

The Education-Implementation Gap: Why AI Courses Do Not Translate to Operational Outcomes for Small Business Owners

Humberto Inciarte

May 2026

Abstract

The proliferation of artificial intelligence (AI) education products marketed to small and medium-sized business (SMB) owners rests on an implicit premise: that knowledge acquisition is the primary bottleneck preventing AI value capture in owner-operated firms. This paper examines that premise against the empirical literature on training transfer, online course completion, and workplace AI deployment. Three convergent bodies of evidence are reviewed: (1) the forty-year transfer-of-training literature documenting that classroom-acquired skills decay to between thirty and fifty percent application within twelve months, with weakest transfer in environments lacking post-training support (Baldwin & Ford, 1988; Blume et al., 2010; Saks & Belcourt, 2006); (2) the open online course completion record, where the most prestigious platforms — MITx and HarvardX — recorded a 3.13% completion rate across 2017–2018 with no improvement over six years (Reich & Ruipérez-Valiente, 2019); and (3) controlled workplace evidence demonstrating that AI tool *deployment*, embedded in workflow, produced productivity gains of 14% on average and 34% for novice workers without requiring formal AI training (Brynjolfsson et al., 2025). The analysis applies the Amabile–Kramer progress principle and Bjork’s fluency illusion to argue that course purchase generates a subjective signal of advancement that is dissociated from operational change. The paper proposes the Agentes Para Tu Negocio framework — a done-with-you implementation model rooted in implementation science — as a structurally different intervention category, and identifies the empirical study of such interventions in owner-operated SMBs as a research priority.

Author: Humberto Inciarte

Affiliation: Department of Research, Arquitectos de la Felicidad / NeuroFlow 30H™

ORCID: 0009-0006-0379-7328

Corresponding Author: hola@humbertoinciarte.com

Date: May 2026

1. Introduction

1.1 The Prevailing Narrative

A dominant assumption shapes the AI education market directed at small business owners: that the bottleneck preventing operational AI adoption is one of knowledge. The implicit causal chain is straightforward — owners who do not understand AI cannot deploy it; therefore, owners who learn AI will deploy it; therefore, the optimal first investment is education. This logic is reinforced by an expanding ecosystem of online courses,

prompt-engineering bootcamps, certificate programs, and self-paced tutorials, the majority of which position learning as the precondition for value capture.

The narrative has commercial coherence. It places the owner-operator at the center of the implementation responsibility. It treats AI as a generalizable skill to be acquired rather than a system to be embedded. And it frames the friction between intent and outcome as a competence gap, solvable through additional content consumption.

The Organisation for Economic Co-operation and Development (OECD, 2025), in a discussion paper for the G7 surveying more than five thousand SMEs across seven countries, found that 50% of SMEs identify skills gaps as a primary barrier to generative AI adoption. The volume of AI training products responding to this perceived gap has expanded accordingly. Yet the same OECD data reveals that fewer than 30% of SMEs already using generative AI provide formal AI training to staff — and the firms that do not train report adopting AI to *compensate* for skills gaps rather than as the result of closing them.

1.2 The Problem

The empirical record contradicts the prevailing narrative in three places. First, the transfer-of-training literature, accumulated over four decades, documents that even effective classroom training produces application rates that decay sharply within months of delivery, with the rate of decay most pronounced in environments lacking the structural support to sustain new behavior (Saks & Belcourt, 2006). Second, the open online course record, which constitutes the dominant delivery format for self-directed AI education, has stabilized at single-digit completion rates with no measurable improvement over a six-year observation window across the most prestigious platforms (Reich & Ruipérez-Valiente, 2019). Third, the most rigorous workplace AI study to date — a field deployment across 5,179 customer-support agents at a Fortune 500 firm — produced its largest gains (34% productivity for novice workers) through tool *deployment integrated into workflow*, not through prior training (Brynjolfsson et al., 2025).

These three findings, taken together, suggest that the operational gap experienced by owner-operators is not, at least primarily, a gap of knowledge. The phenomenon requiring explanation is therefore the persistence of the knowledge-first assumption in the face of contrary evidence, and the cost of that assumption to the firms acting on it.

1.3 Research Question and Thesis

This paper examines whether the assumption that AI courses produce operational outcomes for small business owner-operators is supported by available evidence, and proposes an alternative implementation-first framework — *Agentes Para Tu Negocio* — that addresses the structural causes of the gap identified in the literature.

The thesis advanced is twofold. First, that the act of course purchase and consumption activates the same psychological signal of progress (Amabile & Kramer, 2011) that operational advancement activates, while producing demonstrably less of it; the felt experience of advancement is therefore a poor instrument

for predicting operational change. Second, that the population most affected by this misalignment — owner-operators in small firms — faces precisely the post-training environment characteristics (low slack time, no peer modeling, no organizational scaffolding) that the transfer literature identifies as predictive of failed transfer (Blume et al., 2010). The remedy this paper proposes is structural rather than informational: a reduction of the friction between business judgment and operational result, achieved through done-with-you implementation rather than additional content delivery.

2. Literature Review

2.0 Methodological Note

This review synthesizes peer-reviewed empirical studies, organizational case reports, and institutional data published between 1988 and 2026, sourced from PsycINFO, Google Scholar, SSRN, the National Bureau of Economic Research (NBER), and OECD and U.S. federal databases. Inclusion criteria prioritized studies with measurable training-to-practice transfer outcomes, completion rates from large-scale online education datasets, and empirical evidence on workplace AI deployment in non-technical populations. Grey literature from established institutional sources (McKinsey, BCG, OECD, U.S. Small Business Administration) was included where peer-reviewed evidence on small and medium-sized business AI adoption was limited, and is identified as such throughout.

2.1 Transfer of Training: A Forty-Year Empirical Problem

The disjunction between training delivery and workplace application has been a documented concern in industrial-organizational psychology since at least the 1980s. Baldwin and Ford (1988), in their foundational review in *Personnel Psychology*, defined transfer of training as the generalization and maintenance of trained knowledge and skill in the work context, and proposed that transfer is a function of trainee characteristics, training design, and the work environment in which application occurs. Their review opened with the observation, sourced from Georgenson (1982), that not more than ten percent of training expenditures may produce on-the-job transfer. The figure has been widely cited in subsequent literature, though Fitzpatrick (2001) and Saks (2002) have noted that the original estimate was a practitioner observation rather than an empirical finding.

The empirical record that has accumulated since Baldwin and Ford has provided more defensible numbers. Saks and Belcourt (2006), surveying training professionals across 150 organizations, found that an estimated 62% of trained content is applied immediately following training, dropping to 44% at six months and 34% at one year. The decay pattern is consistent across studies and is driven less by trainee characteristics than by post-training environmental support. Burke and Hutchins (2007), in an integrative review of the field, identified work environment factors — supervisor support, peer modeling, transfer climate — as among the strongest predictors of sustained transfer.

The most authoritative quantitative synthesis is Blume, Ford, Baldwin, and Huang (2010), a meta-analysis

of 89 empirical studies that confirmed the dominant role of trainee motivation, cognitive ability, and supportive work environment in determining transfer outcomes. Blume and colleagues drew an important distinction between *closed skills* (procedural tasks with defined correct execution) and *open skills* (judgment-dependent tasks with multiple acceptable applications), finding that effect sizes for transfer are smaller and more environment-dependent for open skills than for closed ones. AI implementation in a small business — the integration of a general-purpose technology into specific commercial workflows — is paradigmatically open. The transfer literature predicts, with strong empirical support, that classroom or course delivery of such content into an environment lacking post-training scaffolding will produce minimal sustained application.

2.2 The Cognitive Architecture of Felt Progress

The persistence of course-buying behavior despite poor transfer outcomes invites a cognitive explanation. Two literatures are relevant.

The first is Amabile and Kramer's (2011) progress principle, derived from analysis of approximately twelve thousand diary entries collected from 238 employees across seven companies and 26 project teams. The authors found that the single most powerful driver of positive inner work life — defined as the constellation of emotions, perceptions, and motivations that workers experience during a work day — was the experience of making progress in meaningful work, even in small steps. The phenomenology of progress, in their formulation, is the proximal psychological reward that sustains engagement. The progress principle has been generative for organizational scholarship, but its application to *information consumption* warrants examination. The principle was developed in the context of bounded, observable work tasks; the question of whether the felt experience of progress can be elicited by activities that produce no operational change has not been directly tested.

The second relevant literature is the cognitive psychology of learning illusions. Bjork and Bjork (2011, 2020) have documented the phenomenon of *fluency illusion*: the subjective experience of competence produced by passive exposure to material, which is dissociable from — and a poor predictor of — durable retention or capability. Karpicke and Blunt (2011), in a study published in *Science*, found that retrieval practice produced learning approximately 1.5 standard deviations greater than elaborative study, and that students were *unable to predict* this benefit. Asked which method produced more learning, students systematically selected the less effective method. Dunlosky and colleagues (2013), in a comprehensive review for the Association for Psychological Science, classified the techniques most associated with course consumption — rereading, highlighting, summarization — as among the lowest-utility learning techniques in the cognitive science consensus.

These two literatures, taken together, suggest a mechanism. The act of consuming a course delivers the proximal psychological reward of felt progress (Amabile & Kramer, 2011) while simultaneously producing the proximal cognitive reward of felt fluency (Bjork & Bjork, 2020) — two signals the brain treats as evidence of operational advancement, neither of which is a reliable indicator of it. The signal-to-substance dissociation explains how a course buyer can sustain repeated purchase behavior across years without observable change

in operational outcome.

2.3 Online Education at Scale: The Completion Reality

The dominant delivery format for self-directed AI education today — the online course or massive open online course (MOOC) — has accumulated more than a decade of completion data. The pattern is consistent.

Reich and Ruipérez-Valiente (2019), in *Science*, analyzed all MITx and HarvardX courses delivered through the edX platform between 2012 and 2018. Across the full participant population, the completion rate stabilized at 3.13% in the 2017–2018 window, with no improvement over the six-year observation period. Among verified (paying) students, completion rates declined from 56% to 46% over the same window. More than half of all registrants — 52% — never began the course they had enrolled in.

The pattern is corroborated across the broader MOOC literature. Jordan (2015), in a meta-dataset of 221 courses, found a median completion rate of 12.6%, with a range from 0.7% to 52.1%; the lower end was concentrated in courses lacking assessment requirements or characterized by extended length. Chuang and Ho (2016), in HarvardX–MITx working paper data covering 290 courses and 4.5 million participants, reported a certification rate of approximately 5.4% across the population, rising to 36% among those who accessed at least half of the course content. Differential completion between paying (59%) and non-paying (5%) tracks confirms that financial commitment correlates with completion — but provides no evidence that completion correlates with operational outcomes downstream.

The completion data are particularly significant because they describe behavior in conditions of self-selection: individuals who enrolled in the course had already expressed intent. The transfer step from completion to operational application is unmeasured in this literature, and is presumptively further attenuated. The course-completion record therefore represents an upper bound on the proportion of course buyers who plausibly could implement what they consumed; the operational-implementation record is necessarily smaller.

2.4 Research Gap

The literature establishes three points: that classroom training transfers poorly to workplace practice in the absence of environmental scaffolding (Section 2.1); that the cognitive and emotional signals associated with course consumption are dissociated from actual capability (Section 2.2); and that the dominant delivery format for AI education completes at single-digit rates, with operational application unmeasured (Section 2.3). What the literature does not adequately address is the population most affected by the convergence of these three patterns: small business owner-operators, whose post-training environment is uniquely impoverished in transfer-supporting features.

The owner-operator profile is characterized by structural absence of the supports the transfer literature identifies as predictive: no organizational training mandate, no peer cohort modeling new behavior, no protected time blocks for application, and no supervisor or coach to provide post-training feedback. The owner-operator both delivers and applies the training, with no separation of role. The implication is that the transfer

rate predicted for an owner-operator who completes an AI course is, by structural argument, lower than the population-average rates already documented in the literature — and the population-average rates are already low.

This paper contributes by synthesizing the transfer, cognitive, and online education literatures into a single argument about the small-business AI education market, and by proposing an implementation-first alternative framework grounded in the implementation science literature.

3. Analysis and Discussion

3.1 The Implementation Premium: Direct Evidence from Workplace AI

The most rigorous empirical study of generative AI in a real workplace, to date, provides direct evidence of the implementation premium. Brynjolfsson, Li, and Raymond (2025), in a paper published in *The Quarterly Journal of Economics*, analyzed data from 5,179 customer-support agents at a large software firm following the staggered rollout of an AI-based conversational assistant. The study design — staggered access to a deployed tool, integrated into the agents' existing workflow — is methodologically the cleanest available approximation of an implementation-first intervention.

The findings are notable in two respects. Average productivity, measured by issues resolved per hour, increased by approximately 14% across the workforce. Among novice and low-skilled workers, productivity gains reached 34%; among experienced, high-skilled workers, gains were minimal. Customer sentiment improved, supervisor escalations decreased, and worker retention increased. The intervention required no formal training; the tool was embedded in the workflow and offered context-sensitive recommendations during live customer interactions.

The pattern is significant because it inverts the assumption embedded in the AI education market. The workers who gained most were not those who had acquired the most knowledge about AI; they were those who had been given a deployed system that operated within their workflow. The capability uplift was not knowledge transfer — it was friction reduction between the worker's existing judgment about customer needs and the system's contribution to acting on that judgment.

The macro evidence converges. The Boston Consulting Group (2024), in a survey of one thousand executives across fifty-nine countries, found that 74% of companies have not unlocked AI value, with implementation challenges concentrated in people and process (70% of the difficulty by BCG's classification) rather than in technology or algorithms (10%). The McKinsey Global State of AI report (2025) found that 88% of organizations now use AI in at least one function but that nearly two-thirds have not begun scaling AI across the enterprise; AI high performers were 2.8 times more likely than peers to have redesigned workflows, suggesting that the differentiator is structural integration rather than literacy.

The implication consistent across these findings is that the operational return on AI is captured in the act of deployment, not in the act of learning. The empirical curve relating training-hours-consumed to operational

outcome is, on the available evidence, flat or near-flat for the majority of populations. The empirical curve relating workflow-integrated deployment to operational outcome is steep for the populations that need the gain most.

3.2 The Owner-Operator Constraint

The analysis above describes the average enterprise context. The owner-operator context is more constrained. Three structural features of small owner-operated businesses make the education-first model especially poorly fitted.

The first is the absence of separation between learner and applier. In the transfer literature, the population most likely to apply trained content is not the most able trainee but the trainee whose post-training environment provides scaffolding — supervisor follow-up, peer cohort, dedicated application time (Blume et al., 2010). The owner-operator is, by definition, alone in the post-training environment; there is no supervisor to follow up because the owner is the supervisor, no peer cohort because the owner is the only employee in their category, and no protected time because the unprotected time is the work that generates the firm's revenue.

The second is the open-skill nature of AI implementation. Blume and colleagues (2010) document smaller transfer effects for open skills than for closed skills, and AI implementation in a specific commercial context is paradigmatically open: it requires translation from general technique to particular firm circumstance — a translation that depends on judgment about the firm rather than knowledge about the tool. The entrepreneurship education literature, the closest analogue, has consistently produced small effect sizes. Martin, McNally, and Kay (2013), in a meta-analysis of 42 samples involving 16,657 participants in *The Journal of Business Venturing*, found a correlation of 0.159 between entrepreneurship education and entrepreneurship outcomes — statistically significant but operationally small. The authors observed that academically focused entrepreneurship education produced larger effects ($r = 0.238$) than training-focused programs ($r = 0.151$), inverting the intuition that practical training would outperform theoretical content.

The third is the cost asymmetry of failed implementation. For an owner-operator, time invested in self-directed learning that does not produce a deployed system is not recoverable through other organizational mechanisms; the organization is the owner. Macnamara, Hambrick, and Oswald (2014), in a meta-analysis of deliberate-practice effects in *Psychological Science*, found that deliberate practice — the most effective form of effortful learning identified in the cognitive science literature — accounts for less than 1% of the variance in professional performance. The implication is not that practice does not matter, but that variance in professional outcomes is dominated by factors other than the volume of effortful practice: context, fit, environment, and the structure of the work itself. For an owner-operator weighing the marginal hour, the empirical case for additional content consumption is weaker than the case for a structural change in the work environment that captures value from existing judgment.

In quantitative terms, the cost asymmetry has direct implications. An owner-operator allocating ten hours per week to self-directed AI experimentation, valued at a conservative hourly opportunity cost of \$100, expends \$4,000 per month and \$16,000 over four months. The empirical literature predicts that this allocation,

under owner-operator post-training conditions, produces transfer rates at the low end of the documented range, with operational outcomes further attenuated. The cost-effectiveness of the alternative — purchasing a deployed system that operates within the firm’s existing workflow — is bounded by whether the system performs.

3.3 Proposed Framework: Done-With-You Implementation

The implementation science literature provides a structural alternative. Damschroder and colleagues (2022), in their update to the Consolidated Framework for Implementation Research (CFIR 2.0), catalog the determinants under which a given intervention transitions from knowledge to applied practice in real organizational settings. The framework identifies five domains: characteristics of the intervention, the inner setting (the organization), the outer setting (the context), the individuals involved, and the implementation process itself. CFIR has been adopted across health, education, and management literatures because it formalizes a finding implicit in the transfer literature: implementation is its own discipline, distinct from training, and requires its own design.

The *Agentes Para Tu Negocio* framework, advanced in this paper as a candidate structural intervention for SMB AI adoption, is grounded in this distinction. The framework’s central operating principle is that the small business owner does not require additional knowledge of AI; the owner requires reduced friction between their existing judgment about their business and the operational systems that translate that judgment into result. The intervention is therefore not a course but a co-built system: a single, narrowly scoped operational system, identified through structured diagnosis of the firm’s primary bottleneck, constructed with the owner’s tacit knowledge encoded into it, and delivered functioning rather than instructed.

The framework operationalizes four moves consistent with the implementation science literature. First, *bottleneck-first selection*: rather than addressing AI as a general capability, the intervention isolates the single operational constraint whose resolution produces the largest immediate change in firm-level outcome — sales conversion, response latency, qualification accuracy, fulfillment friction. The selection is judgment-driven, not technology-driven. Second, *judgment encoding*: the system is constructed with the owner’s existing operational knowledge embedded in its decision rules, rather than asking the owner to acquire the meta-knowledge of how to encode such rules themselves. Third, *workflow integration*: the system is delivered as a deployed component of the firm’s existing operating cadence, consistent with the workflow-integration finding from Brynjolfsson and colleagues (2025), rather than as a discrete tool to be invoked. Fourth, *post-deployment scaffolding*: the framework includes ongoing access to the implementer for system iteration, paralleling the supervisor-and-peer scaffolding the transfer literature identifies as critical for sustained application.

The framework is presented here as a research model rather than a tested intervention. It synthesizes principles whose individual components have empirical support — workflow integration (Brynjolfsson et al., 2025), implementation science scaffolding (Damschroder et al., 2022), the knowing-doing gap diagnosis (Pfeffer & Sutton, 1999) — into a service-delivery configuration appropriate to the owner-operator population. Whether the framework, as a configured intervention, produces measurable operational outcomes

superior to course-based alternatives in randomized comparison is an empirical question that has not yet been addressed and which the discussion below identifies as a research priority.

3.4 Practical Implications

For the owner-operator population, the analysis above suggests three practical implications.

The first concerns allocation. The marginal dollar an owner-operator can allocate to AI capability building has a measurably different expected return depending on the form of the allocation. The empirical record on self-directed online courses, applied to owner-operator post-training conditions, predicts low operational return. The empirical record on workflow-integrated deployment, in the closest available study (Brynjolfsson et al., 2025), produces large gains for the population most analogous to the SMB owner-operator profile (novice users in the worker classification). The implication is not that learning has no value, but that the order in which learning and implementation occur is consequential: implementation that produces a working system creates an environment in which subsequent learning has somewhere to land.

The second concerns the assessment of progress. Pfeffer and Sutton (1999), in *California Management Review*, documented what they termed the *smart talk trap*: the tendency in organizations to substitute the discussion, planning, and articulation of action for the action itself. The trap is particularly acute when the substitute behaviors produce the same psychological signals as the original behavior. For an owner-operator, the discipline of distinguishing between operational change and the felt experience of operational change is structurally absent unless deliberately constructed. The literature reviewed in Section 2.2 implies that this discipline cannot be constructed through introspection alone; the cognitive signals are misleading. It can be constructed through external accountability to operational metrics — the firm’s actual conversion, response, or fulfillment rates — measured before and after intervention.

The third concerns the structural diagnosis the implementation produces. A system designed to address the owner-operator’s primary bottleneck reveals, through its design and deployment, the firm’s dependency structure: which decisions still flow only through the owner, which information still lives only in the owner’s head, which approvals still require the owner’s presence. This diagnostic byproduct is, in some respects, more durable than the system itself. It identifies the structural conditions under which the firm depends on the owner — the conditions that determine whether the owner is operating a business or operating *as* a business.

4. Conclusions

4.1 Summary of Findings

The available empirical literature does not support the assumption that AI courses produce operational outcomes for small business owner-operators. The transfer-of-training literature, accumulated across forty years, documents that classroom-acquired skills decay to between thirty and fifty percent application within twelve months, with the most pronounced decay in environments lacking post-training support — precisely the en-

vironment of the owner-operator. The online course completion literature documents single-digit completion rates across the most prestigious platforms, with no measurable improvement over six years of observation. The cognitive science of learning illusions documents that the felt experience of learning is dissociated from actual capability. The most rigorous workplace AI study to date documents that the productivity gains from AI in real workplaces are produced through tool deployment integrated into workflow, not through prior training, with the largest gains accruing to the population most analogous to the SMB owner-operator profile.

These findings suggest a reformulation of the problem. The structural feature of course consumption that explains its commercial persistence — its capacity to deliver felt progress and felt fluency without operational change — is the same feature that makes it an unreliable intervention for owner-operator firms whose objective is operational change. The course is not failing because it is poorly designed; it is succeeding at what it does, which is producing a psychological signal of advancement. The signal is, on the evidence, dissociable from the substance.

The *Agentes Para Tu Negocio* framework, advanced in this paper, treats the operational gap as a friction problem rather than a knowledge problem. The owner does not require additional knowledge to act on their existing business judgment; the owner requires a reduction in the friction between that judgment and the system that executes it. The framework is presented as a research model whose configuration is consistent with established implementation science principles and which awaits direct empirical testing.

The synthesis offered here can be expressed compactly: course purchase functions as a substitute progress signal — the felt experience of advancement without the substance of operational change. A functioning system, by contrast, is operational advancement itself. The distinction is not rhetorical; it is structural, and the literature reviewed above is, on balance, consistent with treating it as such.

4.2 Limitations

This review is subject to the limitations inherent in narrative literature synthesis. Selection of the literature reflects the author's editorial judgment about the relevance of the included studies, and alternative selections could plausibly produce different emphases. The most rigorous workplace AI study cited — Brynjolfsson, Li, and Raymond (2025) — is set in a Fortune 500 customer-support context and is not directly transferable to owner-operated SMB conditions; its inclusion is justified by its methodological rigor and by the population analogue argument advanced in Section 3.1, but the analogue is imperfect. The proposed *Agentes Para Tu Negocio* framework has not been validated in controlled comparison with course-based alternatives; the empirical case for it rests on the principles it embodies rather than on outcomes it has demonstrated in randomized comparison. The transfer-of-training estimates discussed in Section 2.1 derive from organizational populations that are not exclusively SMB and do not exclusively address AI implementation; their application to the owner-operator population is structural rather than directly empirical. Latin American owner-operator data, of particular relevance to the population the framework addresses, is sparse in the international literature and is identified below as a research priority.

4.3 Future Research Directions

Three lines of research are warranted by the gaps identified above.

The first is a direct comparison of course-based and implementation-based interventions in matched samples of owner-operated SMBs, measured against operational outcomes (revenue, conversion, response latency, fulfillment time) over twelve-month observation windows. The transfer literature predicts that the implementation arm will outperform the course arm; the magnitude of the difference, and the conditions under which it holds, is empirically open.

The second is the systematic measurement of completion-to-operational-outcome attrition in self-directed AI courses targeting SMB owners. The MOOC literature reports completion rates but does not report what completers subsequently implement. Closing this gap would clarify the operational ceiling of the course-first model and inform allocation decisions for owner-operators considering the format.

The third is the empirical study of done-with-you implementation models in Latin American SMB contexts, where the population the framework addresses is concentrated and where peer-reviewed evidence is currently sparse. The structural arguments advanced in this paper transfer across geographies in principle, but the empirical literature should not be assumed to do so without verification.

The convergence between cognitive psychology, transfer-of-training research, and workplace AI evidence identified here suggests an applied research agenda for the SMB AI implementation field. The discipline of distinguishing felt from operational progress, and the design of interventions that produce the latter rather than the former, are problems amenable to empirical investigation. They are also, at present, underrepresented in the literature relative to their commercial and economic significance.

Resumen en Español

La proliferación de productos educativos de inteligencia artificial (IA) dirigidos a dueños de pequeñas y medianas empresas (PYMES) descansa sobre una premisa implícita: que la adquisición de conocimiento es el cuello de botella principal que impide la captura de valor de IA en empresas operadas por su dueño. Este artículo examina dicha premisa contra la literatura empírica sobre transferencia de capacitación, completación de cursos en línea, y despliegue de IA en entornos laborales. Se revisan tres cuerpos convergentes de evidencia: (1) la literatura de cuarenta años sobre transferencia de capacitación, que documenta que las habilidades adquiridas en aula decaen a entre treinta y cincuenta por ciento de aplicación en doce meses, con la transferencia más débil en entornos sin soporte post-capacitación (Baldwin & Ford, 1988; Blume et al., 2010); (2) el registro de completación de cursos abiertos en línea, donde las plataformas más prestigiosas — MITx y HarvardX — registraron una tasa de completación de 3.13% durante 2017–2018 sin mejora a lo largo de seis años (Reich & Ruipérez-Valiente, 2019); y (3) evidencia controlada en el lugar de trabajo que demuestra que el *despliegue* de IA, integrado al flujo de trabajo, produjo aumentos de productividad de 14%

en promedio y 34% para trabajadores novatos sin requerir capacitación formal (Brynjolfsson et al., 2025). El análisis aplica el principio del progreso de Amabile y Kramer y la ilusión de fluidez de Bjork para argumentar que la compra de un curso genera una señal subjetiva de avance disociada del cambio operacional. El artículo propone el marco Agentes Para Tu Negocio — un modelo de implementación done-with-you fundamentado en ciencia de implementación — como una categoría estructuralmente distinta de intervención.

Palabras clave: transferencia de capacitación, completación de cursos en línea, brecha educación-implementación, aprendizaje en PYMES, ROI de educación en IA, resultados de desarrollo profesional, brecha conocimiento-acción, práctica deliberada, Agentes Para Tu Negocio

References

- Amabile, T. M., & Kramer, S. J. (2011). The power of small wins. *Harvard Business Review*, 89(5), 70–80.
- Amabile, T. M., & Kramer, S. J. (2011). *The progress principle: Using small wins to ignite joy, engagement, and creativity at work*. Harvard Business Review Press.
- Baldwin, T. T., & Ford, J. K. (1988). Transfer of training: A review and directions for future research. *Personnel Psychology*, 41(1), 63–105. <https://doi.org/10.1111/j.1744-6570.1988.tb00632.x>
- Bjork, R. A., & Bjork, E. L. (2020). Desirable difficulties in theory and practice. *Journal of Applied Research in Memory and Cognition*, 9(4), 475–479. <https://doi.org/10.1016/j.jarmac.2020.09.003>
- Blume, B. D., Ford, J. K., Baldwin, T. T., & Huang, J. L. (2010). Transfer of training: A meta-analytic review. *Journal of Management*, 36(4), 1065–1105. <https://doi.org/10.1177/0149206309352880>
- Boston Consulting Group. (2024). *Where's the value in AI?* BCG Research Report.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.3386/w31161>
- Burke, L. A., & Hutchins, H. M. (2007). Training transfer: An integrative literature review. *Human Resource Development Review*, 6(3), 263–296. <https://doi.org/10.1177/1534484307303035>
- Chuang, I., & Ho, A. D. (2016). HarvardX and MITx: Four years of open online courses — Fall 2012–Summer 2016. *SSRN Working Paper*. <https://doi.org/10.2139/ssrn.2889436>
- Damschroder, L. J., Reardon, C. M., Widerquist, M. A. O., & Lowery, J. (2022). The updated Consolidated Framework for Implementation Research based on user feedback. *Implementation Science*, 17, 75. <https://doi.org/10.1186/s13012-022-01245-0>
- Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improv-

- ing students' learning with effective learning techniques: Promising directions from cognitive and educational psychology. *Psychological Science in the Public Interest*, 14(1), 4–58. <https://doi.org/10.1177/1529100612453266>
- Fitzpatrick, R. (2001). The strange case of the transfer of training estimate. *The Industrial-Organizational Psychologist*, 39(2), 18–19.
- Georgenson, D. L. (1982). The problem of transfer calls for partnership. *Training and Development Journal*, 36(10), 75–78.
- Jordan, K. (2015). Massive open online course completion rates revisited: Assessment, length and attrition. *International Review of Research in Open and Distributed Learning*, 16(3), 341–358. <https://doi.org/10.19173/irrodl.v16i3.2112>
- Karpicke, J. D., & Blunt, J. R. (2011). Retrieval practice produces more learning than elaborative studying with concept mapping. *Science*, 331(6018), 772–775. <https://doi.org/10.1126/science.1199327>
- Macnamara, B. N., Hambrick, D. Z., & Oswald, F. L. (2014). Deliberate practice and performance in music, games, sports, education, and professions: A meta-analysis. *Psychological Science*, 25(8), 1608–1618. <https://doi.org/10.1177/0956797614535810>
- Martin, B. C., McNally, J. J., & Kay, M. J. (2013). Examining the formation of human capital in entrepreneurship: A meta-analysis of entrepreneurship education outcomes. *Journal of Business Venturing*, 28(2), 211–224. <https://doi.org/10.1016/j.jbusvent.2012.03.002>
- McKinsey & Company. (2025). *The state of AI: Global survey 2025*. McKinsey QuantumBlack.
- Organisation for Economic Co-operation and Development. (2025). *AI adoption by small and medium-sized enterprises: OECD discussion paper for the G7*. OECD Publishing.
- Pfeffer, J., & Sutton, R. I. (1999). Knowing “what” to do is not enough: Turning knowledge into action. *California Management Review*, 42(1), 83–108. <https://doi.org/10.1177/000812569904200101>
- Reich, J., & Ruipérez-Valiente, J. A. (2019). The MOOC pivot. *Science*, 363(6423), 130–131. <https://doi.org/10.1126/science.aav7958>
- Saks, A. M. (2002). So what is a good transfer of training estimate? A reply to Fitzpatrick. *The Industrial-Organizational Psychologist*, 39(2), 29–30.
- Saks, A. M., & Belcourt, M. (2006). An investigation of training activities and transfer of training in organizations. *Human Resource Management*, 45(4), 629–648. <https://doi.org/10.1002/hrm.20135>