



Why AI Training Fails

A Three-Layer Framework for SME Adoption

Chris Duffy

Founder and Chief AI Officer, IgniteAI Solutions Ltd

August 2026

Executive Summary

Based on my experience, most UK SMEs are on their second or third attempt at AI adoption now in August 2026. The Copilot licences have been purchased for 12 months, the ChatGPT enterprise seats got rolled out. The Lunch & Learn or “AI Day”. Three months later, a handful of people are using it well, most people have drifted back to how they worked before, and the investment case is looking thin.

The industry tells a story about why this happens. The story is usually about tools, training, or change management. The story is incomplete.

AI adoption fails because it sits across three layers at once, and most interventions only address one of them.

The first layer is neurological. Under pressure, the emotional brain overrides deliberate practice. The habits that fire fastest are the ones that win, and the habits people come in with are search-engine habits.

The second layer is cognitive. Without a reframe of what the tool actually is and how it should be used, employees default to treating it like Google with better manners.

The third layer is organisational. Individual behaviour change does not survive contact with a team that has not been sequenced through it.

Address one layer, the other two collapse your effort. Address all three, the adoption rate compounds.

This paper synthesises three bodies of published work, credited to their authors. Professor Steve Peters' [Chimp Model](#) explains the neurological layer. Conor Grennan's [AI Mindset](#) method, used as the cognitive training arm in Microsoft Research's 2026 field experiment at Gap Inc., provides the cognitive layer. Prosci's [ADKAR Model](#), now being applied to AI adoption by Prosci themselves, provides the organisational sequencing.

The synthesis is mine. The application to UK SMEs, through the fractional Chief AI Officer model I run at IgniteAI Solutions, is mine. The case studies are from live client work.

If you are an MD, Ops Director, or founder looking at an AI investment that is not landing, or about to make one and trying to get it right first time, this paper describes why most of the advice in the market is addressing the wrong layer. The section on what this means for your business is the one to read first.

The adoption failure pattern

The headline failure statistics are consistent across every major research body. Gartner puts 80% of AI projects failing to scale beyond pilot. MIT Sloan found 16% of AI pilots achieve full enterprise rollout. BCG's 2024 work put the figure at 74% of enterprises failing to scale AI pilots. McKinsey's 2025 survey found that over 80% of companies using generative AI report no material earnings contribution from it.

The Institute of Directors' research is sharper for UK SMEs specifically. 92% of SME leaders recognise AI's potential. 16% feel prepared to implement it. The gap between recognition and readiness is where the investment is getting wasted.

Forrester's 2026 AIQ research lands in the same place from a different angle. J.P. Gownder, Forrester's principal analyst on the future of work, [reported in March 2026](#) that the proportion of employees who know what prompt engineering is and how to use it moved from 22% in 2024 to 26% in 2025. A 4-point increase over a year that saw tens of billions of pounds poured into enterprise AI rollouts. Gownder's conclusion is direct. Training alone is not sufficient to raise AIQ. The problem is not the employee. The problem is the employer who has not built a learning and engagement environment that actually changes behaviour.

Something more useful sits underneath these numbers. AI adoption does not just fail. It fails in a pattern.

The pattern goes like this. Tools get bought or rolled out. A training session happens. For a few weeks, usage is high because novelty is high. Then usage drops. A few champions keep going. Most people quietly revert to how they worked before, either because the tool did not save them as much time as promised, or because the outputs were not as good as they hoped, or because the method they were taught was shallow enough that they stopped getting value from it once the easy use cases ran out.

The J-curve

The productivity data confirms this. MIT's research on AI adoption in US manufacturing firms, published in 2025 by [McElheran, Yang, Kroff and Brynjolfsson](#), mapped the full adoption trajectory against output. The pattern is a J-curve. Productivity drops first, then recovers, then rises. The initial drop was measurably larger in established firms, where legacy processes created more friction with the new tools. Companies that did not push through the dip did not get to the upside.

The mapping problem

There is a separate piece of evidence that explains why even organisations with good tools and good training struggle to convert investment into value. [Kim, Kim and Koning](#), in a 2026 field experiment published through INSEAD and Harvard Business

School, randomly assigned high-growth firms to two arms. Both had AI access. The treatment arm received workshops that helped them map where AI could fit across their business functions. The treatment firms discovered 44% more use cases, completed 12% more tasks, generated 1.9x the revenue, and had 39.5% lower capital demand. The tools were the same. The mapping was different.

Access to AI does not produce value. Method around AI produces value. The method has to work at multiple levels, because the friction points sit at multiple levels.

The easy interpretation of all this is “we need better training.” That is the interpretation the market tends to sell. It is not wrong, but it is not complete. Better training addresses one of three layers. If you only address one, the other two will drag your adoption rate back down.

The three layers are neurological, cognitive, and organisational.

The neurological layer

The Chimp Model

Professor Steve Peters is a consultant psychiatrist with over thirty years in the clinical field. He has worked with British Cycling, Liverpool FC, the England football team, Ronnie O'Sullivan, Sir Chris Hoy, and multiple Olympic teams. [The Chimp Paradox](#), published in 2012, has sold close to two million copies and remains one of the best-selling self-help book in UK history. Peters' [Chimp Model](#) is now supported by NHS-backed research trials. It is the most rigorously applied psychological framework in elite performance.

The model is a way of describing how three parts of the brain work together. Peters is explicit that it is a model, not a hypothesis. It simplifies neuroscience into something operational.

The Human brain is the frontal lobe. Rational. Logical. Evidence-based. Weighs options, considers consequences, plans ahead. Slow.

The Chimp brain is the limbic system. Emotional. Impulsive. Runs on feelings and impressions. Fires fast, jumps to conclusions, fills in gaps with emotion rather than evidence. Five times faster than the Human.

The Computer is the parietal region and associated systems. It stores beliefs, habits, and automatic responses. Both the Human and the Chimp use it as a reference library. Most importantly, the Computer runs roughly four times faster than the Chimp and twenty times faster than the Human. When the Computer has a pattern laid down, that pattern fires before either the Human or the Chimp gets a look in.

Peters' central insight is that under pressure, the Chimp wins. The Human is slower to engage, the Chimp is already reacting, and the Computer has fired an autopilot before anyone has thought about it. The work of managing your mind is the work of nurturing the Chimp, exercising the Human, and deliberately programming the Computer so the autopilots that fire are the ones you actually want.

Applied to AI use

A well-briefed prompt is Human work. It involves context, framing, examples, expected outputs, and a sense of what good looks like. It takes effort. It is slow. It produces good results.

A reflex prompt is Chimp work. It is fast, emotional, shaped by whatever impression the user has of what AI is. For most users, that impression is "it is a search engine that talks back." So they type a search-engine query, receive an output, judge it quickly,

and move on. The judgement is emotional rather than analytical. If it feels right, it gets accepted. If it feels wrong, the tool gets dismissed.

The Computer is where most of the problem lives. Every time a user interacts with AI in a particular way, they are reinforcing a pattern. The pattern compounds. By the time someone has used ChatGPT reflexively for six months, they have a Computer that fires search-engine prompts faster than they can consciously choose a better approach. Asking them to prompt differently at that point is not a training issue. It is a reprogramming issue.

The Lee finding

A recent study from Microsoft and Carnegie Mellon gives this a hard empirical edge. [Lee et al.](#) published at CHI 2025, surveyed 319 knowledge workers across 936 real-world AI tasks. They found that higher confidence in AI was associated with less enacted critical thinking. Higher self-confidence was associated with more. In Peters' language, the first pattern is the Chimp running the Computer unchecked. Confidence in the tool suppresses the Human's analytical check. The output gets accepted on feel, not on evidence.

Why users abandon the J-curve

The J-curve dip also becomes readable through Peters. When a user first starts using AI seriously, they experience a short period where everything feels slower, harder, and more frustrating than the old way. This is the Chimp reacting to loss of competence. The Chimp hates uncertainty, fears incompetence, and will push the user back towards the familiar autopilot of how they used to work. Most users abandon the tool at exactly this point. The few who push through are the ones who have either an external motivator, a well-designed social environment that reinforces the new behaviour, or enough Human discipline to override the Chimp's resistance.

The Troop

Peters has a concept called the Troop. The Troop is the social context. In the natural world, the Chimp survives by being part of the Troop, and the Troop shapes which behaviours are reinforced and which are punished. In business terms, the team around the user either reinforces new AI habits or lets them die. A Troop that normalises reflex prompting will pull even a well-trained user back into reflex prompting. A Troop that normalises deliberate AI use will pull even a weaker user upwards.

Goblins and Gremlins

Peters also describes two categories of negative pattern stored in the Computer. Goblins are destructive autopilots that can be overwritten with effort. Gremlins are deeply embedded, often from childhood, and are much harder to shift. In an AI

context, a Goblin is the search-engine-style prompting habit picked up in the last two years. It is overwritable, with the right intervention. A Gremlin is something deeper, such as a fundamental belief that “I am not a technical person” or “I don't trust machines to do important work.” These need different, slower interventions.

None of this is speculative. It is an application of a published model that has been used in elite sport, clinical psychiatry, and NHS-backed trials for over a decade. What is new here is the mapping to AI.

Cognitive training without neurological understanding will produce short-term behaviour change that evaporates under pressure. The Chimp overrides the new pattern. The Computer fires the old autopilot. The Troop tolerates the reversion. Three months later, the adoption numbers are where they were before.

This is why so many well-designed AI training programmes deliver weak long-term adoption. The training is correct. The layer it addresses is incomplete.

The cognitive layer

Grennan's AI Mindset method

Conor Grennan is the CEO of [AI Mindset](#) and former Chief AI Architect at NYU Stern. His work has been used as a training intervention at OpenAI, Microsoft, Google, McKinsey, NASA, JP Morgan, Unilever, and the US Department of the Treasury. In 2026, Microsoft Research and Gap Inc. ran a field experiment at Gap where his curriculum was the cognitive scaffolding arm of the study.

The study, [Farach et al.](#) published on arXiv as 2604.08678, found that participants who received Grennan's training had directionally better outcomes on a strategic communications task than participants who received standard feature-based Copilot training. 77% of the trained group produced top-quality documents against 62% of the control group. I want to be careful with this finding. The pre-specified continuous outcome in the study was null. The 77-versus-62 headline comes from an exploratory binary analysis so I treated it as directional, not definitive. The study's stronger finding is that rigid, mandated workflow protocols reduced document quality by five points on a 22-point scale and made teams over eight times less likely to complete the task.

The three traps

What Grennan teaches is not about tools. It is a cognitive reframe of what AI is and how it should be engaged.

His central move is to reframe the user's mental model of AI. Most people treat AI as a search engine, a software feature, or a tool to pull out for specific tasks. Grennan teaches it as a thought partner, operationalised through the metaphor of a smart intern. You brief it. You give it context. You iterate. You correct.

He names three traps that stop users getting value. The Search Engine Mindset is treating the tool transactionally, one query for one answer. The Use Case Trap is thinking of AI as a shelf item you pull out for specific tasks, rather than a new way of working across your whole job. The Big Red Button Trap is expecting a perfect finished artefact from a single prompt.

The four habits and the G-CEO framework

The behavioural habits he teaches are specific. His workhorse prompting framework is G-CEO: Goal, Context, Examples, Outputs. He teaches a three-part work pattern he calls Learn, Execute, Strategize. Learn is using AI to build understanding of complex topics. Execute is using it for everyday workflow. Strategize is using long, iterative sessions on back-burner, high-value work. Each has its own prompting cadence.

He also teaches specific daily habits, some of which are counterintuitive enough to be distinctive. He advocates for starting every session with a greeting, because the greeting shifts the user's cognitive register away from query-mode and towards conversation-mode. The change is in the user, not the AI. He advocates for asking the model what question to ask next, for role-ing in an expert to critique the output, and for pushing back on first drafts rather than accepting them. He is explicit that the first draft is almost never the final draft, and the iteration is where the value lives.

For responsible use, he teaches an acronym called R.E.S.P.E.C.T., covering hallucinations, data protection, and critical review.

Where Grennan's methodology can be strengthened

When I map Grennan's method against Peters' model, the fit is clean. Grennan's four habits are Human brain work. They require deliberate effort, context, iteration, and critique. What Grennan does well, and what Peters explains why it works, is give the Human brain specific moves to make. What Peters adds, and what Grennan's method alone does not address, is why these habits break down under pressure and what needs to be in place in the Computer and the Troop to stop that happening.

This is the fundamental limit of cognitive training alone. You can teach someone G-CEO on Tuesday. On Wednesday morning under deadline, the Chimp is already typing a search-engine query before the Human has remembered the framework.

Grennan's work is the right content for the cognitive layer. The cognitive layer is not the whole problem.

The organisational layer

ADKAR

Prosci's [ADKAR Model](#) was developed by Jeff Hiatt in the late 1990s, based on studying over 700 organisations going through change. It is now used by thousands of change management practitioners worldwide. Its insight, and the reason it has outlasted dozens of competing frameworks, is that organisational change only happens if individual change happens. Teams do not change. People change, and teams change because enough people changed.

ADKAR is five sequential milestones every individual has to pass through for a change to stick.

Awareness is understanding why the change is necessary. Not awareness that a change is happening, which is a different thing. Prosci's own research consistently shows lack of awareness as the number one reason for employee resistance. People resist because nobody made a compelling case for why.

Desire is personal motivation to participate in and support the change. Awareness without desire produces compliance, not adoption. People show up to the training with their arms folded.

Knowledge is understanding how to change. The content of what needs to be learned. This is where most AI training programmes live, and where most fail, because training is being delivered without the prior two milestones having been established.

Ability is the demonstrated capacity to actually perform the new behaviour under real working conditions. Training produces knowledge. Practice produces ability. Prosci's own work emphasises that the gap between the two is where adoption silently dies.

Reinforcement is the organisational context that keeps the new behaviour in place after the novelty wears off. Recognition, management feedback, measurement, and consequences. Without reinforcement, the old autopilot wins.

Prosci on AI adoption

[Prosci has published their own work](#) on applying ADKAR to AI adoption. Their research points at the specific barriers at each stage. Awareness barriers cluster around fear, uncertainty, and concerns about job security. Desire barriers cluster around the “what's in it for me” question, particularly when leadership frames AI as a cost-reduction tool rather than a capacity-building one. Knowledge barriers emerge when training is generic rather than role-specific. Ability barriers develop when there is no practice environment between training and real work.

Mapping ADKAR to Peters and Grennan

Awareness and Desire are Chimp work in Peters' terms. The Chimp decides whether to engage before the rational brain gets involved. You cannot reason an emotional resistance into submission. You have to nurture the Chimp first. This is why poorly run AI announcements backfire. Leadership makes the case for AI efficiency, the workforce hears “they want to reduce us,” the Chimp reacts with loss-aversion, and the Desire stage of ADKAR is dead before training begins.

Knowledge is where Grennan's cognitive method lives. It is the content of what people need to be taught. The three traps, the four habits, the G-CEO framework, the Learn, Execute, Strategize pattern. This is the teachable material.

Ability is the convergence zone. It is where Grennan's method and Peters' Computer work meet. Knowing the habits is not the same as having the Computer fire the habits under pressure. Ethan Mollick, professor at Wharton and author of Co-Intelligence, has argued that most users need roughly ten hours with a frontier model before they stop being beginners. This is not a training threshold. It is a Computer-reprogramming threshold. The new autopilots have to be laid down through repetition in real working conditions.

Reinforcement is Troop work. It is the organisational system that either locks in the new Computer pattern or allows the old Goblins to fire back. Management feedback, peer norms, visible champions, measurement of the right behaviours, consequences for backsliding. When the Troop reverts, the individual reverts, almost without exception.

Rolling out AI by sending everyone to a half-day training is a structurally failed strategy. The training might hit the Knowledge milestone. It does almost nothing for Awareness, Desire, Ability, or Reinforcement. The four milestones where the adoption actually lives or dies.

Prosci's original research found that when people progress through all five ADKAR stages, change sticks. When they skip stages, it does not. Their data on this is decades deep. The model is not speculative. What I am adding is the mapping to Peters and Grennan, because the ADKAR literature tells you that the stages matter, and Peters and Grennan tell you what specifically has to happen at each stage for AI adoption in particular.

The synthesis

Three bodies of published work. Three layers. Each necessary, none sufficient on its own.

Peters explains why AI adoption is difficult at the individual neurological level. The Chimp resists, the Computer runs on old patterns, the Human is too slow to win the race. This is the layer that explains why even well-trained users revert.

Grennan explains what the new pattern should be. The cognitive reframe, the habits, the specific moves that separate a user who extracts value from AI from one who does not. This is the content layer.

ADKAR sequences how the change gets embedded organisationally. Awareness and Desire in the emotional brain before Knowledge in the rational brain, then Ability through deliberate practice, then Reinforcement through the organisational context. This is the sequencing layer.

None of these sources contradicts the others. They address different layers of the same problem. My contribution is to hold all three in one frame and apply them together to a specific context. That context is UK SMEs of 20 to 500 staff, £5M to £50M turnover, where the leadership is close enough to the work that a well-designed intervention can propagate through the whole business in weeks rather than quarters.

The framework table

ADKAR stage

Peters layer

Grennan layer

The work

Awareness

Chimp nurture

Not yet active

Emotional preparation. Address fear, uncertainty, and the 'am I being replaced' question before anyone is asked to do anything.

Desire

Chimp

Starting to seed

Make the case that lands emotionally, not just rationally. Show people what they gain, not what the company saves.

Knowledge

Human exercise

Fully loaded

Teach the cognitive reframe and the behavioural habits. The AI Mindset content applied to the chosen tool.

Ability

Computer programming

Practice

Hours of deliberate, supervised, real-work practice until the new autopilots are laid down. Ten hours minimum, with feedback.

Reinforcement

The Troop

Habit maintenance

Build the organisational structure that keeps the new behaviours alive. Champion networks, governance, measurement, consequences for reversion.

A company can be in Awareness on one team and Ability on another. Individuals within the same team can be in different stages. The framework is a diagnostic before it is a plan. Before you intervene, you find out where each individual and each team actually is, and you address the bottleneck rather than the symptom.

There is also a premium application of this framework for senior leaders. The principles are the same. The calibration is different. A founder or MD needs a version of this tuned to how they personally think, their failure modes, their triggers, and their decision-making patterns. I call this a Personal Operating Protocol. It is bespoke, intensive, and reserved for fractional CAIO engagements.

How I apply this at IgniteAI

At IgniteAI Solutions, I deliver this framework through a set of productised and bespoke services that map to the three layers and the five ADKAR stages. The architecture is deliberate. Each service addresses a specific part of the adoption failure pattern. The whole system is designed so the pieces reinforce each other.

SPARK: the readiness diagnostic

SPARK is the readiness diagnostic. Eighteen dimensions, run over two weeks. It is built to identify where in the ADKAR sequence each part of the business actually sits, and which of the three layers has the biggest gap. Some businesses come to me with a pure Awareness problem masquerading as a training problem. Some have the opposite, a workforce ready to use AI and a leadership team still in Awareness. SPARK surfaces that before anyone commits to a specific intervention.

Discovery: the use-case mapping

The Discovery process sits immediately after SPARK when the finding warrants it. This is the use-case mapping work that the [Kim, Kim and Koning study](#) at INSEAD and Harvard identified as the single highest-value pre-investment activity. Treatment firms in their study discovered 44% more use cases and generated 1.9x the revenue of control firms. The discovery process I run is built on a specific nine-advisor adversarial review protocol, with evidence grading, split-panel scoring, and a pre-mortem stage. It sits between Awareness and Knowledge in ADKAR terms. It is how you figure out what specifically to teach before you try to teach anything.

Claude Built For You: the cognitive layer

Claude Built For You is the cognitive layer delivery. This is Grennan's method applied specifically to Claude, shaped for UK SMEs, delivered as a productised onboarding programme. Module 1 teaches the three traps, the four habits, and the Learn, Execute, Strategize pattern. Modules after that apply it to the real work of the specific business. The cognitive content is attributed to Grennan and AI Mindset. The Claude-specific application, the UK SME shaping, and the productised format are IgniteAI's.

Fractional CAIO: the oversight layer

The fractional CAIO engagement model is the oversight layer. I work with SME leadership teams on a retainer basis, providing the strategic context, the methodological rigour, and the independent challenge function that most SMEs cannot justify building in-house. This sits across all five ADKAR stages. It is where the Troop-

level work happens at senior level, because the leadership team is itself a Troop that either reinforces or undermines whatever is being rolled out below them.

AI Governance: the reinforcement layer

AI Governance is the Reinforcement layer. This is the organisational infrastructure that keeps the new behaviours alive after novelty wears off. Acceptable use policies, validation matrices, human-in-the-loop protocols, champion networks, governance forums, and measurement frameworks. Without this, the Goblins fire back. With it, the new autopilots stabilise.

ENGAGE, IGNITE, and the frameworks I credit

The ENGAGE and IGNITE frameworks sit as structured methods for leadership alignment (ENGAGE) and implementation execution (IGNITE), running in parallel with Nufar Gaspar's C-H-A-N-G-E cultural operating system, which I credit explicitly to her every time I use it. Dr Markus Schmidberger's AI-EAF maturity measurement framework sits alongside, again credited to him. I do not take authorship of work that is not mine. I build the system around it.

The SME J-curve is steeper

One thing that matters for UK SMEs specifically. The J-curve productivity dip is steeper in smaller firms because the ratio of early adopter to total workforce is less forgiving. If you have 2,000 staff and 100 are struggling during the dip, the rest of the business continues. If you have 40 staff and 15 are struggling, the business is struggling. This makes the Reinforcement layer more important, not less, in an SME context. The Troop is tighter. Reversion happens faster. The cost of getting the emotional preparation wrong at the Awareness and Desire stages is proportionally higher because there is less slack in the system.

Where the experience comes from

The experience I draw on for this work has two sources. One is direct. Twenty-three years in the British Army, including six years with 22 SAS, and most recently as Technical Programme Manager on the UK MoD's [Project ASGARD](#) in Q1 2025, the British Army's AI-enabled digital targeting web, delivered in six months against a two-year norm. Defence does not get a J-curve. The cost of adoption failure is not measured in P&L. That context shapes the discipline I bring to commercial work, and the posture I take with clients who are underestimating what rigorous implementation requires.

The second source is the current client base. The next section covers four of those engagements and what they illustrate about the framework in practice.

What this looks like in practice

Hart's of Stur

Hart's is a UK cookware retailer established in 1919. The problem they came to me with was a manual product-listing process that was eating their team's week. Before we started, a full product listing took hours of manual work per SKU, across research, copywriting, image specification, and metadata. The bottleneck was limiting how fast they could expand their catalogue.

We built an automated pipeline, branded as FORGE, that handled the listing workflow end to end. The measurable outcome, verified against their baseline, is a 94% reduction in time spent on the listing process. Not on revenue. Not on conversion. Specifically on the time cost of the listing workflow. The FORGE engagement was calculated as a 400% productivity gain in the product listing workflow.

What Hart's illustrates is the Ability and Reinforcement layers in a bounded, high-frequency task. The team was ready in Awareness and Desire terms. The Knowledge was specific to the workflow, not general. The productised automation is what carried the Ability, and the fact that it runs every day is the Reinforcement. No Chimp resistance, because nothing is being replaced that the team wanted to do. The drudgery went. The team focused on the work that actually needed judgement.

Woodhouse Surveys

Woodhouse Surveys is a surveyor I worked with on a Claude Built For You proof of concept. The problem they had was a bespoke workflow around a specific part of their reporting process that was manual, slow, and repetitive. The engagement was built specifically to prove the cognitive layer of the framework in a real, small-team environment, with real work, over a short period.

The outcome was a working Claude setup, tuned to their actual process, that the team was using independently inside two weeks. What Woodhouse illustrates is that when you deliver the cognitive layer properly, with the right tool setup, a small team can cross the J-curve quickly. The full J-curve dip never really materialised because the onboarding addressed the emotional and Knowledge stages in parallel. The Troop was small enough that one champion was enough to carry the rest.

KASP Solutions and Informa PLC

KASP Solutions is a training and consulting firm that delivers to Informa PLC and other enterprise clients. I worked with them on AI training delivery specifically for their Informa engagement. The problem was not their technical understanding, which was strong. The problem was translating AI capability into a methodology that a large, distributed enterprise workforce could actually adopt.

What KASP illustrates is the organisational layer at a larger scale, delivered through a partner channel. The cognitive content had to be compressed without losing rigour. The organisational sequencing had to account for a workforce that did not work for me and whose Troops I could not directly shape. The mechanism for landing this was to equip KASP's own trainers with the framework, so they carried it into the Informa environment with integrity, rather than trying to deliver it myself in a context I did not know well enough.

A Security Cleared defence contractor

The fourth client in this set is a Security Cleared defence contractor I have been supporting on AI transformation. The nature of the work means I cannot describe it in detail, and the client's name is not public. What I can say is that the framework in this paper is being applied in an environment where the Reinforcement layer is formally mandated through governance, where Awareness and Desire must be managed carefully because of the sensitivity of the domain, and where the Ability stage requires substantially more oversight than in a commercial setting. The defence context is where this framework originally clarified itself for me, because the cost of adoption failure in defence is measured in operational capability, not quarterly revenue. What works in defence works commercially. The reverse is not always true.

Four engagements, four different shapes of the same framework. The layers are the same. The application varies with the client, the problem, and the sector. That is what a framework is for.

What this means for your business

If you are running an SME and looking at AI investment, or looking at one you have already made that is not producing the returns you expected, a few things follow from the framework in this paper.

Diagnose first

The first thing is diagnostic. Before you buy more tools, run more training, or hire more consultants, figure out which layer is actually broken. Most SMEs I see have the layers mixed up. They describe a Knowledge problem and it is actually a Desire problem. They describe an Ability problem and it is actually a Reinforcement problem. The symptoms look similar from the outside. The interventions are completely different.

The financial cost of misdiagnosis is meaningful. An average UK SME rolling out Copilot or ChatGPT at scale is looking at licence costs running into tens of thousands of pounds a year. Layer on training and implementation costs and the total investment is easily £50K to £150K for a mid-sized business. If the investment produces a year of patchy adoption, you have not just wasted the money. You have burned the workforce's willingness to try the next thing. Second attempts at AI adoption are

harder than first attempts, because the Chimp is now scarred. Desire at the second attempt is starting below zero.

There is also evidence that untrained AI use is worse than no AI use. [Forrester's 2026 AIQ research](#) by J.P. Gownder found that employees with low AI competence often produce negative productivity. They use the tool, get wrong outputs, redo the work manually, and lose time overall. That finding should sit alongside any licence-cost calculation. The downside of a botched rollout is not a zero return on the investment. It is a return below zero.

The stages are load-bearing

The second thing is sequential. The stages of ADKAR cannot be reordered. You cannot drive Ability without Knowledge. You cannot drive Knowledge without Desire. You cannot drive Desire without Awareness. Every time I see a rollout failing at Ability, I trace it back and usually find Desire was never properly addressed. The stages are load-bearing.

Leadership behaviour is the signal

The third thing is organisational. The Troop layer, which is Peters' concept and which maps to Reinforcement in ADKAR, is where adoption either stabilises or reverts. In an SME, the Troop is the senior team and the direct peers around each user. If the senior team is using AI visibly, talking about their use of it, and feeding it back into decisions, the Troop reinforces the new behaviour. If senior people are letting their EAs or direct reports handle AI output, the Troop reads that as the real signal and adoption stalls at middle management. Leadership behaviour in this is not optional. It is the signal the workforce is actually reading.

Compound advantage

The fourth thing is about what happens when you get it right. Compound advantage. SMEs that get the three layers working together move faster than peers on comparable AI investment, because the adoption rate compounds instead of leaking. The firms in the [Kim, Kim and Koning study](#) that got use-case mapping right saw 1.9x the revenue impact. [Brynjolfsson, Li and Raymond](#), in their Quarterly Journal of Economics paper on AI at call centres, found gains averaged 14% overall but compressed the skill distribution sharply. Novice workers gained 34% in issue-resolution rate. Experienced workers saw minimal benefit. The compression is valuable for SMEs specifically, because an SME does not have a deep bench of senior people to carry the firm. Lifting the median up is how you grow.

This is a senior leadership decision

The final thing is about the decision you are making. Rolling out AI is a senior leadership decision, not a delegation to IT. The research on this is consistent. McKinsey's 2025 work found that organisations with active executive champions were three times more likely to succeed. WRITER Research found that middle-management champions drove 3.5x faster scaling. This is Troop work. The MD who delegates AI strategy to a technical team is signalling to the workforce that it is a technical matter, and the Chimp reads that as “this does not affect me until it is forced on me.” Desire dies at the top before the Knowledge stage even begins.

If you are an MD reading this, the question is not whether your technical team has chosen the right tools. The question is whether you have personally done the Awareness and Desire work with your own Troop, which includes you, your fellow directors, and your senior team. If you have not, the rest will not stick.

Honest posture

Some caveats matter for a paper like this.

Peters' Chimp Model is, in his own words, a model, not a hypothesis. It is a simplification of neuroscience into something operational. The mapping of AI behaviour onto the model in this paper is my application, not Peters'. I have not consulted Peters or Chimp Management about it. If the mapping is useful, the usefulness is in the application. If it is wrong in specifics, the correction improves the framework.

The Farach et al. 2026 Gap study, which gives Grennan's method its freshest empirical support, is exploratory rather than definitive. The pre-specified continuous outcome in the study was null. The 77-versus-62 percentage-point finding that features in Grennan's public communications, and which I have cited in this paper, comes from a post-hoc binary analysis. Differential attrition could nullify the effect under stricter analysis. I treat it as directional evidence consistent with a hypothesis I was already operating on, not proof of a doubling effect. The study's stronger finding is the negative one. Rigid, mandated workflow protocols reduced output quality and made teams far less likely to complete the task. That finding is the more useful rhetorical asset for anyone making the case that heavy-handed top-down AI process design is worse than no intervention.

ADKAR is an empirically grounded framework built on Prosci's research across thousands of organisations. The research is proprietary and the samples are not always disclosed in detail. It is the dominant framework in the change management field, which is both an argument in its favour and a reason to be alert to selection effects in how it gets reported.

The broader academic base, including the Dell'Acqua et al. Jagged Frontier study, the Brynjolfsson, Li and Raymond work, the Kim, Kim and Koning mapping paper, and the Lee et al. critical thinking survey, is well-sourced and peer-reviewed. Where I cite these, I have cross-checked the specific numbers and study designs.

What this framework does not cover, and what any MD reading it should be aware of, is the strategic portfolio question of where AI investment should be made in the first place. The paper assumes an AI investment has been decided on or is underway. Deciding whether AI is the right investment for a given business, relative to process redesign, hiring, or operational change, is a separate question. The framework here is about making an AI investment succeed. Not about deciding whether to make one.

The framework is also not a substitute for specific technical judgement about which tools, models, or implementations to select. Those choices depend on use case, data posture, regulatory context, and integration architecture. This paper describes the human and organisational scaffolding around AI adoption. The technical scaffolding is

separate, and often best decided after the readiness diagnostic is complete, not before.

Finally, the framework is not a closed system. It is a working synthesis. Future research, including follow-up studies on the Farach 2026 finding and continued work on the Dell'Acqua Jagged Frontier, will sharpen or challenge specific claims. Treat it as the current best working version, not the final word.

What to do next

If this paper has been useful, the practical implications are these.

If you are early in AI adoption, start with the readiness diagnostic before you commit to tools or training. SPARK is how I do that. Any equivalent diagnostic that addresses all three layers would serve the same function.

If you are in the middle of a rollout that is not landing, stop adding more training. Find out which layer has broken. The symptoms and the causes are different things.

If you are a senior leader and you have not personally done the work the framework asks of the whole workforce, you are the bottleneck. The Troop reads its senior people. Nothing below you will be better than the signal you are sending.

If you want to talk about how this applies specifically to your business, [get in touch](#).

About the author

Chris Duffy is the Founder and Chief AI Officer (CAIO) of IgniteAI Solutions Ltd, a UK fractional CAIO practice helping SMEs implement AI that works at every layer of the problem. Before founding IgniteAI, he spent 23 years in the British Army, including six years with 22 SAS as a Special Forces Communicator, and in Q1 2025 served as Technical Programme Manager on the UK Ministry of Defence's [Project ASGARD](#), the British Army's AI-enabled digital targeting web, delivered in six months against a two-year norm. He works with defence, professional services, manufacturing, e-commerce and financial services clients across the UK SME landscape to prove that his AI methodologies can be sector agnostic.

Website: igniteaisolutions.co.uk

LinkedIn: linkedin.com/in/christopher-duffy-caio

Sources

Brynjolfsson, E., Li, D. and Raymond, L. (2025). Generative AI at Work. Quarterly Journal of Economics 140(2): 889-942.

<https://academic.oup.com/qje/article/140/2/889/7990658>

Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F. and Lakhani, K. (2023). Navigating the Jagged Technological Frontier. Harvard Business School Working Paper 24-013.

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321

Farach, A., Cambon, A., Tankelevitch, L., Hsueh, C. and Janssen, R. (2026). Scaffolding Human-AI Collaboration: A Field Experiment on Behavioral Protocols and Cognitive Reframing. arXiv:2604.08678.

<https://arxiv.org/abs/2604.08678>

Gownder, J.P. (2026). Your Employees Aren't Ready For AI — And It's A Problem. Forrester Research, based on AIQ 2.0: Employees (Still) Aren't Ready To Succeed With Workforce AI.

<https://www.forrester.com/blogs/your-employees-arent-ready-for-ai-and-its-a-problem/>

Grennan, C. AI Mindset.

<https://www.ai-mindset.ai>

Hiatt, J. (2006). ADKAR: A Model for Change in Business, Government and our Community. Prosci Research.

<https://www.prosci.com/methodology/adkar>

Institute of Directors (2024). SME leadership and AI readiness research.

Kim, H., Kim, D. and Koning, R. (2026). Mapping AI into Production: A Field Experiment on Firm Performance. SSRN 6513481.

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6513481

Lee, H., Sarkar, A., Tankelevitch, L. et al. (2025). The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. CHI 2025.

<https://www.microsoft.com/en-us/research/publication/the-impact-of-generative-ai-on-critical-thinking-self-reported-reductions-in-cognitive-effort-and-confidence-effects-from-a-survey-of-knowledge-workers/>

McElheran, K., Yang, M-J., Kroff, Z. and Brynjolfsson, E. (2025). The Rise of Industrial AI in America: Microfoundations of the Productivity J-Curve(s). SSRN 5036270.

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5036270

McKinsey (2025). State of AI survey.

Mollick, E. (2024). Co-Intelligence: Living and Working with AI. Portfolio.

Peters, S. (2012). The Chimp Paradox: The Mind Management Programme for Confidence, Success and Happiness. Vermilion.

<https://chimpmanagement.com/the-chimp-paradox/>

Peters, S. (2021). A Path Through The Jungle. Mindfield Media.

<https://chimpmanagement.com/a-path-through-the-jungle/>

Prosci (2024–2025). Applying ADKAR to AI Adoption.

<https://www.prosci.com/blog/adkar-model>

This paper is an independent application of the published work of Professor Steve Peters, Conor Grennan, and Prosci. It is not endorsed by any of them. Attribution throughout is intended to be clean and to direct the reader to the primary sources.

© 2026 IgniteAI Solutions Ltd. Distribution permitted with full attribution to source.



Where to start

If you are looking at an AI investment that is not landing, the first job is to find out which of the three layers is failing. That is a diagnosis, not a sales conversation.

Send me the problem you are trying to solve and I will tell you straight whether AI is the right tool for it. If it is not, I will say so.

chris.duffy@igniteaisolutions.co.uk

www.igniteaisolutions.co.uk

United Kingdom

Ignite AI Solutions Ltd
66 Paul Street, London EC2A 4NA
Company No. 16194166 · VAT 498925711

United Arab Emirates

Ignite AI Solutions FZCO
Building A1, Dubai Digital Park, Dubai Silicon Oasis, Dubai
Company No. 80682

© 2026 IgniteAI Solutions Ltd. Distribution permitted with full attribution to source.