NextGen Scientific Publication
American International Journal of Computer Science and Technology

Volume 8 Issue 5, Pg. No. 15-29, AIJCST-V8I5P102, 2026
https://doi.org/10.63282/3117-5481/AIJCST-V8I5P102

Original Article

Designing AI-Agent-Enabled Entrepreneurship Training Programs: An Experiential Learning Approach to Business Automation, Opportunity Recognition, and Entrepreneurial Decision-Making

***Daniela Hayoshi**
Founder, Mado & Co., Brazil.

Abstract:

Entrepreneurship education increasingly employs generative artificial intelligence, yet proficiency in prompting a chatbot does not necessarily prepare learners to design accountable AI-agent workflows or make defensible business decisions. This article develops an experiential training framework that integrates opportunity recognition, business process automation, and entrepreneurial decision-making within a human-governed AI-agent environment. Using a theory-informed conceptual design approach, the framework synthesizes experiential learning, entrepreneurship-as-practice, effectuation, entrepreneurial competence, AI literacy, and responsible AI governance. It specifies six learner competencies, an eight-week studio curriculum, three interconnected practice strands, and a repeatable cycle of experience, reflection, conceptualization, and bounded experimentation. Learners are asked to interrogate market claims, compare alternative venture actions, prototype low-risk automations, and document why they accept or reject agent recommendations. An agent architecture separates information retrieval and proposed actions from human authorization and outcome verification. The article also proposes observable performance measures, assessment rubrics, and a future comparative evaluation design for estimating competency gains without conflating task completion with entrepreneurial learning. The principal contribution is a practical and testable instructional design that connects technological capability with entrepreneurial judgment. Its effectiveness remains a proposition to be examined through implementation and empirical assessment rather than an established result. The framework is intended for higher education, incubators, and professional entrepreneurship programmes with varied levels of technical preparation.

Keywords:

AI Agents, Entrepreneurship Education, Experiential Learning, Business Automation, Opportunity Recognition, Entrepreneurial Decision-Making, Human-AI Collaboration.

Article History:

Received: 22.07.2026

Revised: 21.08.2026

Accepted: 29.08.2026

Published: 05.09.2026

Copyright by the Author(s). This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-NC-ND 4.0) License (https://creativecommons.org/licenses/by-sa/4.0/)

1. Introduction

Entrepreneurship education is expected to prepare people to notice unmet needs, mobilize limited resources, test uncertain assumptions and act responsibly when evidence is incomplete. Conventional coursework often covers planning, marketing and finance, but entrepreneurs increasingly encounter artificial intelligence (AI) systems that can search across information sources, prepare material, use approved software tools and recommend sequences of action. A learner may now ask an AI assistant to classify customer feedback, prepare alternative pricing arguments or suggest an outreach workflow in minutes. The educational challenge is therefore not simply to teach faster content production. It is to teach when the apparent speed and coherence of an AI output should be trusted, challenged, tested or declined.

Entrepreneurship has long been understood as the study of opportunities and their discovery, evaluation and exploitation (Shane & Venkataraman, 2000). Identifying opportunities requires attention to changing conditions and the capacity to connect dispersed signals into meaningful patterns (Baron, 2006). Digital tools alter the amount of information available and the practical cost of exploring an idea, but they do not eliminate uncertainty about customer behavior, willingness to pay, operational feasibility or entrepreneurial ethics. Digital entrepreneurship also changes where venture agency resides and how entrepreneurs navigate uncertainty, making technology part of the entrepreneurial process rather than merely an efficiency aid (Nambisan, 2017). A polished model-generated market report is not proof that a viable market exists.

Recent work makes AI's relevance concrete. Renz et al. (2025), in an experiment involving 63 participants, reported improved opportunity-recognition performance among those using large language models (LLMs), with stronger gains among those who had lower initial capability. Triantoro et al. (2025) examined experiential AI instruction for entrepreneurs and small-business owners and reported improvements in AI competency. Ahmad et al. (2026), however, found that a bounded, human-led AI arrangement could preserve active discussion and judgment without establishing improvements in judgment or creativity. These studies support pedagogical exploration but do not warrant the assumption that all AI-enabled training produces better entrepreneurial decisions.

An additional distinction matters. A chatbot ordinarily responds to successive user requests; an AI agent may plan intermediate steps, access selected tools, consult stored information and propose or execute actions within a defined environment (Wang et al., 2024). Agentic systems are attractive for business automation, but they introduce risks involving erroneous tool use, unauthorized communications, confidential information and over-delegation. They also generate a new learning problem: successful automation can conceal whether the student understands the commercial assumptions behind a workflow.

This paper addresses three connected gaps. First, many AI entrepreneurship activities emphasize idea generation rather than the disciplined comparison of hypotheses and evidence. Second, agent-based automation skills are rarely integrated into entrepreneurial learning objectives concerning opportunity quality and managerial responsibility. Third, evaluations often privilege confidence, intention or completion instead of observable decision quality, transfer and the safe conduct of business experiments (Nabi et al., 2017). To address these gaps, the article proposes an integrated, practice-oriented curriculum and a structured evaluation pathway.

The guiding research question is: How can AI-agent-enabled entrepreneurship training be designed to develop opportunity recognition, accountable business automation and entrepreneurial decision-making through experiential learning? Three subsidiary questions concern the competencies to teach, the structure of supervised practice, and the evidence required to assess whether learning transfers beyond a classroom demonstration. The contribution is a design and its evaluable propositions, not a claim of completed field validation.

2. Theoretical Foundations and Related Literature

2.1. Experiential Entrepreneurship Education

Kolb (1984) described learning as a process in which concrete experience, reflective observation, abstract conceptualization and active experimentation inform one another. The sequence is not a rigid checklist; meaningful learning requires movement between action and interpretation. A. Y. Kolb and D. A. Kolb (2005) further emphasized learning environments that support different forms of engagement. Applied to entrepreneurship, this implies that a student must do more than submit a finished business plan. Cope (2005) likewise conceptualized entrepreneurial learning as a dynamic process shaped by experience and critical reflection. Learners should

encounter a customer or operational problem, examine the reasoning behind possible responses, formulate working principles and return to practice with new questions.

Neck and Greene (2011) argued for teaching entrepreneurship as a method of thinking and acting rather than a collection of facts. Practice-based teaching includes simulations, design activities, experiments and reflection, which are particularly appropriate when decisions must be made before reliable predictions are possible. Neck et al. (2014) translated this approach into a portfolio of teaching practices, including experimentation, empathy and reflection. Pittaway and Cope (2007) further emphasized the experiential and collaborative features of entrepreneurial learning. Nabi et al. (2017) found, in a systematic review of 159 entrepreneurship education articles, that studies often used short-term subjective outcome measures and insufficiently described pedagogy. Their findings strengthen the case for clearly specified learning tasks, observable artifacts and delayed assessment. They do not imply that experiential methods automatically outperform all alternatives.

Generative AI can intensify the experiential cycle by providing rapid alternative explanations, mock customer objections and draft analytical artifacts. Problem-based learning scholarship stresses that effective learning arises from facilitated inquiry, self-directed investigation and reflection rather than the possession of an answer alone (Hmelo-Silver, 2004). A rapid review of ChatGPT in education similarly highlighted both opportunities for support and concerns about accuracy and educational misuse (Lo, 2023). Darnell and Gopalkrishnan (2024) describe practical ways to integrate generative AI into entrepreneurship exercises while calling attention to the importance of educational purpose. Yet a digitally mediated experience becomes educational only when the learner scrutinizes what has been generated. The curriculum proposed here therefore grades the quality of questions, verification and reflection, not merely the attractiveness of a generated output.

2.2. Opportunity Recognition and Entrepreneurial Action

Opportunity recognition is not synonymous with idea quantity. Baron (2006) conceptualized recognition as the ability to detect meaningful patterns connecting technological, demographic and market developments. Entrepreneurial learners must distinguish a genuine unmet need from an interesting technology, assess whose problem is being solved, and identify evidence that would change their judgment. DeTienne and Chandler (2004) showed that opportunity-identification processes can be taught and assessed in the entrepreneurship classroom. Davidsson (2015) also cautioned against treating an opportunity as a single unproblematic object, distinguishing external conditions, venture ideas and confidence about them. AI can help structure search and compare patterns across cases, but it may repeat widely circulating assumptions and produce confident, unsupported claims.

Effectuation provides a complementary perspective. Sarasvathy (2001) distinguished predictive, goal-driven causation from a more control-oriented logic that begins with available means, affordable loss and stakeholder commitments. Both approaches can inform training. The distinction has been operationalized in validated measures of entrepreneurial causation and effectuation, including experimentation, affordable loss and flexibility (Chandler et al., 2011). Entrepreneurial action also requires a person's willingness to act under uncertainty, not simply awareness of a possible opportunity (McMullen & Shepherd, 2006). An agent may forecast demand or compare scenarios, but a learner should also ask what can be tested inexpensively with existing contacts and resources. The proposed program deliberately requires participants to convert AI suggestions into bounded experiments rather than treating predicted returns as facts. EntreComp likewise frames entrepreneurship through interconnected competence areas of ideas and opportunities, resources and action (Bacigalupo et al., 2016).

Studies of AI use in entrepreneurship are promising but qualified. Duong and Nguyen (2024) linked ChatGPT adoption with opportunity recognition, entrepreneurial knowledge and downstream attitudes or behavior in their research model. Renz et al. (2025) provide more direct experimental evidence for support in opportunity recognition. Neither result establishes that AI will reliably judge customers' authentic preferences, or that an algorithmically attractive venture is commercially or ethically sound. For this reason, the curriculum uses competing hypotheses, primary customer evidence and explicit disconfirmation criteria.

2.3. AI Agents, Automation and Human Decision Authority

LLM-enabled agents combine language generation with planning, working memory and tool interaction; implementations differ in how much freedom an agent receives (Wang et al., 2024). Park et al. (2023) demonstrated how generative agents can maintain contextual memory, reflection and plans in simulated environments. Yao et al. (2023) developed a reasoning-and-acting approach in which language models interact with external sources and environments, while Shinn et al. (2023) used verbal feedback and memory

to support iterative agent behavior. These are technical demonstrations, not evidence that autonomous business decisions are inherently trustworthy. These advances offer teaching opportunities for workflow design and decision rehearsal. They also increase the importance of separating recommendations from authorization. A proposed email sequence, for example, is a different action from sending customer communications or changing records in a live customer relationship management system.

Human-AI performance does not follow a simple 'two heads are better than one' rule. Vaccaro et al. (2024) synthesized experimental evidence and found that human-AI combinations, on average, did not outperform the stronger of the two acting separately; results varied across decision and content-creation tasks. The implication for training is to teach calibrated reliance, not habitual approval of machine advice. Participants need to explain what evidence supports a recommendation, identify plausible error mechanisms and recognize circumstances in which human expertise or consultation should dominate.

AI literacy offers further foundations. Ng et al. (2021) organized AI literacy around understanding, application, evaluation and ethics. Long and Magerko (2020) identified AI-literacy competencies that include critical evaluation and understanding of what AI systems can and cannot do. Ouyang and Jiao (2021) contrasted AI-directed, AI-supported and AI-empowered educational paradigms, highlighting the importance of learner agency. UNESCO's AI competency framework emphasizes human agency, ethical awareness, technical application and system design (Miao et al., 2024). The NIST generative AI risk profile specifies a risk-management orientation relevant to source quality, privacy and oversight (Autio et al., 2024). These resources inform the boundaries of agent use in the proposed program without implying that a training exercise itself certifies legal or regulatory compliance.

2.4. Evidence-To-Design Implications

The literature points to a specific design problem: AI helps learners produce more possibilities, but entrepreneurship education must cultivate the ability to choose among them, validate assumptions and retain responsibility for consequential actions. The most defensible synthesis combines practice cycles, decision documentation, tool limitations and assessment of observable performance. Table 1 traces the proposed design choices to those theoretical commitments. It distinguishes cited evidence from recommendations developed in this article; the mapping is an analytical synthesis, not a new empirical finding.

Table 1. Evidence-To-Design Synthesis

Literature Perspective	Relevant Evidence or Concept	Curricular Design Implication
Experiential learning	Learning through experience, reflection, concepts and testing (Kolb, 1984).	Require recurrent action-reflection cycles and revised experiments.
Entrepreneurship-as-practice	Action, simulation and reflection as methods (Neck & Greene, 2011).	Assess applied tasks and decision records rather than plans alone.
Opportunity recognition	Pattern recognition and source-supported evaluation (Baron, 2006; Renz et al., 2025).	Separate plausible ideas from validated customer problems.
Human-AI complementarity	Human-AI teams are not universally superior (Vaccaro et al., 2024).	Teach calibrated reliance, error detection and justified rejection.
Responsible AI literacy	Human agency, ethics and risk governance (Miao et al., 2024; Autio et al., 2024).	Limit permissions and document human approval of consequential actions.

3. Research Approach and Framework Development

This article uses a conceptual, theory-informed design approach. Its product is an educational artifact comprising a competency model, training sequence, human-governed agent workflow and evaluation protocol. Design-science reasoning is appropriate because the research question asks how an intervention could be constructed and subsequently evaluated, rather than whether a completed intervention has already produced a measurable effect (Hevner et al., 2004). No participant recruitment, intervention delivery, survey administration or statistical testing was undertaken for this article.

The design process involved four analytic steps. First, foundational perspectives were selected to represent the core entrepreneurial learning problem: experience and reflection, opportunity evaluation, resource-constrained action and entrepreneurial competence. Second, recent peer-reviewed research and institutional frameworks were consulted to identify the capabilities and

limitations of generative AI, AI agents and human-AI cooperation. Third, these concepts were translated into teachable competencies and supervised tasks, with evidence requirements specified for each task. Fourth, an evaluation protocol was constructed to make the framework open to correction through future pilot implementation.

This is a purposive, narrative synthesis, not a systematic review. Sources were chosen for their theoretical relevance, direct applicability to entrepreneurship or educational practice, credible publication venue, or relevance to responsible AI use. The manuscript does not claim exhaustive database coverage, a predefined screening strategy or pooled quantitative effects. That limitation matters because a framework built from selective literature may overlook competing theories or emerging evidence. Its propositions should therefore be considered provisional.

The design assumes participants have basic digital competence but not software engineering experience; access to an institutionally approved LLM or restricted agent tool; and the opportunity to work with public, synthetic or permissioned business data. The proposed 8-week schedule is a design parameter rather than an empirically optimized duration. The model is intended to be portable across universities, incubators and professional training settings, but implementations should adapt examples, data safeguards and support levels to their contexts.

A framework is useful only if another educator can reproduce its essential decisions. Accordingly, the article identifies competencies, instructor inputs, student outputs, human approval rules, feedback points and assessment criteria. It also separates what the framework specifies from what later research must determine: effectiveness, accessibility, cost, transfer and unintended consequences.

4. Proposed AI-Agent-Enabled Experiential Framework

4.1. Learning Architecture and Competencies

The framework, designated E-ACT (Experiential AI-Agent Competence Training), joins three practice strands: recognizing opportunities, automating bounded business processes and making evidence-based entrepreneurial decisions. Across these strands, learners develop six competencies: framing a viable problem; evaluating opportunities; specifying workflows; supervising agent actions; making traceable decisions; and reflecting on outcomes. Each competency requires both a technical action and an entrepreneurial justification. A successful workflow demo without a credible business purpose remains incomplete; an appealing opportunity statement without validation and risk analysis is similarly insufficient.

Figure 1: Describes the recurrent instructional cycle. A business challenge initiates experience; learners then reflect on outputs and assumptions, develop transferable decision criteria and test a revised approach. The loop creates an explicit teaching place for surprise. Failed retrieval, contradictory customer comments or an inefficient automation becomes material for analysis rather than an embarrassment to hide. In this sense, failure is evidence about the fit between a proposed solution and its environment, provided the activity is low risk and appropriately supervised.

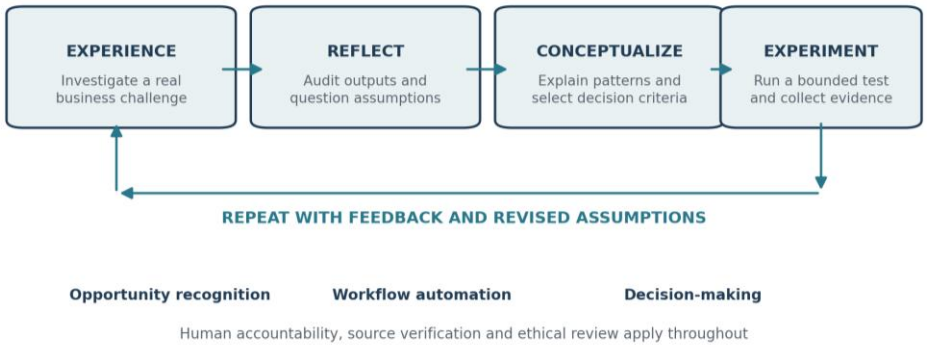


Figure 1. E-ACT Learning Cycle Connecting Experience, Reflection, Conceptualization and Bounded Experimentation

4.2. Opportunity-Recognition Strand

The opportunity strand begins with a narrow problem space, such as missed appointment bookings among local services or unpredictable replenishment among small retailers. Learners construct an evidence inventory listing the stakeholder, observed inconvenience, current workaround, possible causes and information still missing. An AI agent may search an approved collection of public materials, cluster comments or propose interview questions. Students must classify each output as observed, source-supported, inferred or speculative. This distinction operationalizes Davidsson's (2015) separation of external enabling conditions, entrepreneurial ideas and confidence in an opportunity. Only source-supported claims can enter the evidence table without a qualification.

Learners subsequently create two competing problem explanations, not just variations of the same product idea. They develop a minimal customer discovery plan, formulate potential disconfirming evidence and record whether an opportunity should be pursued, reformulated or abandoned. This protects against a common failure in AI-supported ideation: accepting fluent text as customer validation. Outputs include an opportunity hypothesis, source audit, stakeholder map and validation plan. A useful instructor intervention is to ask what evidence would make the team change its mind.

4.3. Business Automation Strand

The automation strand applies agent tools to a repeatable, low-risk process, such as categorizing synthetic inbound queries, preparing draft responses for approval or converting a spreadsheet into a weekly inventory exception report. Teams identify triggers, inputs, permissions, decision rules, expected outputs and error conditions before configuring any automation. The design uses agent reasoning and tool interaction as capabilities to be tested, rather than trusting their outputs as inherently correct (Shinn et al., 2023; Yao et al., 2023). They compare the proposed agent workflow against a simple manual or rule-based process. Introducing an agent is not an achievement when a deterministic form, spreadsheet rule or conventional script would be cheaper, safer and sufficiently reliable.

Every prototype must include a test dataset, at least three realistic failure cases, a clear rollback or stop procedure, and an audit trail. Examples of failure cases include duplicate records, missing customer context, contradictory source documents, ambiguous instructions and sensitive data accidentally embedded in a prompt. The student submits a workflow specification showing who may view information, which actions are read-only, and which require human consent. Execution outside the sandbox is prohibited unless an institution has separately approved a controlled deployment.

4.4. Decision-Making and Accountability Strand

The decision-making strand connects model outputs to a business choice with uncertain consequences. Learners compare alternatives, such as whether to invest in an appointment reminder workflow, improve the booking interface or allocate resources to customer research first. They must define evaluation criteria before viewing an agent's ranking. Criteria can include customer value, feasibility, cash requirements, data exposure, reversibility, ethical consequences and uncertainty. This ordering reduces the temptation to retrofit criteria to justify a persuasive recommendation.

A decision log records alternatives, evidence, assumptions, dissenting interpretations, the recommended action, the authorized decision-maker and the review date. The log also distinguishes predictions from commitments. The instructor can introduce a deliberately flawed agent suggestion containing a fabricated competitor statistic or an unrealistic market assumption, then evaluate whether the student identifies the defect. The educational target is the quality of critical judgment, not agreement with the machine or with the instructor.

4.5. Human-Governed Agent Configuration

Figure 2 locates the agent inside a controlled sequence: the entrepreneur defines the brief; the system proposes steps using permitted information and tools; a person verifies assumptions and authorizes any consequential action; and the team checks outcomes. The approval gate is a design requirement, not a guarantee of technical compliance. Instructors should preconfigure permissions, disable external communications and spending, use test credentials and monitor logs. Higher-risk domains such as financial transfers, legal commitments, hiring, credit decisions or medical advice are unsuitable for autonomous student execution.

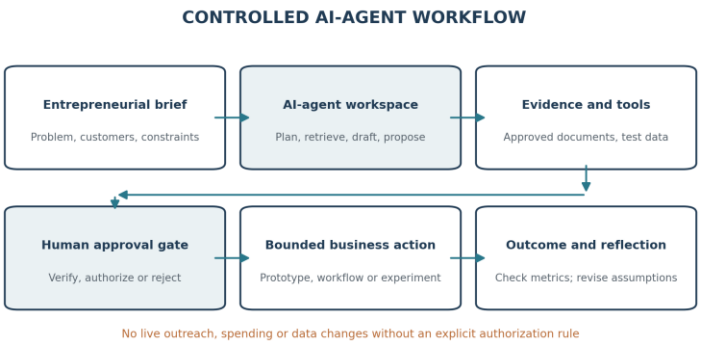


Figure 2. Proposed Human-Governed Architecture for Bounded Entrepreneurship Agent Exercises

Human oversight is substantive only when the person can understand the proposed action and has adequate time and information to refuse it. A nominal approval button with opaque evidence would not satisfy the framework's accountability principle. For this reason, every agent proposal must show its sources where applicable, intended recipient or data target, likely failure mode and proposed validation step. Table 2 maps the six competencies to artifacts and assessment signals.

Table 2. Competencies, Authentic Tasks and Assessment Artifacts

Competency	Applied Task	Evidence Assessed
Problem framing	Define stakeholder pain and competing causes.	Problem brief; assumption inventory.
Opportunity evaluation	Triangulate public evidence and discovery questions.	Source audit; disconfirmation plan.
Workflow specification	Map triggers, data, rules, exceptions and approvals.	Process map; agent specification.
Agent supervision	Challenge incorrect outputs in a sandbox.	Error log; permissions test.
Entrepreneurial decision	Rank feasible actions against predeclared criteria.	Decision log; affordable-loss rationale.
Reflection and transfer	Adapt approach to an unfamiliar business case.	Reflective memo; revised experiment.

5. Eight-Week Entrepreneurship Training Programme

5.1. Programme Design and Delivery

The proposed intervention consists of eight weekly four-hour studio meetings and 16 additional hours of independent project work, producing a designed total of 48 hours. Within each studio, approximately one hour addresses essential concepts, two hours support coached practice and one hour is reserved for peer critique and reflective synthesis. These are planned allocations, not attendance or performance results. A cohort may work individually or in teams of two to four, but every learner must submit an individual decision log so that collective outputs do not obscure personal judgment.

Participants begin with a familiar venture problem selected from a vetted case bank, a partner organization's nonconfidential challenge, or their own business setting where permission has been obtained. Instructors provide a data-use statement, assessment rubric, technical access guide and baseline exercise before the first agent task. Beginners receive templates and no-code tools; more advanced learners may design tool connectors in a sandbox. Both groups are assessed against the same entrepreneurial criteria, while technical complexity is evaluated relative to clearly stated learning objectives.

Table 3 presents the teaching sequence. Each week ends with a deliverable that becomes input to the next studio. The sequence intentionally moves from evidence collection toward controlled action, rather than opening with unrestricted agent construction. Figure 3 displays the designed distribution of learning hours. The chart represents curriculum specification only, so it should not be interpreted as observed time use or an effectiveness result.

Table 3. Proposed Eight-Week Programme and Outputs

Week	Studio Focus	Required Output
1	Business problem framing; AI literacy; baseline task	Problem brief and baseline response

2	Opportunity discovery and source evaluation	Evidence register and competing hypotheses
3	Customer discovery and opportunity testing	Interview guide; validation plan
4	Process mapping and automation selection	As-is/to-be workflow; agent-or-rule decision
5	Bounded agent prototype and failure tests	Demo/storyboard; failure and approval log
6	Entrepreneurial decision scenarios	Weighted qualitative decision log
7	Peer audit and redesigned experiment	Risk critique; revised prototype
8	Capstone defense and transfer reflection	Action brief; verification and reflection

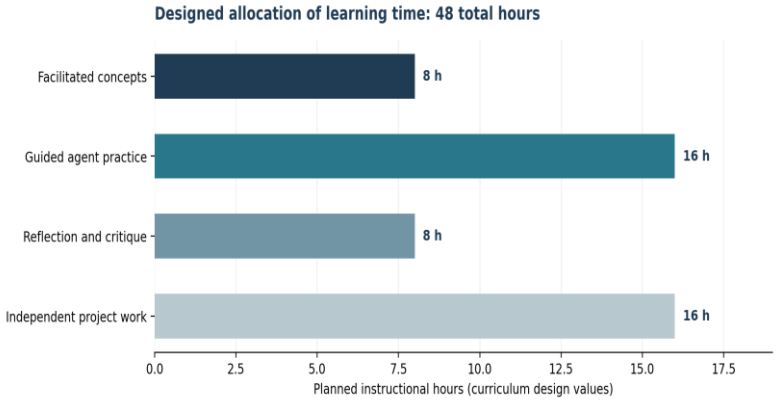


Figure 3. Planned Hours for the Proposed 8-Week Course; Design Values, Not Observed Data

5.2. Learning Activities and Formative Feedback

Week 1 establishes entrepreneurial problem framing and AI literacy. Students map where an agent might help and where a human must remain responsible. Initial exercises draw on AI-literacy competencies and problem-based facilitation to require explanation, verification and self-directed inquiry (Hmelo-Silver, 2004; Long & Magerko, 2020). They conduct a baseline opportunity-analysis task without AI, then discuss differences between automation, assistance and decision authority. In Week 2 they perform structured opportunity search with sources and create an uncertainty register. In Week 3 they conduct or rehearse customer discovery, using agent-generated questions only as provisional materials. An instructor reviews interview consent, source attribution and the difference between authentic feedback and simulated personas.

Weeks 4 and 5 shift to process mapping and the safe implementation of an agent prototype. Teams calculate a basic workload estimate using their own stated task assumptions, define acceptance criteria and test a bounded workflow. The comparison of predictable workflow rules and experimental adjustments parallels the distinction between causation and effectuation examined by Chandler et al. (2011). Feedback centers on the quality of the specification and error analysis, not the number of software integrations achieved. Week 6 introduces scenario-based entrepreneurial decisions under resource constraints. Each team receives conflicting information and must defend its choice using a decision log. Week 7 requires a revised experiment and peer audit; the reviewing team attempts to find hidden assumptions, privacy risks and misleading metrics.

Week 8 culminates in a venture action brief, a documented prototype or storyboard, a test report and a reflection explaining when the agent should not be used. The capstone is intentionally smaller than a complete startup plan. Its purpose is to show credible movement from a problem to a tested decision. A participant who rejects automation after comparing cost, uncertainty and risk may meet the learning objectives more effectively than one who builds a complex but unnecessary agent.

5.3. Illustrative Instructional Scenario

Consider an illustrative small service firm receiving repetitive customer booking questions. This is a teaching case, not a real organization or reported implementation. A student team identifies slow replies as a possible source of lost bookings, but competing explanations include confusing appointment availability and inconvenient payment options. The agent first summarizes a synthetic inquiry dataset and drafts response categories. The team inspects examples, corrects misclassification and maps the process. A human gate is placed before any customer-facing reply.

The entrepreneurial decision is whether to build a reply-drafting workflow, redesign the booking page or interview customers before spending. Students score each choice against feasibility, expected customer value, reversibility, privacy exposure and evidence gaps. They then choose one low-cost test and state the observation that would invalidate their rationale. The final submission includes the revised problem statement, source and data provenance, process map, agent error log, approval protocol and a dated decision record. This scenario illustrates how commercial reasoning and automation engineering can be assessed together without presupposing that AI delivers a favorable outcome.

5.4. Instructor Preparation and Equity of Access

The intervention requires an instructor to understand the approved environment, guardrails and evaluation criteria before students begin. A pre-course rehearsal should test prompts, tool permissions, connectivity, export logs, and the case materials for hidden sensitive information. When technical support is limited, use staged demonstrations or a mock agent interface rather than asking novices to configure external integrations. A low-bandwidth or offline alternative should allow equivalent assessment of problem framing, workflow logic, source criticism and decision quality.

Equity also concerns whose business experiences are recognized. Examples should include ventures beyond technology startups and should avoid requiring access to paid datasets, expensive subscriptions or personal business contacts. Teams should receive multilingual support where feasible and learn to check whether an agent's recommendations presume a particular income, infrastructure or regulatory context. These accommodations are not evidence of equal outcomes; they are design features to be tested for accessibility and fairness.

6. Assessment and Proposed Empirical Evaluation

6.1. Assessment Logic and Performance Indicators

A credible evaluation must determine whether learners can perform entrepreneurial work more thoughtfully, not simply whether they enjoy AI tools. This criterion is consistent with DeTienne and Chandler's (2004) emphasis on teachable opportunity-identification performance rather than attitude alone. The primary assessment should consist of blind-scored, scenario-based tasks covering opportunity quality, evidence handling, automation design, risk identification and reasoned decisions. At least two trained assessors should use a common rubric on anonymized artifacts, with disagreements resolved through documented adjudication. Each criterion can be scored on a four-level anchored scale from inadequate to advanced; scale anchors must specify observable behavior rather than broad labels such as 'excellent innovation'.

Opportunity recognition is evaluated by the specificity of the target problem, the distinctiveness and plausibility of alternatives, the quality of supporting evidence, and the articulation of disconfirming tests. Automation competence is evaluated through workflow clarity, error handling, permission boundaries, auditability and an appropriate choice between agentic and simpler approaches. Decision quality is judged against declared criteria, uncertainty recognition, attention to tradeoffs, and the adequacy of the proposed next experiment. A learner should not be rewarded for reaching the same substantive choice as an assessor when different options remain defensible.

Complementary measures include a pre/post AI literacy test, a self-efficacy measure adapted only after checking its validity for the context, learner reflection and instructor observation. McGee et al. (2009) developed a multidimensional entrepreneurial self-efficacy measure, but high self-reported confidence cannot substitute for demonstrated competence. Table 4 specifies proposed measures and their limitations. No scores, coefficient estimates or significance tests are presented because the programme has not yet been implemented as a study.

Table 4. Proposed Evaluation Measures and Interpretation Limits

Outcome domain	Suggested assessment	Important limitation
Opportunity quality	Blind-rated problem and validation plan	Strong prose can hide weak evidence.
Automation competence	Workflow and failure-case rubric	A sandbox demo is not real-world reliability.
Decision quality	Scenario-based decision log	Several choices may be defensible.
Calibrated reliance	Identify erroneous agent suggestions	Performance can depend on task difficulty.
AI literacy and agency	Pre/post knowledge and applied test	Self-report does not prove competence.

Transfer and durability	Delayed unfamiliar case assessment	Attrition and repeated testing can bias results.

6.2. Proposed Comparative Study

A future pilot could use a controlled, mixed-methods design across comparable entrepreneurship course sections. After baseline measurement, classes could be assigned to an E-ACT condition or a comparison condition providing conventional experiential entrepreneurship activities with a standard AI-literacy briefing but no agent workflow. Random assignment at class level would reduce contamination where feasible, although the number of clusters must be sufficient for meaningful inference. If randomization is not possible, matching on prior entrepreneurial exposure, technical experience, initial opportunity-recognition performance and programme level would improve comparability but would not remove unobserved selection bias.

Assessments should occur before instruction, immediately after the programme and after a delayed interval, such as eight to twelve weeks. Parallel task forms and blind grading are necessary to reduce rehearsal effects and expectation bias. An analysis plan should specify the primary outcome in advance, report uncertainty intervals, account for baseline performance and clustering, and distinguish exploratory outcomes from confirmatory tests. The required sample size should be calculated from the expected effect, clustering, attrition and desired precision rather than invented for this paper. Qualitative interviews and artifact analyses could explain why some participants improve, over-rely on agents or reject automation.

Transfer should be tested through an unfamiliar problem rather than a replay of the capstone case. One condition might require students to diagnose a failed agent workflow; another might present an apparently attractive opportunity supported by weak evidence. To test calibrated reliance, investigators can compare students' performance with reliable and deliberately erroneous agent suggestions, with disclosure and consent appropriate to the research protocol. Such tasks can expose whether students are learning the purpose of verification or simply memorizing a classroom approval sequence.

6.3. Fidelity, Governance and Research Ethics

Implementation fidelity is a separate construct from learning outcomes. Researchers should document instructor training, module completion, actual contact time, platform differences, availability of feedback, permission settings and deviations from the syllabus. Agent logs can record proposed tool calls and approval decisions, but they should avoid capturing unnecessary personal information. The programme should specify retention limits, role-based access and a process for reporting unexpected tool behavior. A risk assessment must precede any live connection to organizational systems. Table 5 summarizes the associated safety controls for student exercises.

Table 5. Safety Controls for Student Agent Exercises

Risk Scenario	Preventive Control	Required Response
False external claim	Source citation and cross-check	Withhold claim; correct or mark uncertain.
Personal/customer data leakage	Synthetic or permissioned data; least privilege	Stop workflow and report exposure.
Unauthorized outreach	External send functions disabled by default	Human review before any publication or sending.
Erroneous tool action	Sandbox data, permission boundary, event log	Revert test state; document cause.
Automation bias	Independent criteria before agent recommendation	Record reasons to reject or modify.
Unnecessary system complexity	Compare with manual/rule-based alternatives	Select simpler approach when appropriate.

Participants should be informed of which activities are assessed, whether agent interactions are recorded, how data will be used, and how they may complete equivalent tasks without disclosing business-sensitive material. Institutional ethical review may be required if the intervention is undertaken as human-participant research. Instructors should prohibit fabricated interview transcripts

presented as real fieldwork and should require declaration of meaningful AI assistance in submitted artifacts. These safeguards support educational integrity while allowing productive use of AI.

6.4. Testable Propositions

The framework yields four propositions suitable for preregistered investigation. P1: learners receiving integrated agent-enabled experiential training will show greater gains in independently scored opportunity-evaluation tasks than learners receiving otherwise comparable conventional instruction. P2: requiring source audits and decision logs will reduce acceptance of unsupported agent recommendations in transfer tasks. P3: supervised workflow prototyping will improve the clarity and safety of process specifications, although it may not improve all business performance indicators. P4: structured reflection will improve learners' ability to identify occasions when no AI intervention is warranted. These propositions should be treated as hypotheses, not study findings. An appropriately designed evaluation may reject any of them.

7. Discussion and Implications

7.1. Theoretical Contribution

The proposed framework contributes by positioning AI agents within the mechanism of experiential learning instead of treating their presence as an educational intervention in itself. Tools expand the set of actions students can attempt, but learning depends on how experience is interpreted, challenged and revisited. The emphasis on reflection and collaborative experimentation draws on dynamic entrepreneurial learning and simulated practice (Cope, 2005; Pittaway & Cope, 2007). The transition from generating an idea to defending a business decision is particularly important because the apparent sophistication of an agent may obscure superficial opportunity analysis. E-ACT therefore links the entrepreneurial practice tradition of Neck and Greene (2011) with Kolb's (1984) learning cycle and the competence orientation of EntreComp (Bacigalupo et al., 2016).

A second contribution is the explicit distinction between decision support and delegated authority. Entrepreneurial tasks differ in reversibility and stakes. Asking an agent to summarize nonconfidential feedback is unlike approving spending, publishing claims, handling personal records or contacting customers. In the latter settings, the entrepreneurial act includes accepting responsibility for consequences. The framework makes responsibility visible through decisions, permissions and post-action verification. This aligns with the cautious interpretation of human-AI complementarity suggested by Vaccaro et al. (2024) and the bounded partnership examined by Ahmad et al. (2026).

A third contribution concerns evidence of learning. Because entrepreneurship involves uncertainty, assessment cannot be reduced to matching the instructor's preferred answer. A defensible rejection of an AI proposal may indicate stronger capability than a technically polished but unsupported implementation. The proposed assessment accordingly values articulated assumptions, verification, risk discrimination and appropriate next experiments. It also allows measurement of negative transfer, including automation bias and deskilling, which enthusiasm-based surveys may overlook.

7.2. Practical and Policy Implications

For universities, the framework offers a modular route to embed agentic AI within existing entrepreneurship courses without creating a separate programming requirement. Departments can introduce the opportunity and decision strands before offering agent tools. Such sequencing recognizes that digitalization transforms entrepreneurial agency and uncertainty and therefore should be taught as a strategic issue (Nambisan, 2017). Incubators can substitute venture-specific cases while preserving the same approval and evidence standards. Professional training providers can use the modules to teach owners and managers how to evaluate an automation proposal before committing scarce resources.

Institutions should establish a small set of approved use cases, safe datasets and reusable scoring rubrics rather than relying on individual instructors to negotiate privacy and tool permissions independently. Partnership with information technology staff and local businesses can improve realism, provided the educational setting is not used to pilot unreviewed systems on customer data. Funding should cover instructor development and assessment calibration, not only tool subscriptions.

The framework is intentionally vendor-neutral. Learners should be taught to describe inputs, decisions, validation and human authority in a way that remains meaningful when a commercial platform changes. This reduces the risk that entrepreneurship education becomes short-lived product training. In resource-constrained settings, paper-based process maps, synthetic examples and

limited-access demonstrations can preserve much of the reasoning pedagogy, although actual automation proficiency will depend on hands-on tool access.

7.3. Boundary Conditions and Potential Unintended Effects

Agentic tools may produce uneven benefits. Participants with stronger domain knowledge may detect flawed market assumptions more readily, while learners with limited experience may be more vulnerable to fluent but inaccurate recommendations. Technical barriers, language support, connectivity, subscription costs and disabilities may affect participation. The same student can show strong prompt fluency but weak evidence appraisal. This is consistent with the need to distinguish empowered learners from AI-directed instruction and to assess both the benefits and the limits of general-purpose generative AI in education (Lo, 2023; Ouyang & Jiao, 2021). Programme designers should therefore avoid treating tool-use frequency as a proxy for capability.

There are organizational boundary conditions as well. Firms differ in process standardization, data quality, consent requirements and tolerance for errors. A workflow that appears efficient in a simulated case may be unsuitable when customer expectations, cybersecurity obligations or regulatory constraints are introduced. The curriculum should teach a 'do not automate' decision as a legitimate outcome. Cost-benefit judgments must include monitoring, correction, maintenance, failure recovery and the opportunity cost of entrepreneurial attention.

Finally, the design may produce risks of its own. A capstone centered on AI can crowd out face-to-face customer discovery. Standardized agent prompts may homogenize venture ideas. Automated feedback may favor conventional reasoning, and team projects may conceal unequal contribution. Instructors can mitigate these risks through authentic field evidence, independent reflection, varied cases and evaluator scrutiny; whether those measures are sufficient requires investigation.

8. Limitations and Future Research

The article proposes and rationalizes a curriculum but does not demonstrate its effectiveness. The literature synthesis is purposive rather than exhaustive, and the resulting model may reflect assumptions from higher education settings more strongly than those of informal enterprises or vocational programmes. The eight-week duration, training-hour allocation and six-competency structure have face validity as design choices, but they have not been optimized or psychometrically validated. The figures depict a proposed process, and the chart depicts planned instructional hours; none reports participant data.

A second limitation is the speed of change in AI-agent platforms. Tool interfaces, data policies, available safeguards and operating costs can shift faster than curricula. No single platform configuration should be assumed to meet an institution's privacy, security or accessibility requirements. The approach also assumes that educators can supervise learners' judgment and inspect agent outputs, which may be difficult in large cohorts without additional support.

Future research should test the framework across institutions, sectors and levels of prior experience, compare agent-enabled instruction with less complex AI or manual alternatives, and examine delayed transfer. Additional work should validate scoring rubrics, estimate educator workload and explore whether learners become appropriately skeptical rather than uniformly distrustful. Studies could investigate which forms of human approval genuinely improve decision quality and when automated scaffolds reduce independent reasoning. Publication of task materials, rubrics and deidentified protocols would support reproducibility.

9. Conclusion

AI-agent-enabled entrepreneurship education should not be judged by how many tasks an agent performs for a student. Its value depends on whether the student can recognize a credible opportunity, specify an appropriate business process, challenge unreliable information and retain responsibility for consequential decisions. This paper has developed E-ACT, a theory-informed experiential framework combining six competencies, three practice strands, an eight-week curriculum, a bounded agent workflow and a proposed evaluation protocol.

The framework treats AI output as material for inquiry rather than a substitute for entrepreneurial judgment. Its central teaching sequence links real or realistic problems with verification, reflection, controlled action and reassessment. The proposed rubrics and comparative design make its claims open to empirical examination. Until such testing is completed, improved

performance, transfer and safe adoption remain hypotheses. The practical task for educators is to create conditions in which learners can use AI agents competently while still learning to question them and decide for themselves.

References

[1] Ahmad, A., Amry, D., Oakley, D., McInerny, J., & Napitupulu, M. (2026). AI as intern, not oracle: A design logic for human-AI partnership in experiential entrepreneurship education. *Entrepreneurship Education*. https://doi.org/10.1007/s41959-026-00196-7

[2] Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1

[3] Bacigalupo, M., Kampylis, P., Punie, Y., & Van den Brande, G. (2016). *EntreComp: The entrepreneurship competence framework* (EUR 27939 EN). Publications Office of the European Union. https://doi.org/10.2791/593884

[4] Baron, R. A. (2006). Opportunity recognition as pattern recognition: How entrepreneurs 'connect the dots' to identify new business opportunities. *Academy of Management Perspectives, 20*(1), 104–119. https://doi.org/10.5465/amp.2006.19873412

[5] Chandler, G. N., DeTienne, D. R., McKelvie, A., & Mumford, T. V. (2011). Causation and effectuation processes: A validation study. *Journal of Business Venturing, 26*(3), 375–390. https://doi.org/10.1016/j.jbusvent.2009.10.006

[6] Cope, J. (2005). Toward a dynamic learning perspective of entrepreneurship. *Entrepreneurship Theory and Practice, 29*(4), 373–397. https://doi.org/10.1111/j.1540-6520.2005.00090.x

[7] Darnell, J. A., & Gopalkrishnan, S. (2024). Entrepreneurship teaching exercises: Integrating generative AI. *Discover Education, 3*, Article 172. https://doi.org/10.1007/s44217-024-00261-0

[8] Davidsson, P. (2015). Entrepreneurial opportunities and the entrepreneurship nexus: A re-conceptualization. *Journal of Business Venturing, 30*(5), 674–695. https://doi.org/10.1016/j.jbusvent.2015.01.002

[9] DeTienne, D. R., & Chandler, G. N. (2004). Opportunity identification and its role in the entrepreneurial classroom: A pedagogical approach and empirical test. *Academy of Management Learning & Education, 3*(3), 242–257. https://doi.org/10.5465/amle.2004.14242103

[10] Duong, C. D., & Nguyen, T. H. (2024). How ChatGPT adoption stimulates digital entrepreneurship: A stimulus-organism-response perspective. *The International Journal of Management Education, 22*(3), Article 101019. https://doi.org/10.1016/j.ijme.2024.101019

[11] Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–105. https://doi.org/10.2307/25148625

[12] Hmelo-Silver, C. E. (2004). Problem-based learning: What and how do students learn? *Educational Psychology Review, 16*(3), 235–266. https://doi.org/10.1023/B:EDPR.0000034022.16470.f3

[13] Kolb, A. Y., & Kolb, D. A. (2005). Learning styles and learning spaces: Enhancing experiential learning in higher education. *Academy of Management Learning & Education, 4*(2), 193–212. https://doi.org/10.5465/amle.2005.17268566

[14] Kolb, D. A. (1984). *Experiential learning: Experience as the source of learning and development*. Prentice-Hall.

[15] Lo, C. K. (2023). What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences, 13*(4), Article 410. https://doi.org/10.3390/educsci13040410

[16] Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). Association for Computing Machinery. https://doi.org/10.1145/3313831.3376727

[17] McGee, J. E., Peterson, M., Mueller, S. L., & Sequeira, J. M. (2009). Entrepreneurial self-efficacy: Refining the measure. *Entrepreneurship Theory and Practice, 33*(4), 965–988. https://doi.org/10.1111/j.1540-6520.2009.00304.x

[18] McMullen, J. S., & Shepherd, D. A. (2006). Entrepreneurial action and the role of uncertainty in the theory of the entrepreneur. *Academy of Management Review, 31*(1), 132–152. https://doi.org/10.5465/amr.2006.19379628

[19] Miao, F., Shiohira, K., & Lao, N. (2024). *AI competency framework for students*. UNESCO. https://doi.org/10.54675/JKJB9835

[20] Nabi, G., Liñán, F., Fayolle, A., Krueger, N., & Walmsley, A. (2017). The impact of entrepreneurship education in higher education: A systematic review and research agenda. *Academy of Management Learning & Education, 16*(2), 277–299. https://doi.org/10.5465/amle.2015.0026

[21] Nambisan, S. (2017). Digital entrepreneurship: Toward a digital technology perspective of entrepreneurship. *Entrepreneurship Theory and Practice, 41*(6), 1029–1055. https://doi.org/10.1111/etap.12254

[22] Neck, H. M., & Greene, P. G. (2011). Entrepreneurship education: Known worlds and new frontiers. *Journal of Small Business Management, 49*(1), 55–70. https://doi.org/10.1111/j.1540-627X.2010.00314.x

[23] Neck, H. M., Greene, P. G., & Brush, C. G. (2014). *Teaching entrepreneurship: A practice-based approach*. Edward Elgar Publishing. https://doi.org/10.4337/9781782540564

[24] Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence, 2*, Article 100041. https://doi.org/10.1016/j.caeai.2021.100041

[25] Ouyang, F., & Jiao, P. (2021). Artificial intelligence in education: The three paradigms. *Computers and Education: Artificial Intelligence, 2*, Article 100020. https://doi.org/10.1016/j.caeai.2021.100020

[26] Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (Article 2, pp. 1–22). Association for Computing Machinery. https://doi.org/10.1145/3586183.3606763

[27] Pittaway, L., & Cope, J. (2007). Simulating entrepreneurial learning: Integrating experiential and collaborative approaches to learning. *Management Learning, 38*(2), 211–233. https://doi.org/10.1177/1350507607075776
[28] Renz, P. A., Martin, J., & Flad, M. (2025). Large language models as a catalyst for opportunity recognition capability. *International Journal of Entrepreneurial Behavior & Research*. https://doi.org/10.1108/IJEBR-08-2024-0815
[29] Sarasvathy, S. D. (2001). Causation and effectuation: Toward a theoretical shift from economic inevitability to entrepreneurial contingency. *Academy of Management Review, 26*(2), 243–263. https://doi.org/10.5465/amr.2001.4378020
[30] Shane, S., & Venkataraman, S. (2000). The promise of entrepreneurship as a field of research. *Academy of Management Review, 25*(1), 217–226. https://doi.org/10.5465/amr.2000.2791611
[31] Shinn, N., Cassano, F., Gopinath, A., Narasimhan, K., & Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems, 36*, 8634–8652. https://doi.org/10.52202/075280-0377
[32] Triantoro, T., Rua, T., & Lang, G. (2025). AI skills for entrepreneurs: A practical experiential learning approach. *Information Systems Education Journal, 23*(3), 31–40. https://doi.org/10.62273/MLUM6680
[33] Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour, 8*, 2293–2303. https://doi.org/10.1038/s41562-024-02024-1
[34] Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science, 18*, Article 186345. https://doi.org/10.1007/s11704-024-40231-1
[35] Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*. https://openreview.net/forum?id=WE_vluYUL-X