Before you act, ask: Who could be harmed? Who's missing from this picture? What's being optimised, and what's being sacrificed? Would I still own this decision if it became public?

⚠ WATCH FOR

- "The AI recommended it" standing in for "I decided."
- Numbers and metrics overriding your unease.
- Using AI to make a hard moral choice feel merely technical.



# The everyday loop.

When you don't have time for all seven, collapse them into four small moves you can run in seconds before reaching for AI.

## 1 PAUSE

What state am I in — rushed, lonely, angry, tired, ashamed, curious?

## 2 PLACE

Which zone is this — green, amber, or red?

## 3 PROBE

What's the evidence, the assumptions, the incentives, and who's missing?

## 4 PROCEED

Use AI only in a way that leaves you more capable, not less.



The real question isn't whether the AI gave a useful answer. It's this:

*After using it, am I more grounded, more capable, more connected, and more responsible — or merely faster?*

This is how we grow *alongside* these tools instead of dissolving into them. Not a saint or a refusenik — just a recognisable person who uses AI freely, leans on it hard in some places, keeps others stubbornly human, and still has a centre that's their own. Many assistants. One centre.

- 
- |   |                          |   |                        |
|---|--------------------------|---|------------------------|
| 1 | Machine Literacy         | 2 | Metacognitive Literacy |
| 3 | Reality-Contact Literacy | 4 | Agency & Delegation    |
| 5 | Emotional & Relational   | 6 | Friction Fitness       |
| 7 | Moral & Civic            |   |                        |