Rieda edishan | 2026-11-20

LAC AI Playbook. From regional knowledge to everyday work

Adrian Dunkley | 2026-11

Di guide join local information, weh di business know an one clear task. Di country example dem show how a team can plan a pilot. Dis review edition nuh report field-study result.

Adrian Dunkley, President a di Caribbean AI Association. November 2026, review edition 1.2. Policy synthesis an a proposed implementation method. Di paper nuh report a new regional field study.

Executive summary. Start wid di work

Adrian Dunkley playbook join a regional context prompt, business knowledge base an one defined task. Di country an sector guide help reader find possible use. Di safety tool ask who can act, weh can go wrong an weh evidence justify release. [1–3]

Dem part ya form a proposed adoption method fi small business an di organisation weh support dem. A public agency, educator or investor can use it if dem adapt di control to di people an decision affected.

Fund an assess di whole workflow, including preparation, staff training, review, correction, supervision an fallback if di tool fail. Compare it wid di existing way fi do di same task.

Di paper draw pon di practical guide, regional evaluation proposal an safety toolkit. Outside research provide context, no proof it work fi every country or sector. Country example show possible pilot setting. Dem nuh report completed deployment or promise financial return, fewer error or time saving.

Owner an manager can use di pilot plan. Business-support agency can use di institutional priority dem. Technical reviewer shoulda read di testing an permission requirement. Original map: regional evidence p. 3, context, knowledge an task p. 4, test an safety pp. 5–6, institutional support an pilot pp. 7–8, limit an source p. 9.

Research question. Can a locally specified, human-reviewed AI workflow cut total task time widout exceeding a agreed error or harm threshold? Compare wid di same task under di existing workflow. Count review an rework time.

1. Regional evidence. Exposure an outcome

Inna 2024, ILO an World Bank estimate seh 26–38% a job inna Latin America an di Caribbean coulda face generative AI exposure. Dem distinguish potential productivity gain from full automation potential. Digital access can keep plenty worker from getting di benefit. [4]

Figure 1. Share a employment from published 2024 estimate. Total exposure 26–38%. Productivity-enhancing transformation 8–14%. Full automation potential 2–5%. Scale 0, 10, 20, 30 an 40%. Source: Gmyrek, Winkler an Garganta, ILO Working Paper 121, 31 July 2024. Bar show published range, no confidence interval. Di two lower category part a exposure but dem nuh add to di total. Di model estimate potential. It nuh measure adoption, realised productivity or predicted job loss, an it nuh gi a separate estimate fi every Caribbean economy.

A owner can select task fi controlled comparison. Di figure nuh justify replacing staff or assuming a vendor general benchmark apply to di business. Policymaker shoulda put access, training an job quality inna di same programme. Count who can use di service an who cyaan. Report small-island an rural result separate when di sample support it.

2. Operating method. Gi di tool a local brief

Regional context. Include country, language, currency, payment method, delivery constraint an relevant authority. Use fact bout dis operation. Avoid assumption bout di whole region.

Business knowledge. Use approved product detail, opening hour, service term, price list an document template. Name owner an review date. Remove record weh di tool nuh need.

Task instruction. Specify di output, reader, allowed source an decision weh stay wid a person. State weh fi do if information missing or contradictory.

Keep di brief small enough fi maintain. If staff cyaan tell which price list current, di knowledge base become a risk. Start wid a short approved set. Mark effective date, superseded version an who can change it. A local prompt can reduce assumption. Check law, arithmetic an language against di source or wid a qualified reviewer.

Show uncertainty. Require di tool fi flag missing detail before it prepare a customer-ready answer. A booking reply cyaan invent availability, a quote cyaan invent tax treatment, an a grant application cyaan invent achievement. Keep confidential customer an staff record outside general-purpose tool unless di organisation approve di specific data use, access an retention.

Guesthouse example. Draft a reply from di current room list, approved rate an cancellation policy. Use di guest requested language. Flag missing date. Nuh reserve a room, promise a refund or send di reply. Di owner check di draft against di booking system. Dis a proposed workflow from di hospitality an context guide. It nuh have case-study evaluation yet.

3. Evaluation. Test di task inna it setting

Di proposed regional benchmark cover language, regional context, small-business task, cost an low-bandwidth operation. Dem dimension guide test selection. Ranking need inspectable prompt, output, score record an reviewer procedure. [2]

Figure 2, step 1. Time di current workflow. Keep sample a correct output an common error. Define acceptable result before di AI trial.

2. Use representative case: routine work, ambiguous request, missing information an case weh di system shoulda refuse or refer. Use authorised or synthetic record.

3. Keep difficulty an scoring rule comparable. Record version, instruction an setting. Keep familiarisation separate from measured work.

4. Count preparation, checking an correction time. Check failure by language, task an affected group. Decide fi stop, revise or extend. Di sequence a proposal from di task guidance an regional evaluation design [1–2]. It never get empirical validation yet.

Track total time per accepted task, correction rate an cost per accepted output. Keep serious error separate from small wording change. A lower average time cyaan excuse a harmful failure. Ask people weh know di language an task fi review output. Record disagreement an resolution. Translation score alone cyaan establish seh a reply suitable fi customer.

4. Safety an accountability. Control weh di system can do

Adrian Dunkley proposed educational S4 method organise review under Specify, Stress-test, Safeguard an Scrutinise. A person enter di toolkit risk index fi organise review. It nuh gi probability a harm or security certificate. [3, 5]

Specify. Record purpose, user, data, permitted action an named owner. Example: allow drafting but block sending.

Stress-test. Record case, output, failure an unresolved question. A outside document must never change payment instruction.

Safeguard. Record access limit, approval step an tested manual fallback. Restrict access to di record weh di task need.

Scrutinise. Record review date, incident, change an retest result. Repeat relevant test after model or permission change.

Separate content from authority. OWASP explain how indirect prompt injection bring instruction through outside content like a webpage or file. A connected assistant can meet it inna ordinary work. Limit privilege an require independent approval fi consequential action. A sentence inna prompt nuh control access. Di application haffi enforce which record an tool available. [6]

Name who can suspend di workflow. Record how fi remove access, preserve relevant log an return to di manual process. Protect personal data inna incident record. NIST voluntary framework support context-specific risk management through di system lifecycle. Apply national an sector duty separate. [7]

Proposed release condition. Nuh enable a consequential action while authority, review or failure response unresolved. A low average score cyaan cancel critical control failure.

5. Institutional priority fi real constraint

Small-business support agency. Run task clinic weh owner bring one recurring problem an leave wid a bounded pilot. Help dem prepare approved record, compare total cost an pick review method. Nuh count tool demo as business outcome.

Government an regulator. Publish usable procurement requirement fi data handling, access, performance evidence, human review an exit. Ask supplier fi demonstrate local language an connectivity. Keep complaint an redress route clear. A system explanation cyaan remove di public body responsibility fi a decision.

Educator an training provider. Teach learner fi define task, inspect evidence, recognise error an know when fi seek help. Assess independent judgement alongside di ability fi produce output wid AI.

Investor an lender. Ask fi evidence from completed workflow instead a headline model score. Examine recurring cost, customer retention, supervision effort an supplier dependence. Tell a successful demo apart from a repeatable operation.

Regional institution. Coordinate reusable test material an reviewer training. Maintain country-specific entry wid date an source link. A shared format can reduce duplication widout assuming identical law, infrastructure or customer expectation.

Resilience an access. Include low-bandwidth route an offline fallback where feasible. Local hosting can reduce some transfer but still need secure device, maintenance, access control an task test. Smaller nuh automatically mean safer or cheaper.

Start wid a service weh already have reliable record an manageable review load, like booking enquiry, product description or document formatting. Keep clinical, credit an other consequential decision outside simple starter pilot. Dem a proposed choice weh need local feasibility an legal review. No finding seh one sector must always go first.

6. Implementation. Run a pilot wid a clear decision

Week 1–2, define. Select one task an name owner. Map data an permission. Record di manual baseline. Produce approved scope, baseline record, cost limit an stop condition.

Week 3–4, prepare. Clean di knowledge base, train reviewer an assemble routine an failure case. Produce dated source, test register an agreed scoring rule.

Week 5–8, compare. Run a limited supervised trial. Record review time, correction an incident. Keep paired or otherwise comparable task record an explain exclusion.

Week 9–12, decide. Review result an cost. Retest change. Stop, revise or extend. Keep signed decision, unresolved limit, next review date an rollback owner.

Set threshold before trial. Adjust di illustrative schedule fi risk, staffing an available representative case. Define acceptable accuracy fi each task, serious error, maximum review load an budget. Cover important failure mode. No universal sample size. If sample too small fi estimate rare harm, say so. A spell widout incident nuh prove safety.

Keep a short decision record. State weh get tested, how it compare wid baseline, weh fail an weh still unknown. Explain who benefit an who face extra work or risk. Fi expansion specify added user, data an permission. Check if di change invalidate previous evidence.

Continue only if di workflow meet pre-agreed quality an control requirement an have worthwhile benefit after checking, correction an operating cost. Otherwise revise or return to di old process. Documented local trial wid comparable outcome can supply di next regional evidence. Publishing failure an boundary condition help other organisation avoid repeating dem.

Source an evidence limit

Di conceptual policy synthesis join context/knowledge/task, regional evaluation design an S4 safety. Di employment figure reproduce dated external estimate. No new experiment, systematic review or regional adoption survey done. Model superiority, universal time saving an country-wide readiness each need separate validation. Di paper use di regional test design widout di unverified performance table. Pilot stage an release condition a author recommendation, no official regulatory standard.

[1] Author-maintained practical guide, regional context prompt an business knowledge-base template. Check 11 November 2026. Source file: guide/README.md, guide/part-1.md, guide/context-prompt.md an guide/knowledge-base.md.

[2] Author-maintained LAC Benchmark v0.1, April 2026, rankings/lac-benchmark.md. It test dimension inform original page 5. Dis paper never reproduce or validate it model score. Find it through di Playbook rankings resource.

[3] Human-reviewed ordinal screening tool. Documentation an engine inspect 11 November 2026.

[4] Joint ILO–World Bank publication, 31 July 2024. Primary page verified 11 November 2026. Modelled exposure nuh measure observed job loss.

[5] Author S4 framework an editorial source register dated 3 November 2026. Proposed educational method, no independent effectiveness claim. Link point to original issue.

[6] OWASP LLM01:2025 primary guidance pon prompt injection, verified 11 November 2026.

[7] Voluntary NIST guidance from 2023 an 2024. Overview checked 11 November 2026 said AI RMF 1.0 under revision. Source list include framework page an profile DOI.

Original design adapt di type hierarchy, A4 page structure an publication furniture a CAIA Learning on Our Terms white paper fi black an white. No United Nations affiliation or endorsement implied.

Suos dem

[1] Dunkley, A. AI Playbook for the Caribbean and Latin America.

[2] Dunkley, A. The LAC Benchmark, version 0.1 (April 2026). Playbook repository document rankings/lac-benchmark.md; see Playbook rankings.

[3] Dunkley, A. AI Safety & Alignment Toolkit.

[4] Gmyrek, P., Winkler, H. and Garganta, S. (2024). Buffer or Bottleneck? Employment Exposure to Generative AI and the Digital Divide in Latin America. ILO Working Paper 121.

[5] Dunkley, A. S4 AI Safety Framework.

[6] OWASP Gen AI Security Project (2025). LLM01:2025 Prompt Injection.

[7a] NIST (2023). AI Risk Management Framework 1.0.

[7b] NIST (2024). Generative Artificial Intelligence Profile, NIST AI 600-1.

Inglish edishan (PDF)