

Borrowed From an Art School

Where the competency framework came from, and why its provenance outside medicine is the reason it transfers to clinical work.

Dr Neal Aggarwal

Medicine · Surgery · Engineering · Information Technology · Artificial Intelligence

10 AUGUST 2026

The competency framework underneath my clinical AI blueprint was not built for medicine. It came out of an art school in Sarasota and a business school in Cork, developed by a novelist-turned-AI-coordinator and an information systems professor. That provenance is not a curiosity — it is the reason the framework transfers to clinical work at all. A framework built for medicine would have hard-coded medicine into it. This essay sets out what Dakan and Feller built, what I kept, what I reframed, and the one thing I had to add because creative work does not kill anyone.

Contents

The two people	3
Why the provenance is the point	4
It is grounded, not armchair	4
The internal structure is tidier than the four names suggest — but only partly	5
What I kept	5
What I reframed	6
What I reweighted	6
What I added	6
The detail I liked most	7
Two licences, not one	7
Sources	8

When I published the blueprint for a national clinical AI institute last week, the competency framework underneath it carried an attribution line that most readers will have skimmed past:

Domains adapted from the AI Fluency Framework by Prof. Rick Dakan (Ringling) and Prof. Joseph Feller (UCC) / Anthropic.

I want to spend a post on that line, because the most interesting thing about the framework I built on is where it came from — and because the provenance turns out to be the reason it works.

It did not come from medicine. It came from an art school and a business school.

The two people

Rick Dakan is Professor of Creative Technologies at Ringling College of Art and Design in Sarasota, Florida, where he is also the college's AI Coordinator and Interim Co-Director of its Center for the Creative Economy. He oversees Ringling's Undergraduate Certificate in Artificial Intelligence and its Professional Certificate in Fundamentals of AI for Creatives.

His background is not what a clinician would expect. He is a game designer and an author, with more than thirty games and books behind him — video games, tabletop games, novels, non-fiction, comics.

Joseph Feller is Professor of Information Systems and Digital Transformation at Cork University Business School, University College Cork. His research sits on AI-human hybrid creativity, innovation and learning, and it has been published in the journals that field takes seriously — *Information Systems Research*, *Journal of Management Information Systems*, *Journal of the AIS*, *Journal of Information Technology*, *Information Systems Journal*, *European Journal of Information Systems*, *Journal of Strategic Information Systems*. His work has been funded by the European Commission, the Irish Research Council and the Irish Higher Education Authority.

Their collaboration explores the intersection of human creativity, generative AI, and teaching in higher education. The framework that came out of it was supported by Ireland's Higher Education Authority through the National Forum for the Enhancement of Teaching and Learning, and was then elaborated into an open course series in partnership with Anthropic.

So: a novelist and games designer running AI at an art school, and an information systems professor at an Irish business school, funded by an Irish teaching-and-learning body, producing a framework that I have made the spine of a proposal to train Kenyan surgeons and ward nurses.

That sounds like an accident. It is not.

Why the provenance is the point

Here is the counterfactual that convinced me.

A framework built for medicine would have hard-coded medicine into it. It would talk about diagnosis and prescribing and differential reasoning. It would read beautifully to a physician — and it would need rewriting the moment you handed it to a hospital administrator evaluating a procurement, or a health informatician validating a model against the local case mix. Both of those are tracks in the blueprint. Both would have been second-class citizens in a framework built around the consultation.

Dakan and Feller went the other way. They claim their framework is **platform and technology agnostic**, adaptable to emerging tools and use cases, and that it characterises *effective action* rather than prescribing rigid process. They define AI fluency as the ability to work “effectively, efficiently, ethically, and safely” within the emerging modalities of human-AI interaction.

That agnosticism is not a nicety. It is precisely the property that lets the teaching survive a vendor disappearing overnight — which happens to be design commitment number one in my blueprint:

We teach judgement, not tools. Test every module against one question: *if the vendor disappeared overnight, would this teaching still be worth anything?* If no, it is training, not education, and it belongs in a vendor manual.

I did not have to add that property. I inherited it. A framework written from inside one profession, about one generation of tools, would have failed that test on the day it was published.

It is grounded, not armchair

The other thing that persuaded me is that this is not a model somebody drew on a whiteboard.

Dakan and Feller have been running it through **actual student courses and faculty seminars at both institutions across three consecutive academic years** — 2023/24, 2024/25 and 2025/26. It has been taught, revised, taught again. The version that reached me had already survived contact with real learners who were free to find it useless.

For anyone building a curriculum, that matters more than the elegance of the diagram. Most competency frameworks in circulation have never been taught from.

The internal structure is tidier than the four names suggest — but only partly

Everyone quotes the four Ds. Fewer people notice that each one has three subcategories underneath it, set out in the authors' Practical Summary Document.

Two of the four follow the same clean pattern — **Product, Process, Performance**:

	Product	Process	Performance
Description	defining the output you want	dialogic, iterative prompting	defining future AI behaviour
Discernment	judging output quality	judging whether the collaboration is working	judging AI-driven behaviour in use

And that triad lines up with the framework's three **modalities** of interaction — automation, augmentation and agency — which is a genuinely satisfying piece of design once you see it.

But I should be accurate about how far it goes, because I nearly wasn't. The other two competencies have their own triads and do not follow the Product/Process/Performance pattern. Delegation breaks into Goal & Task Awareness, Platform Awareness and Task Delegation. Diligence breaks into Creation, Transparency and Deployment Diligence. The mapping is elegant and it is partial. Anyone building on the framework should know that before they design a curriculum grid around a symmetry that only half exists.

The figure at the top of this post sets it out honestly, with the two patterned rows solid and the two unpatterned rows dashed.

What I kept

Working through the adaptation, I kept more than I changed.

All four competency names, unaltered. The three modalities. The two-loop structure that the authors' teaching course adds — Delegation and Diligence as the outer loop, what you decide to hand over and what you answer for afterwards; Description and Discernment as the inner loop, turning many times within a single encounter. And, importantly, their framing of **ethics as fundamental rather than bolted on** — the framework treats responsible use as constitutive of competence, not as a module you bolt to the end.

Clinicians recognise the two loops immediately when you name them as consent-and-audit on the outside, history-and-examination on the inside. That recognition is doing a lot of teaching work, and it is not mine.

What I reframed

One change is substantive rather than cosmetic.

In the original, Delegation includes **Platform Awareness** — understanding the capabilities and limits of the available tools, and evaluating them against project requirements, budget, and operational and regulatory needs. That is an **optimisation problem**. Which tool serves the goal best?

In clinical work, delegation is not an optimisation problem. It is a **safety boundary**. So I replaced the criterion. The question is not *which tool serves the goal best* but:

What is the cost if this is wrong and nobody notices?

Everything downstream follows from that substitution. Reversibility and severity of error become the test of what may be delegated, rather than convenience or capability. And out of that comes the **non-delegable list** — the clinical acts that never leave your hands regardless of how good the tool becomes: obtaining consent, breaking significant news, the decision to operate, the final diagnosis committed to the record, the prescription, the signature.

A creative professional who delegates badly produces worse work. A clinician who delegates badly produces a harmed patient. Same competency, different physics.

What I reweighted

In the original the four Ds are broadly co-equal. In my clinical adaptation they are not: **Discernment takes 40% of the assessment blueprint and two of the twelve teaching hours** in the Level 1 common core.

That is not a refinement of Dakan and Feller's design. It is a response to evidence they were not writing about.

A 2025 randomised trial published in *NEJM AI* gave clinical cases to physicians who had already completed twenty hours of AI-literacy training, and exposed half of them to deliberately erroneous model output. They deferred to it. Consultation was voluntary and they retained full autonomy to reject the suggestion. They accepted it.

If familiarity with these systems does not confer the capacity to catch them being wrong, then the hours have to go into the reflex of checking rather than into fluency with the tools. The weighting is the pedagogical claim expressed as a timetable.

What I added

And here is the sharpest edge between the source and the adaptation.

The framework has no concept of conjunctive failure, and it does not need one. Nothing in it says a competency cannot be compensated by strength elsewhere — because in creative and problem-solving work, it can. Get the delegation wrong on a design brief and good discernment rescues you. Describe the task badly and iteration recovers it. The competencies genuinely do trade off against each other.

Clinical work does not permit that everywhere. Enter identifiable patient data into a system you do not control, and nothing rescues you. Fail to detect a seeded clinical error at threshold, and no amount of elegant prompting compensates.

So the competency standards specify two **conjunctive** requirements — criteria that cannot be traded against aggregate performance. The precedent is ordinary in medical assessment: a candidate does not pass an OSCE by offsetting a fatal drug error with excellent communication. But it is genuinely absent from the source framework, and it is mine.

That distinction is the honest answer to “how much of this is yours?” The architecture is theirs. The safety physics is mine.

The detail I liked most

Their Practical Summary Document carries a **Diligence Statement**. They disclose that they used Claude to help draft and refine the text, affirm that all AI-generated content was vetted, edited and curated by the human authors, and take full responsibility for the content, its accuracy and its presentation.

They practise the competency they are describing, in the document describing it.

I have mirrored that. Each of the three annexes to the blueprint — the business case, the competency standards and the Level 1 common core — closes with an AI-assistance disclosure rather than staying quiet about it. It costs nothing, it demonstrates the Diligence domain rather than merely asserting it, and I would rather a Council registrar read it from me than infer it.

Two licences, not one

A practical note for anyone thinking of building on this, because it caught me out and it matters.

The two artefacts carry **different** Creative Commons licences.

- The **open course materials**, including video, are released under CC BY-NC-SA 4.0 — attribution, non-commercial, share-alike. Adaptation is permitted provided you credit the authors and release your derivative on the same terms.
- The authors’ **Framework for AI Fluency (Practical Summary Document)**, v1.1, is released under CC BY-NC-ND 4.0 — **NoDerivatives**. Redistribution is permitted; adaptation is not.

In practice: you may build a clinical adaptation from the course materials, and you must cite rather than remix the summary document. My annexes do exactly that, and carry the share-alike obligation forward.

Building an institution about diligence on a careless reading of somebody else's licence would be an unpromising start.

Sources

- **AI Fluency Framework** — the authors' own site, with biographies, the framework definition, and the open course series: aifluencyframework.org
- **Framework for AI Fluency (Practical Summary Document), v1.1** — Dakan, R. and Feller, J., Ringling College of Art and Design, 2025. The source for the modalities, the four competencies and their subcategories: ringling.libguides.com/ai/framework
- **The open course series on Anthropic Academy** — AI Fluency: Framework & Foundations · Teaching AI Fluency · AI Fluency for Educators · AI Fluency for Students
- **Also hosted openly** at [OpenCourses.ie](https://opencourses.ie)
- **Institutions and funders** — Ringling College of Art and Design · Cork University Business School, UCC · Higher Education Authority, Ireland · National Forum for the Enhancement of Teaching and Learning · Anthropic
- **The evidence behind the reweighting** — *Automation Bias in Large Language Model-Assisted Diagnostic Reasoning among Physicians Trained in AI Literacy — A Randomized Clinical Trial*, NEJM AI, 2025
- **Licences** — CC BY-NC-SA 4.0 (course materials) · CC BY-NC-ND 4.0 (summary document)

Suggested citation for the framework: Dakan, Rick and Feller, Joseph. "Framework for AI Fluency (Practical Summary Document)," Version 1.1, [Ringling.edu/ai](https://ringling.edu/ai), 2025.

The clinical adaptation, with the four domains written out as numbered performance criteria at three tiers, is in [Annex B — Competency Standards](#). The teaching material that puts them into practice across twelve hours is in [Annex C — the Level 1 Common Core](#). Both are drafts, and I would rather they were argued with than admired.

In the spirit of the framework's own Diligence competency: this post was drafted with AI assistance. The argument, the judgements and any errors are mine.