

Painted Footsteps

There is a certain kind of dog owner who will explain at length why their dog cannot be trained. The dog is exceptionally intelligent. Strong willed. Too independent for treats, too clever for the usual methods. Previous trainers didn't understand him. The breed is notoriously difficult. Each explanation arrives with the confidence of hard-won insight, which in a sense it is — the owner has developed a remarkably detailed theory of why nothing works, refined through long experience of nothing working.

The dog in question was large, enthusiastic, and dragging the neighbor down the street at a pace that suggested he had somewhere important to be. Henry walked beside me on a loose leash, happy to be a dog out for a walk, happy to be by my side. The contrast was not subtle.

She noticed. Most people do, eventually. Her dog, she explained, was just too smart for that kind of thing. Too independent. Wouldn't respond to the usual methods. I asked if she'd mind if I gave it a try.

Five minutes later her dog was walking quietly alongside me. No raised voice, no drama. She took the leash back, and off they went, her dog in the lead, resuming his previous agenda almost immediately.

I've done this little demonstration before. I usually confess that it's something of a stupid pet trick — and mean it. I can show someone how to do this in just a few minutes because the technique is real. What it implies is not. Getting a dog to walk calmly alongside once, with a stranger, in a quiet moment, is not the same thing as having a dog that does it reliably, everywhere, under any circumstances. That requires everything that happens in the other twenty three hours of the day — the accumulated texture of daily life, the thousand small interactions that either reinforce or undermine what happens on the leash.

The failure wasn't the dog's. The dog in that moment was already capable of the behavior. What was missing was the owner's understanding of how to build it, maintain it, make it stick across contexts and circumstances. That understanding doesn't come from watching a five minute demonstration. It comes from something slower and less visible — daily contact, accumulated adjustment, feedback that arrives not as instruction but as consequence.

Contact

Training a dog is not easy. But the world of dog training is forgiving in one specific sense: the feedback is honest. A dog doesn't care about credentials, theory, or explanation for why the previous approach should have worked. Either the approach works or it doesn't. There is no successful performance of dog training. There is only dog training.

Other worlds, other knowledge domains, are not like this. In most the feedback is slower, noisier, more ambiguous. It is possible to be wrong for a long time before the wrongness becomes undeniable, and sometimes it never does. The dog story illustrates something familiar: there is a significant difference between a casual understanding of how something works and the ability to actually do it — a gap that only sustained practice under real conditions can close. Harry Collins and Robert Evans, sociologists who spent decades studying how expertise actually develops, call this contributory expertise: the ability to work within a domain in ways that survive contact with reality. It is built through failure, correction, and the accumulated experience of something that pushes back.

There is a second kind of expertise — interactional expertise — built not through doing but through sustained engagement with a domain's language, debates, and practitioners. It develops through reading, conversation, careful attention to how those who do the work actually think. It lets its holder follow expert arguments, challenge expert conclusions, and translate domain knowledge for those outside it. It is genuinely useful, but it is not the same thing.

Contributory and interactional expertise can develop together or independently. Some of the most genuinely expert practitioners are most comfortable interacting within their own domain — their knowledge lives in their hands, their instincts, their pattern recognition, not in explanations aimed at outsiders. Interactional expertise without contributory expertise underneath is also real and valuable — the science communicator, the skilled manager, the policy generalist, the integrator who holds a complex enterprise together precisely because they aren't captured by any single domain's assumptions.

How visible contributory expertise is varies by domain and by what part of the work is visible. A weld either holds or it doesn't. A diagnosis, a policy recommendation, a financial analysis — the work behind these is harder to see, the process harder still. Across that spectrum, interactional expertise tends to become more prominent as the work itself becomes less legible, when the evidence of expertise isn't visible — not because it substitutes for contributory expertise, but because it is often the only available signal by which expertise can be recognized at all.

Most people develop one capacity more fully than the other. In the early days of Apple, Steve Wozniak was the engineer — contributory expertise in the most literal sense, the person who actually designed and built the computers. Steve Jobs understood what Wozniak was building well enough to direct it, shape it, and explain it to the world, but never pretended he could do what Wozniak did. The partnership worked because both knew which side of that boundary they were on.

When the work is less legible, other signals fill the gap. Credentialing systems are designed, among other things, to be one of them — to certify that the person holding the credential has sufficient contributory experience and accomplishment to be trusted. It has always been an imperfect proxy. As I have argued elsewhere, fluent authoritative output has historically functioned as a costly signal — expensive enough to produce that it correlated reasonably well with the real thing.

Psychologist Daniel Kahneman and decision researcher Gary Klein spent careers on opposing sides of a debate about expert judgment — one skeptical of professional intuition, one its defender. When they sat down to work through their disagreements, they found more common ground than either expected. Central to their agreement was the idea that domains — the subject matter areas in which expertise develops — vary significantly in the conditions they provide for that development.

In high validity environments, the domain itself does much of the calibrating work. The approach works or it doesn't, and the domain registers the failure clearly and quickly. Formal work products contribute to this — the architect's drawings must meet specific format requirements before the content is even evaluated, a constraint that filters competence independently of the outcome.

In low validity environments, outcome feedback is slow, ambiguous, and hard to attribute. The work product is often expressive rather than formal — a recommendation, a conversation, a narrative — imposing no prior constraints on form. In those domains, genuine expertise requires immersion in a community of practice — peer review, collegial challenge, shared standards for what counts as reliable judgment. The community provides what the domain cannot: the ethical as well as cognitive calibration that distinguishes genuine expertise from confident practice.

Neither contributory nor interactional expertise travels freely across domains. Contributory expertise is bounded by the domain where it was built — what develops through contact with one domain's real constraints doesn't automatically hold in another. Interactional expertise is more portable but not without limits — facility with one domain's language and debates doesn't automatically extend to another's. And since interactional expertise is not the same as contributory expertise, there is a boundary of sorts between those two as well — one that matters even when the interactional expertise is genuine and valuable.

Expertise does not exist in isolation. The expert, the domain with its feedback characteristics and community of practice, and the people and institutions that depend on expert judgment — patients, policymakers, practitioners, students — form a system. Knowledge moves through that system: from domain to expert, from expert to consumer, from consumer back to practice. When that system works well, expertise gets tested, calibrated, and applied appropriately.

Failure Modes

When the boundaries between domains and between contributory and interactional expertise are respected, expertise functions well. When they aren't, the failures are recognizable. Some failures are about the system — experts, institutions, and audiences — and how knowledge moves between them. Others are about individual experts and the domains in which they operate. Still others are about the expert alone.

Capture

Some failures happen in translation. A finding emerges from careful work, moves into public discourse, and arrives stripped of the qualifications that made it careful. The story that results is simpler, more universal, and more actionable than the evidence supports — and once it finds institutional homes, it becomes extraordinarily difficult to dislodge.

Rudolf Schenkel's 1947 observations of captive wolves — unrelated animals confined together in an artificial environment — produced a genuine finding: under those conditions, dominance hierarchies emerge. The first inferential leap came quickly: captive wolf behavior was taken to describe wolf behavior generally. David Mech's 1970 book popularized this version, and the concept of the alpha wolf entered the culture. The second leap followed from the first: since dogs descended from wolves, wolf social structure must apply to dogs. The dog training community made that inference without much examination. An entire approach to training built itself around the result.

Mech spent decades trying to correct the first leap. His field research from 1986 onward showed that wild wolf packs are family units, not dominance hierarchies. In 1999 he published a corrective paper. He asked his publisher to take his 1970 book out of print. The publisher refused — it was still selling. The scientific community updated within years. The dog training industry has been slower still. A decade after the corrective paper, Cesar Millan built a television career on the original narrative.

The spin step — the transformation of a careful finding into a universal story — is not always cynical. Often it is simply what happens when research meets institutional demand for actionable guidance. The finding says: under these specific conditions, we observed this. The story says: here is what you should do. The distance between those two things is the distance narrative capture travels. Once training schools, certification programs, and television franchises have built themselves around it, the correction cannot simply be published.

The pattern recurs across domains. The relationship between dietary fat, cholesterol, and heart disease. Ten thousand hours to mastery. The role of learning styles in education. In each case a careful finding traveled into public discourse as a universal claim, found institutional homes, and proved resistant to the corrections that followed. While the details may differ, the mechanism is the same.

Migration

Some failures are about individual experts and the nature of the domains in which they operate. Dr. John Campbell — the title earned for a PhD in nurse education — built a substantial following as a YouTuber in the early days of the COVID pandemic, explaining complex health data clearly to a lay audience. That facility — genuine interactional expertise in health communication — was real and valuable. But as his platform grew, so did his ambition. His analysis moved progressively into territory requiring contributory expertise he did not have: epidemiology, biostatistics, immunology. For years his audience grew while his claims moved further from his actual competence. When the BBC eventually debunked his excess mortality analysis, he took the video down and acknowledged he was not a statistician. By then his channel had accumulated half a billion views.

The pattern is recognizable. Campbell did not engage with the established communities of practice in epidemiology or biostatistics — the peer review, the collegial challenge, the standards for what counts as reliable statistical inference. He didn't need to. The audience rewarding him could not evaluate what he was doing and had no reason to. The incentive was the platform, the following, the reception that staying in nurse education would never have provided. In low validity domains — where the feedback on the correctness of expert conclusions is slow, ambiguous, and easily contested — this kind of migration can persist for a very long time before it meets any meaningful correction.

A related boundary failure stays within the domain. Genuine interactional expertise — the facility with a domain's language, debates, and standards that develops through sustained engagement with its practitioners — is real and valuable. But it has a boundary too. The senior manager who understands enough to ask the right questions, coordinate specialists, and evaluate competing claims is doing something legitimate and often essential. The same manager who steps into the specialist's territory — overrides the technical judgment, makes the call that requires doing the work rather than understanding it — has crossed a boundary that interactional expertise doesn't support. The credentials and the fluency are real, but the underlying capability may not be there.

Calibration

Finally, there are failures of expertise that are entirely individual in nature. The expertise may be real, but the practitioner doesn't accurately assess its limits.

Genuine expertise produces calibrated confidence. The expert who has spent years being corrected by a domain develops an accurate sense of where their judgment is valid and where it isn't. The confidence is specific, bounded, honest about uncertainty. True experts, as Kahneman and Klein observe, know when they don't know — not as a matter of temperament but as a product of having been taught by the domain itself.

That calibration requires two things. First, feedback clear and timely enough to distinguish reliable judgment from unfounded confidence. In high validity domains — surgery, structural engineering, chess — the domain provides this directly. The approach works or it doesn't. In lower validity domains, where outcomes are slow, ambiguous, and hard to attribute, the domain can't do that work alone. Second, immersion in a community of practice that has developed, over time, shared standards for what counts as reliable judgment, what warrants confidence, and where uncertainty should be expressed. That community provides the ethical as well as the cognitive dimension of calibration — the commitment to intellectual honesty that genuine expertise requires.

Where feedback and community are both weak or absent or just ignored, confidence can expand beyond what the domain has actually validated. The archetype is the “minor expert” — someone with real but narrow skill who has mistaken that narrowness for comprehensiveness. Not a fraud. Not a boundary crosser. Just a practitioner whose confidence was never corrected to match what they actually know. These tend to accumulate where barriers to entry are low, feedback is weak, and the community of

practice is loose enough to impose few constraints on confidence claims. The dog trainer who knows one technique and applies it to everything. The nutritionist whose genuine knowledge of one metabolic pathway has become a theory of all things dietary. The financial advisor whose success in one market condition has become universal confidence.

The same failure appears at the other end of the expertise spectrum. Where the minor expert expands outward beyond what narrow knowledge supports, deep competence can turn inward — confidence that stops questioning rather than confidence that overclaims. In 1956, Archie Cochrane — a physician and epidemiologist who would become one of the founders of evidence-based medicine — was referred to a surgical specialist after a lesion was found under his arm. The specialist operated, found what he determined to be cancerous tissue extending deep into the chest, removed an entire muscle, and told Cochrane he had limited time to live. The pathologist's report had not yet come back. Neither man waited for it, nor questioned whether they should. Years of competence had produced a confidence that no longer stopped to question itself. When the report arrived, the tissue was not cancerous. Cochrane later reflected that he had never doubted the surgeon's words.

The minor expert and the surgeon are both failures of contributory expertise — its scope misjudged in one case, its certainty in the other. A parallel failure applies to interactional expertise. The person who has absorbed enough surface language to participate in the conversation without genuine community immersion may not know the difference. They can put their feet in the painted footsteps. They are not dancing. In low validity domains this failure can persist without correction for a very long time.

None of these expertise failure modes is new. What is new is the environment they now inhabit.

Amplification

Generative AI has arrived as an “overnight success” built on decades of building blocks that promised more than they were ready to deliver. What it does extraordinarily well, right now, is produce at least the surface appearance of expertise — fluent, confident — in any field, on demand, at essentially zero marginal cost.

Each expertise failure mode responds to this differently. The potential for generative AI to disrupt expertise is greatest in low validity environments, and increases as the work product becomes less formally constrained and less physically instantiated. Generative AI operates entirely in the domain of language and representation. It cannot weld. It cannot perform surgery. The further the work product is from physical reality and formal constraint, the more convincingly AI can simulate it. Purely expressive work — recommendations, narratives, conversations — is easy to simulate with generative AI. Moderately formal work, which mimics the appearance of rigor without requiring its underlying competence, is now considerably easier to fake convincingly. Genuinely formal work — where the product must meet structural constraints before the content is even evaluated — remains harder, though the boundary is moving.

Spin

Narrative capture is where the AI amplification story is most specific. The spin step — the mechanism that turns careful findings into sticky universal stories — previously required specific conditions. A talented popularizer. An institutional appetite for simple actionable guidance. A receptive public moment. The distance between what Schenkel observed in a Swiss zoo and what Cesar Millan was teaching on American television took decades and several human intermediaries to travel.

Generative AI, in conjunction with social media, makes the spin step available at zero marginal cost, in any domain, at any scale, to anyone. Whether the resulting narratives achieve the kind of institutional entrenchment that the alpha wolf story did remains to be seen — that takes time and repetition. What has changed is the supply side. The conditions that previously limited how many findings could be spun into universal stories are less relevant.

Social media is the distribution mechanism. The algorithm selects for exactly the properties the spin step produces — simplicity, universality, emotional resonance — and against the properties of the underlying research — qualification, boundary conditions, uncertainty. What previously took decades and the right combination of human intermediaries can now be attempted in an afternoon. The platform can only see conversance. The product is the post.

Acceleration

The same logic applies to domain migration, which gets considerably cheaper. Before generative AI, moving into new territory required developing some real facility there — specifically, enough interactional expertise to perform credibly to a non-expert audience. Not contributory expertise. But real preparation, real investment, a filter however imperfect. And the natural correction mechanism — the domain expert encounter, where the migration eventually became visible — still functioned, slowly.

Generative AI removes the preparation cost of interactional expertise entirely in written contexts — the essay, the blog post, the social media thread — all fully fakeable at zero cost, in any domain, on demand. The within-domain migration gets easier too. The manager who wants to override the specialist can now generate fluent technical arguments on demand, indistinguishable in written form from genuine interactional expertise developed through years of community immersion. Live performance — where the expert must respond to unpredictable questions without a net — remains a friction point. But nobody gets a platform without a prior written presence. That written presence is now fully AI-assisted. The live appearance is downstream of infrastructure that AI has eroded.

Social media compounds this. The corrective community — the domain experts who would recognize the migration — is present but is not the audience being addressed. The platform audience self-selects for the heterodox credentialed voice. The correction, when it comes, arrives as further evidence of establishment resistance.

Simulation

The self-knowledge failure amplification story is the most consequential. Genuine expertise requires two things that AI can undermine simultaneously: feedback from a domain that corrects, and socialization into a community that calibrates. In high validity domains the first was always the primary mechanism. In low validity domains the second was what made genuine expertise possible at all.

The two self-knowledge failures respond differently. The minor expert — whose narrow contributory expertise was always real — is not directly affected. What AI makes possible is something worse: brand new minor experts with no underlying domain contact at all. In low validity domains where the work product is expressive, AI-generated fluency in a narrow area can be deployed with the same confidence as genuine narrow expertise, and neither the practitioner nor the audience may be able to tell the difference.

The painted footsteps failure gets dramatically worse. Pre-AI, developing even surface interactional fluency required real engagement — sitting through enough meetings, absorbing enough of the community's language to participate convincingly. That engagement, however superficial, was a filter. AI removes it entirely. The footsteps are now provided on demand, with no engagement required.

What replaces community socialization in AI systems is a different kind of formation entirely — one shaped by the design decisions of platform owners rather than by communities of practice. The practitioner using AI to develop or supplement their expertise is being socialized, but not by their domain.

There is a further problem. LLMs are trained on existing content — the accumulated outputs of socialized communities, representing the domain as it was. The emerging challenge to the consensus, the anomaly that doesn't fit, the finding that the expert community is quietly debating — these are statistically invisible in the training data. The genuine expert notices the slight movement against the static background. The LLM reproduces the background. The practitioner using AI to develop interactional expertise in a domain gets yesterday's consensus delivered with today's confidence.

Social media amplifies both problems. The platform selects for confident assertion over calibrated uncertainty, for the closed story over the open question, for the consensus that travels over the correction that complicates. It also rewards the maverick — the provocative, contrarian voice that challenges established expertise. The migrating expert with confident credentials and the AI-assisted simulator who doesn't know what they don't know both find in social media an environment that rewards exactly what genuine expertise would restrain.

Patience

Generative AI does not create these failures. It amplifies them — making each more likely, more frequent, and harder to detect. And yet genuine expertise does not disappear. It becomes harder to see, and harder to develop, but it does not disappear.

The person who has spent years being corrected by a domain has something that fluent output cannot replicate and that fluent output cannot fully disguise the absence of. The knowledge has a texture that comes from accumulated contact with things that don't cooperate with the theory. The system that fails in ways the model didn't anticipate. The patient who doesn't respond the way the narrative said they would. Each is a small correction, a slight adjustment, a revision to the internal model that only happens through genuine contact with a domain that doesn't care about confidence.

That accumulated responsiveness is what expertise actually is. Not the credential. Not the fluency. Not the framework, however elegant. The thing built through years of being wrong in instructive ways, by a domain patient enough to keep correcting.

Beyond the question of recognizing genuine expertise, there is a deeper concern. Generative AI does not merely simulate expertise — in low validity domains where the work is expressive and the feedback slow, it may undermine the conditions through which genuine expertise develops in the first place. The practitioner who delegates the foundational work never struggles with it. The domain never gets a chance to correct them.

The amplification effects of generative AI described here are plausible — but an environment in which genuine expertise becomes permanently indistinguishable from its simulation is not inevitable. The future has a habit of confounding our predictions about it, technological disruption included. The response already underway in academic institutions — rethinking how expertise develops, what AI-assisted work means for learning, how to preserve the developmental pathways that genuine competence requires — suggests that at least some of the adaptation will be deliberate rather than accidental. Whether it moves fast enough, and reaches broadly enough, is an open question.

Strangers sometimes ask how the dog walking quietly beside me came to be that way. It takes a lot of time and dedication, I tell them, but fortunately for me, Henry is a very patient teacher.