AI-SLANDS WORKING PAPER NO. 9

FROM TESTBED TO CHARTER

Innovation, Democratic Risk and the Governance of AI in Island Societies

Andrew Jennings

Professor in Island Studies, University of the Highlands and Islands

September 2026

AI-slands Working Paper No. 9

ABSTRACT

Islands are repeatedly presented as promising sites for artificial-intelligence experimentation. Their bounded jurisdictions, concentrated institutions, identifiable populations and urgent service challenges appear to offer conditions in which new systems can be tested quickly and visibly. The same characteristics can also magnify democratic risk. A pilot conducted across a small public service may reach a large share of the population; a single contract may establish long-term dependence; and close relations among government, business and community can enable coordination while making dissent difficult. This paper examines the difference between an island-led experiment and the treatment of an island as a convenient laboratory. It draws on international work on regulatory sandboxes and AI governance, together with developments in Malta, Singapore, Mauritius and Hawai'i. The cases show that experimentation can build local capability and inform regulation, but only when public purpose, legal responsibility, community participation, independent evaluation and credible exit arrangements are designed before deployment. The paper proposes an Island AI Experiment Test and concludes the AI-slands working-paper series with an Island AI Charter. The Charter brings together ten commitments concerning authority, consent, cultural integrity, data rights, transparency, equitable benefit, resilience, environmental responsibility, reversibility and inter-island cooperation. Its purpose is not to oppose experimentation. It is to ensure that islands enter experiments as political communities with the power to define the problem, govern the process, stop harmful uses and retain a fair share of what is learned and created.

Keywords: artificial intelligence; islands; testbeds; regulatory sandboxes; democracy; public participation; data governance; island sovereignty

1. INTRODUCTION: WHOSE EXPERIMENT?

Islands have long been treated as places where the wider world can observe change in concentrated form. They have served as sites of scientific investigation, military testing, conservation, tourism development, fiscal experimentation and administrative reform. Their apparent boundedness makes relationships easier to trace. Their small scale can shorten the distance between an idea, a decision and its effects. Godfrey Baldacchino described islands as "novelty sites": places at the edge from which dominant assumptions can be questioned and alternative arrangements imagined (Baldacchino, 2007). That capacity for innovation is real. So is the danger that an island becomes valuable mainly because someone elsewhere regards it as manageable.

Artificial intelligence gives the old laboratory metaphor a new institutional form. Governments are encouraged to pilot automated services, create regulatory sandboxes, open public datasets and attract firms seeking a controlled environment in which to demonstrate technology. An island jurisdiction may be able to bring ministers, officials, universities, firms and civil-society organisations around one table more readily than a large state. It may have a national identification system, a concentrated health service, a single electricity grid or a relatively coherent legal framework. These features can support learning. They can also allow one experimental system to become infrastructural before its consequences have been properly contested.

The central question is therefore not whether islands should experiment with AI. Refusal of every experiment would leave island societies dependent on systems designed and tested elsewhere. The relevant question is who defines the experiment and on what terms. An island-led test begins with a problem recognised by island institutions and communities. It assigns public responsibility, protects people who bear risk, produces evidence that can be challenged, and leaves the jurisdiction with greater capability. An extractive test begins with a technology or market opportunity. It treats access to a population, dataset or regulatory environment as the principal asset, while decisions about design, intellectual property and future use remain external.

This distinction completes the argument developed across the AI-slands working-paper series. Paper 1 proposed an island-centred approach to AI and warned that digital connection can reduce

distance while increasing dependence (Jennings, 2026a). Paper 2 examined culture, language, extraction and the right to refuse (Jennings, 2026b). Paper 3 addressed governance and the right of small jurisdictions to shape their digital future (Jennings, 2026c). Papers 4 and 5 considered public services, work and demography (Jennings, 2026d; 2026e). Paper 6 examined infrastructure and governed interdependence, while Paper 7 showed how AI can reproduce inequalities within archipelagos (Jennings, 2026f; 2026g). Paper 8 argued that prediction becomes resilience only when island institutions can act (Jennings, 2026h). The present paper asks what follows when these concerns meet the politics of experimentation.

The argument has two parts. The first examines why islands are attractive testbeds, what can be gained from bounded experimentation, and how democratic risk arises. The second proposes governance tools: an Island AI Experiment Test for proposed pilots and an Island AI Charter for the longer term. The Charter is deliberately practical. Principles matter, but principles without institutions can become a substitute for accountability rather than its foundation (Mittelstadt, 2019). Each commitment is therefore linked to a public obligation, a decision point or a right that an affected person or community can exercise.

This is a working paper rather than a model law. Constitutional arrangements, legal powers and Indigenous rights differ across island jurisdictions. A sovereign small island state, a subnational archipelago and a non-self-governing territory cannot use identical instruments. The shared task is narrower and more fundamental: to make experimentation answerable to the people and places through which it acquires value.

2. WHY ISLANDS BECOME TESTBEDS

The case for island experimentation usually begins with scale. A smaller public administration can sometimes coordinate across departments without the layers found in a large state. A limited number of hospitals, ports, schools or utilities may make a pilot easier to organise. Feedback can travel quickly between frontline staff and senior decision-makers. The whole service may be visible enough for researchers to study interactions that would disappear within national aggregates. These are important advantages when an AI system depends on organisational change rather than software alone.

Boundaries also attract researchers and firms. Islands seem to offer a definable population, network or environment. A transport model can be tested on a finite road and ferry system. A public-service assistant can be introduced within a relatively unified administration. Environmental monitoring can connect observations to a recognisable reef, fishery or watershed. The appeal is methodological: an experiment appears to have clearer inputs, outputs and limits. In practice, island systems are open. People, data, capital, energy, platforms and legal obligations cross the shoreline continuously. The boundary is politically useful, but it is rarely technically complete.

Small jurisdictions may also be motivated to move quickly. Economic diversification is a persistent policy objective in many islands. A government may view a national AI strategy, innovation zone or sandbox as a way to attract investment, retain graduates and gain visibility in international policy networks. Malta's 2019 strategy used the language of the "ultimate AI launchpad" and linked adoption to an innovation ecosystem and regulatory capacity (Government of Malta, 2019). Singapore's national strategies similarly combine public-service adoption, research capability, industry development and international influence, though with resources and state capacity that few island jurisdictions can match (Smart Nation Singapore, 2023; 2026).

Urgency is another driver. Islands face service shortages, climate hazards, language-resource gaps and large maritime responsibilities. These are not invented problems. AI tools may improve forecasting, help allocate scarce professional time or enable analysis that a small team could not perform manually. Waiting for systems designed around large continental populations can reproduce neglect. Local experimentation may be the only way to establish whether a tool works for a small language, an outer island or a distinctive ecological setting.

Regulatory sandboxes formalise part of this impulse. The OECD defines them as temporary, controlled arrangements in which regulators and participating organisations test innovations that may challenge existing rules or compliance processes (OECD, 2023). A sandbox can give a regulator direct knowledge of a developing technology, help a firm identify legal obligations and generate evidence for future policy. The European Union's AI Act requires member states to establish AI regulatory sandboxes and specifies supervision, safeguards, documentation and conditions for testing in real-world settings (European Union, 2024). The legal model is more demanding than the casual use of "sandbox" to mean an innovation programme with few constraints.

For islands, the attraction of a sandbox is easy to understand. It promises a middle route between prohibition and ungoverned deployment. Yet the metaphor can conceal asymmetry. Children play in a sandbox because someone else has defined its boundaries and made the surroundings safe. In a regulatory sandbox, the regulator must have the expertise, independence and legal authority to set those limits. A small administration negotiating with a global technology company may possess formal authority but lack access to the model, training data, security evidence or specialist staff needed to exercise it. The sandbox then protects the experiment from regulation more effectively than it protects the public from the experiment.

An island can therefore be a testbed in at least four different senses. It can be a research site where a system is evaluated. It can be a market demonstration intended to establish commercial credibility. It can be a regulatory experiment in which rules and technologies are tested together. It can also be a political project through which an island community develops its own capability and priorities. These forms may overlap, but they should not be treated as interchangeable. A demonstration that attracts investment is not necessarily a valid evaluation. A technically successful pilot is not necessarily a legitimate public policy. A regulatory waiver is not evidence that a community has consented.

3. THE VALUE OF BOUNDED EXPERIMENTATION

Experimentation is necessary because AI systems cannot be governed adequately through abstract principles or laboratory benchmarks alone. Their performance depends on data, organisational routines, user behaviour and local conditions. A translation system that performs well on a standard dataset may fail with island dialects or administrative terminology. A benefits-screening model may be statistically accurate overall while making serious errors for small groups. A public-sector assistant may save time on routine correspondence while creating new risks around confidential information and records management. Controlled testing can reveal these interactions before a system becomes difficult to remove.

The strongest experiments begin with a decision that people already understand. The objective might be to reduce the time required to find planning guidance, improve detection of a crop disease, translate a public notice or help meteorologists interpret a forecast. The existing process provides a baseline. The pilot can then compare accuracy, cost, workload, accessibility and distributional effects. The question is whether the complete service improves, including the work transferred to people who verify outputs or respond to errors.

Small scale can make such evaluation unusually rich. Staff and service users may be able to describe effects that a large quantitative study would miss. A limited pilot can be paused while problems are corrected. Institutions can bring technical evidence into conversation with professional and community knowledge. When relationships are durable, learning can continue beyond a procurement cycle. This is an advantage of island proximity, provided that familiarity does not prevent criticism.

Experimentation can also build regulatory capability. The OECD notes that sandboxes can create structured interaction between regulators and innovators and help identify how rules apply to new systems (OECD, 2023). This is valuable when legal categories developed before machine learning do not map neatly onto current practice. A regulator may learn what documentation is realistically

available, where responsibility sits across a supply chain and which metrics reveal harm. Firms may learn that public accountability is a design requirement rather than a final compliance exercise.

Regional experimentation can spread costs. Several island states may need to evaluate similar tools for health, disaster management, customs or education. A shared technical unit could test systems across languages and jurisdictions, maintain procurement clauses and publish common incident reports. Regional organisations already coordinate in areas such as meteorology and fisheries. An AI assurance function could extend that practice without requiring every administration to recruit a complete specialist team.

The benefits, however, depend on what the experiment leaves behind. A pilot that ends with a presentation and an expired cloud account has produced little institutional value. A useful experiment leaves documentation, trained staff, reusable data under appropriate governance, an account of failure and a defensible decision about whether to proceed. It should improve the jurisdiction's ability to evaluate the next proposal, including the ability to say no.

Learning must extend beyond the participating organisation. Public money and public data create a strong case for publishing methods, costs, performance and limitations. Commercial confidentiality may protect specific intellectual property, but it should not conceal whether a public pilot worked or harmed particular groups. Small jurisdictions cannot afford to repeat the same failures in isolation. Shared evidence is one way in which island scale can produce influence rather than dependence.

4. DEMOCRATIC RISK IN SMALL JURISDICTIONS

The democratic risks of AI experimentation are not unique to islands, but island conditions can intensify them. A pilot in one agency may affect a large proportion of residents because there are few alternative providers. The same company may supply cloud services, analytics and consultancy across government. Political and professional networks can be dense. These features support coordination, yet they may also reduce the practical distance between the promoter, regulator and evaluator of a project.

Consent is the first difficulty. Individual consent is often an inadequate basis for public-sector experimentation because people cannot realistically decline essential services. A resident cannot opt out of the only hospital, school system or benefits agency without cost. Clicking an acceptance box does not settle whether an experiment is legitimate. Public authority must instead show that the use has a lawful purpose, that less intrusive alternatives were considered and that affected groups had a meaningful opportunity to shape the decision.

Participation can also become ceremonial. A workshop may be held after the technology, vendor and objectives are already fixed. Community members may be asked to comment on interface language while questions about surveillance, ownership or whether the system should exist are excluded. Sloane and colleagues warn that participation can reproduce extraction when institutions invite people into a process without transferring decision-making power (Sloane et al., 2022). Island closeness does not automatically create democratic participation. It can make non-participation more visible and socially costly.

Scale creates problems of anonymity. In a large dataset, removing names may sometimes provide meaningful protection. In a small community, combinations of occupation, location, age, language or medical condition can identify a person even when direct identifiers are absent. A dataset describing an outer island or a small Indigenous group may be effectively identifiable at collective level. Conventional privacy assessments focused on individuals may miss harms to a community's reputation, bargaining position or control over knowledge.

The issue is particularly important for Indigenous and traditional knowledge. The CARE Principles for Indigenous Data Governance emphasise collective benefit, authority to control, responsibility and ethics (Carroll et al., 2020). Indigenous Protocol and Artificial Intelligence likewise rejects the assumption that data and cultural material become free inputs once digitised; governance must arise

through relationships with specific communities (Lewis et al., 2020). An island experiment that uses language, ecological knowledge or cultural collections needs authority beyond a general open-data licence.

Regulatory capture is a second risk. A small government may rely on a vendor for technical explanation, staff training and evaluation. The firm then helps define the problem, supply the system and describe its success. Even without improper influence, this arrangement narrows what the regulator can see. The public body may lack the source code, model documentation or contractual rights needed for independent scrutiny. A pilot can quietly become procurement because switching costs grow while evidence is being collected.

Temporary exceptions can become durable. Sandboxes are meant to have boundaries, supervision and an end point. In practice, a useful service may gain political support before evaluation is complete. Staff reorganise work around it; data pipelines are created; licences are renewed; and withdrawal begins to look like failure. The experiment acquires what might be called institutional momentum. A clear decision gate is therefore required before any transition from pilot to routine use.

Distributional harm may be hidden by national averages. A system can perform well for the capital island while failing in outer islands with weaker connectivity or different language use. It can save professional time while shifting unpaid verification work to citizens. It can create high-skilled jobs in one centre while using communities elsewhere mainly as data sources. Paper 7 described this as algorithmic peripherality: participation in the system without equivalent authority or benefit (Jennings, 2026g). Every island experiment should report results by place and affected group where this can be done safely.

Public debate can also be constrained by promises of economic opportunity. When a project is associated with investment, jobs or international recognition, criticism may be characterised as resistance to innovation. Small communities often experience development proposals through this language. The democratic standard should be the opposite. The greater the claim of strategic importance, the stronger the case for publishing assumptions, alternatives and obligations.

Finally, an experiment may create knowledge that leaves the island. A firm learns how its product performs in a bounded public system, improves the model and uses the demonstration to enter larger markets. The host jurisdiction may receive a short-term service but no ownership, revenue share or durable access. This is not an argument that every public pilot should produce local intellectual property. It is an argument that the value exchanged should be identified. Access to data, institutions, staff time, reputation and a population is a contribution. Contracts should recognise it.

5. DIFFERENT ISLAND PATHS

Island jurisdictions are not one policy category. Malta, Singapore, Mauritius and Hawai'i differ greatly in constitutional status, wealth, population, administrative capacity and relation to external markets. They are considered here because each illustrates a different relationship between experimentation, governance and institutional capability. None should be treated as a model to copy in full.

5.1 Malta and the launchpad model

Malta's 2019 national strategy explicitly presented the country as an AI launchpad. The approach sought to combine investment attraction, adoption, skills and an ethical and regulatory framework (Government of Malta, 2019). The proposition reflects a familiar small-state strategy: use jurisdictional agility and visibility to enter an emerging sector before larger competitors move. Malta also had experience in establishing specialised digital regulation, which gave the ambition an institutional base rather than leaving it as branding alone.

The launchpad model reveals both opportunity and risk. A small jurisdiction can convene actors quickly and create a coherent route from policy to pilot. It can also become dependent on being perceived as permissive. If national competitiveness rests on lowering friction for innovators, the

regulator may face pressure to treat caution as a reputational cost. The relevant safeguard is not a general promise of ethics. It is operational separation between promotion, authorisation and independent review, supported by publication of sandbox conditions and outcomes.

Malta is also an EU member state and is therefore situated within the AI Act's risk-based framework and sandbox requirements (European Union, 2024). This wider legal setting can provide resources and minimum standards that a small regulator could not easily create alone. It also shows that island agility need not mean isolation. Governed interdependence can strengthen local authority when external law supplies enforceable rights, shared expertise and routes for appeal.

5.2 Singapore and capability at scale

Singapore demonstrates a different form of island experimentation. Its AI strategies combine national coordination, public-sector deployment, research, industry support and governance tools. The 2026 update describes broad AI adoption across government, training for public officers, more than seventy company AI centres of excellence, the AI Verify testing framework and Project Moonshot for evaluating large language models (Smart Nation Singapore, 2026). This is experimentation supported by substantial state capacity.

The case matters because it challenges the assumption that speed requires weak governance. Testing infrastructure, public-service capability and technical expertise can be part of an innovation strategy. Regulators and agencies are better placed to question suppliers when they possess their own tools and staff. Open evaluation resources can also contribute to international practice.

Singapore nevertheless shows why technical assurance is not identical to democratic legitimacy. A testing framework can assess declared governance principles and model behaviour; it cannot decide which public purposes justify deployment or whose account of harm should prevail. Those decisions require legal authority, public reasoning and avenues for challenge. High administrative capability may make deployment safer, but it does not remove the need to debate surveillance, worker autonomy, service access and the distribution of economic benefit.

5.3 Mauritius and institution building

Mauritius launched a National AI Strategy and FAIR Guidelines in April 2026. The guidelines are organised around fairness, accountability, inclusiveness and integrity, and responsibility. UNDP reports that the strategy was developed through consultation involving government, business, startups, academia and civil society (UNDP, 2026). This is significant because it places institutional development and a public ethical framework alongside economic ambition.

At the same time, the case remains at an early stage. A published strategy and consultation process are evidence of intention, not of implementation or impact. The difficult questions arise when a ministry procures a system, a startup requests public data or an automated recommendation affects a person. Who performs the assessment? What information is published? Which body can stop the use? How are Rodrigues and other island communities represented within national evidence? These questions translate FAIR from an acronym into governance.

Mauritius also illustrates the relationship between national and regional ambition. Local firms may seek African and international markets, while public institutions draw on external technical support. That can expand opportunity, but it makes exit rights and capability transfer important. An experiment should leave the state and local organisations better able to procure, audit and negotiate, rather than permanently reliant on the partners that helped establish the programme.

5.4 Hawai'i and administrative guidance

The State of Hawai'i has taken a more operational approach through statewide data and AI policies. Its 2025 guidance for generative-AI assistants applies across state agencies and addresses security, data protection, human oversight, intellectual property, disclosure and verification of outputs (State of

Hawai'i State Data Office, 2025). The guidance distinguishes licensed enterprise tools from free public services and warns staff not to place non-public data into insecure systems.

This is a modest but important form of governance. Many AI risks arise before a formal high-profile project exists, when staff begin using widely available assistants in routine work. Clear rules can reduce accidental disclosure and establish that outputs remain drafts requiring human review. The limits are equally instructive. Internal usage guidance does not replace impact assessment for systems that allocate services, profile people or interact directly with the public. Nor does disclosure alone provide a remedy when an automated process causes harm.

Taken together, the cases show four capacities that island experiments require: an explicit public purpose, a regulator able to impose conditions, institutions capable of independent evaluation, and democratic processes that reach beyond the organisations promoting adoption. Jurisdictions possess these capacities in different combinations. The point of comparison is not to rank them. It is to identify what must be strengthened before experimentation expands.

6. FROM PILOT TO PUBLIC DECISION

A well-governed pilot should be organised as a public decision process, not as a technical preview. Before testing begins, the responsible authority should publish a short experiment notice. It should state the problem, the proposed system, the legal basis, participating organisations, affected groups, data involved, duration, success criteria, foreseeable risks and contact for challenge. Commercial details can remain protected where the law permits, but the existence and public purpose of the experiment should not be confidential.

The experiment should have a defined boundary. This includes who may participate, which data may be used, what decisions the system may influence and what it is prohibited from doing. A tool tested for drafting correspondence should not quietly be used to assess eligibility. Data collected for environmental monitoring should not be repurposed for enforcement without a new decision. Scope change is often where a limited pilot becomes an unexamined infrastructure.

Independent review must be proportionate to risk. Low-risk staff tools may require an internal assessment and security review. Systems affecting health, education, employment, policing, migration, benefits or access to essential services require external expertise and representation from affected groups. Independence is weakened if the evaluator depends on the vendor for future work or if results can be withheld by the project sponsor.

Evidence should be agreed in advance. Accuracy is only one measure. Evaluation may need to include false positives and negatives, disparities among islands or groups, staff workload, service delays, user comprehension, accessibility, cost, energy and connectivity requirements, security incidents and the ability to continue without the system. Qualitative evidence is not an optional supplement when impacts concern dignity, trust or cultural meaning.

People need routes to challenge outcomes during a pilot. A human review mechanism must have authority to change the decision, not merely explain it. Complaints should be recorded as evidence about the system. Where a model affects a community collectively, the process should allow collective representation rather than requiring every individual to pursue a separate case.

Exit should be designed before entry. Contracts must specify data return or deletion, model and document access, transition support, continuing service, intellectual-property rights and responsibility for harm. Staff should retain the ability to operate the service manually or through an alternative method. A pilot that cannot be stopped safely is already a deployment.

The final decision should be distinct from the experiment. Continuation should require a published evaluation, legal and financial approval, and a response to public concerns. Success should not be inferred from the absence of complaints or from technical completion. The decision may be to adopt,

redesign, extend testing, restrict the use or stop. Stopping a pilot because the evidence is poor is a successful act of governance.

7. THE ISLAND AI EXPERIMENT TEST

The Island AI Experiment Test is intended for governments, universities, public bodies, community organisations and funders considering a pilot, testbed or sandbox. It should be completed before data are transferred or people are enrolled, and repeated before any move into routine use. The test does not produce a single score. A serious failure in rights, legality, safety or exit cannot be compensated by economic promise elsewhere.

Dimension	Required question
1. Island purpose	Is the experiment responding to a problem defined by island institutions or communities, rather than seeking a convenient population or jurisdiction for a preselected technology?
2. Authority and law	Which public body is accountable, what is the legal basis, and which rules remain fully applicable during the experiment?
3. Participation and consent	Who is affected, how can they shape or refuse the experiment, and what additional collective authority is required for community or Indigenous data?
4. Proportionality	Is real-world testing necessary, are less intrusive alternatives available, and are excluded uses stated clearly?
5. Evidence and independence	What outcomes, harms and distributional effects will be measured, and who can evaluate them independently of the promoter and supplier?
6. Transparency and remedy	What will be published, how will people know AI is involved, and who can correct an outcome or obtain redress?
7. Capability and value	What skills, infrastructure, rights or economic value will remain in the island jurisdiction after the experiment?
8. Continuity and exit	Can the service continue if the system, supplier, power or network fails, and can the jurisdiction leave without losing essential data or capability?
9. Decision after testing	Who decides whether to adopt, redesign or stop, and is that decision separated from the team responsible for promoting the pilot?

The test changes the presumption of an island sandbox. The jurisdiction is not an empty container within which innovation is permitted. It is a political community whose law, knowledge, relationships and public institutions are part of the experiment and must remain capable of governing it.

8. AN ISLAND AI CHARTER

The Island AI Charter brings the series' arguments into a common framework. It is addressed to island governments and public bodies, but its commitments are also relevant to universities, funders, regional organisations and firms seeking to work in island societies. The Charter is not a certification mark. It is a basis for policy, procurement, assessment and public challenge.

8.1 Island authority

Island communities and their legitimate institutions should have the authority to determine whether, where and for what purposes AI is used. External funding or technical ownership must not displace public law or local decision-making. Where constitutional power sits outside the island, responsible authorities should create formal roles through which island institutions can approve, condition or reject uses that materially affect them.

8.2 Democratic consent and participation

People affected by consequential AI should be involved early enough to influence the problem definition, alternatives and decision to proceed. Participation must include groups likely to bear error or surveillance, not only organisations already connected to government and industry. Consultation does not remove the need for lawful authority, and individual consent is not a substitute for democratic justification in essential public services.

8.3 Cultural and linguistic integrity

AI systems should respect island languages, cultural practices and ways of knowing. Language and heritage data should not be treated as unowned raw material. Communities should be able to define appropriate uses, attribution, access and restrictions. A system that increases digital visibility while misrepresenting a culture or weakening the institutions that sustain it has not produced cultural benefit.

8.4 Data rights and collective interests

Island data governance should protect individuals and recognise collective interests. Public, environmental, community and Indigenous data require differentiated rules. The CARE principles provide a strong basis where Indigenous data are involved: collective benefit, authority to control, responsibility and ethics (Carroll et al., 2020). Data access should be limited to an agreed purpose, with clear retention, correction, deletion and onward-sharing provisions.

8.5 Transparency contestability and remedy

People should know when AI materially shapes a public decision or service. Responsible bodies should publish the system's purpose, supplier, data categories, limitations and assessment, subject to narrowly defined security protections. An affected person must be able to obtain human review and an effective remedy. Transparency is useful only when someone has the authority and information to act on it.

8.6 Equitable benefit and value retention

The benefits and burdens of AI should be examined across islands, communities and social groups. Public pilots should identify the value contributed through data, staff time, institutional access and reputation. Contracts and partnerships should provide fair return through service improvement, skills, infrastructure, research access, ownership or revenue where appropriate. Outer islands should not become data sources for benefits concentrated in a capital.

8.7 Resilience and governed interdependence

Island institutions should be able to continue essential functions during power, network, model or supplier failure. Interdependence is unavoidable and can be beneficial, but it must be governed through redundancy, portability, regional cooperation and credible alternatives. AI should strengthen the ability to act under disruption rather than remove local knowledge and manual capability.

8.8 Environmental responsibility

Decisions about AI should account for energy, water, hardware, connectivity and waste across the full service. Environmental cost should not be displaced invisibly to another island or mainland facility. High-compute systems require a demonstrated public value that simpler methods cannot provide. Procurement should consider efficient models, renewable supply, equipment life and the environmental consequences of new infrastructure.

8.9 Reversibility and the right to refuse

Island institutions and communities should be able to stop, limit or reverse an AI use without losing essential services, data or legal rights. Refusal is meaningful only when dependency has not already

made exit impossible. Pilots require sunset dates; contracts require termination and transition provisions; and sensitive cultural or community data may require a continuing right to withdraw permission for specified uses.

8.10 Inter-island cooperation and shared learning

Islands should cooperate to evaluate systems, pool specialist capacity, negotiate with suppliers and publish lessons. Regional action should extend island authority rather than create another distant centre. Shared standards can reduce duplicated effort, but they should allow languages, laws and community priorities to remain visible. Failures as well as successes should enter the common evidence base.

The Charter draws on international frameworks but addresses a distinct geography of power. UNESCO's Recommendation on the Ethics of Artificial Intelligence emphasises human rights, cultural diversity, environmental protection and inclusive governance (UNESCO, 2021). The Council of Europe Framework Convention connects AI explicitly to human rights, democracy and the rule of law (Council of Europe, 2024). The EU AI Act establishes enforceable risk-based duties and conditions for sandboxes (European Union, 2024). The Charter does not replace these instruments. It translates their concerns into questions of island authority, archipelagic inequality, cultural sovereignty and external dependency.

9. PUTTING THE CHARTER INTO PRACTICE

A Charter has value only if it changes decisions. The first step is to embed it in existing institutions rather than create a separate ethics process with no authority. Cabinet guidance, procurement rules, research ethics procedures, data-sharing agreements and sector regulation should refer to the commitments. High-risk proposals should require a published assessment that explains how each commitment is met or why the system will not proceed.

Procurement is a particularly important lever. Many island governments do not build AI systems; they buy access to them. Tender documents should require disclosure of subcontractors, hosting, training-data provenance where relevant, performance by location and group, security testing, incident reporting, audit access, model changes, energy requirements and exit support. Standard clauses developed regionally could improve bargaining power and reduce the cost of specialist legal work.

An island AI register should record consequential public uses and formal experiments. Each entry should identify the responsible body, purpose, supplier, status, data categories, assessment, known limitations, complaints route, evaluation date and decision to continue or stop. The register need not expose security-sensitive detail. Its purpose is to allow elected representatives, journalists, researchers and residents to see where public authority is being delegated or assisted by algorithmic systems.

Independent capability is essential. A small jurisdiction may not need a large standalone AI regulator, but it needs access to technical, legal, social and cultural expertise that is not controlled by suppliers. This may be organised through an auditor, ombuds institution, data protection authority, university network or regional assurance body. Community expertise should be remunerated. Participation that relies repeatedly on unpaid labour transfers the cost of governance to those already bearing risk.

Regional organisations can support shared model evaluation, incident reporting, training and contract negotiation. They can maintain specialist teams for languages, cybersecurity, procurement and impact assessment. They can also represent island priorities in international standards bodies where individual jurisdictions have limited capacity. The design principle is subsidiarity: decisions should remain as close as possible to affected communities, while functions requiring scale are shared.

Evaluation must continue after adoption. Models change, populations change and vendors update services. A system that passed a pilot may drift or be used for purposes not anticipated at approval. Review dates, performance monitoring and public incident reporting should therefore be contractual and regulatory requirements. A material change in model, data or purpose should trigger renewed assessment.

The Charter should also shape research funding. Funders can require evidence of island partnership, community authority, local benefit, data governance and publication of negative results. Universities should distinguish research participation from permission for commercial reuse. Ethics approval should examine collective and place-based harms alongside individual consent. Where a project develops language or cultural technology, sustained support for the people and institutions maintaining the source material should be part of the research design.

Finally, the Charter requires political use. Elected bodies should debate major public deployments and receive evidence in a form that does not depend on vendor expertise. Civil society and media need access to registers, assessments and contracts. Schools and universities can build public understanding without presenting adoption as inevitable. Democratic capacity is not an obstacle added after innovation. It is part of the infrastructure through which an experiment becomes legitimate knowledge.

10. CONCLUSION

Islands can be valuable sites of AI experimentation. Their institutions may coordinate quickly, their challenges may demand place-specific tools and their experience can reveal weaknesses hidden by large national averages. An island can also become a convenient laboratory: bounded enough to market, small enough to influence and distant enough for costs to be overlooked. The difference lies in authority.

The Island AI Experiment Test asks practical questions before a pilot begins. It identifies the public purpose, legal responsibility, affected groups, evidence, remedies, local value and exit route. These questions make it harder for a demonstration to become infrastructure without a public decision. They also make successful experiments more credible because learning is tied to an explicit baseline and an accountable institution.

The Island AI Charter provides the longer-term framework. Its ten commitments bring together the central findings of the AI-slands series: island authority; democratic consent; cultural and linguistic integrity; individual and collective data rights; transparency and remedy; equitable benefit; resilience; environmental responsibility; reversibility; and inter-island cooperation. None can be delivered by a statement of principles alone. They require law, contracts, capable public institutions, funded participation and repeated review.

The series began from the proposition that islands should not be treated as peripheral recipients of a technological future designed elsewhere. The same proposition now applies to experimentation. Islands have good reasons to test new systems and to contribute to global AI governance. They should do so as rule-makers, knowledge holders and political communities. The measure of a successful island testbed is not how easily technology entered. It is whether the island gained the authority and capability to decide what should happen next.

REFERENCES

Baldacchino, G. (2007) 'Islands as novelty sites', Geographical Review, 97(2), pp. 165-174. https://doi.org/10.1111/j.1931-0846.2007.tb00396.x

Carroll, S.R., Garba, I., Figueroa-Rodriguez, O.L., Holbrook, J., Lovett, R., Materechera, S., Parsons, M., Raseroka, K., Rodriguez-Lonebear, D., Rowe, R., Sara, R., Walker, J.D., Anderson, J. and Hudson, M. (2020) 'The CARE Principles for Indigenous Data Governance', Data Science Journal, 19, 43, pp. 1-12. https://doi.org/10.5334/dsj-2020-043

Carroll, S.R., Herczog, E., Hudson, M., Russell, K. and Stall, S. (2021) 'Operationalizing the CARE and FAIR Principles for Indigenous data futures', Scientific Data, 8, 108. https://doi.org/10.1038/s41597-021-00892-0

Council of Europe (2024) Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, CETS No. 225. Strasbourg: Council of Europe. https://rm.coe.int/1680afae3c

European Union (2024) Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Official Journal of the European Union, 12 July. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

Government of Malta (2019) Malta: The Ultimate AI Launchpad: A Strategy and Vision for Artificial Intelligence in Malta 2030. https://oecd.ai/en/dashboards/policy-initiatives/a-strategy-and-vision-for-artificial-intelligence-in-malta-2030

Jennings, A. (2026a) AI-slands: Towards an Island-Centred Approach to Artificial Intelligence. AI-slands Working Paper No. 1. Institute for Northern Studies, University of the Highlands and Islands. https://ai.insularworlds.co.uk/working-paper-1

Jennings, A. (2026b) AI and Island Cultures: Preservation, Extraction and the Right to Refuse. AI-slands Working Paper No. 2. Institute for Northern Studies, University of the Highlands and Islands. https://ai.insularworlds.co.uk/working-paper-2

Jennings, A. (2026c) Small Islands, Big Decisions: AI Governance, Dependency and the Right to Shape the Digital Future. AI-slands Working Paper No. 3. Institute for Northern Studies, University of the Highlands and Islands. https://ai.insularworlds.co.uk/working-paper-3

Jennings, A. (2026d) Bridging the Water: AI, Distance and the Future of Island Services. AI-slands Working Paper No. 4. Institute for Northern Studies, University of the Highlands and Islands. https://ai.insularworlds.co.uk/working-paper-4

Jennings, A. (2026e) Can AI Help Islands Keep Their People? Work, Demography and the Conditions for Durable Island Livelihoods. AI-slands Working Paper No. 5. Institute for Northern Studies, University of the Highlands and Islands. https://ai.insularworlds.co.uk/working-paper-5

Jennings, A. (2026f) Intelligent Islands or Dependent Islands? AI Infrastructure, Energy and the Case for Governed Digital Interdependence. AI-slands Working Paper No. 6. Institute for Northern Studies, University of the Highlands and Islands. https://ai.insularworlds.co.uk/working-paper-6

Jennings, A. (2026g) Algorithmic Archipelagos: Uneven Development Within Island Regions. AI-slands Working Paper No. 7. Institute for Northern Studies, University of the Highlands and Islands. https://ai.insularworlds.co.uk/working-paper-7

Jennings, A. (2026h) AI for Island Resilience: Climate, Disasters, Food and the Sea. AI-slands Working Paper No. 8. Institute for Northern Studies, University of the Highlands and Islands. https://ai.insularworlds.co.uk/working-paper-8

Lewis, J.E., Abdilla, A., Arista, N., Baker, K., Benesiinaabandan, S., Brown, M., Cheung, M., Coleman, B., Cordes, A., Davison, J., Duncan, S., Garzon, S., Harrell, D.F., Jones, P.-L., Kealiikanakaoleohaililani, K., Kelleher, M., Kite, S., Lagon, O., Leigh, J., Leong, J., Lewis, O., Mahelona, K., Moses, C., Nahuewai, I., Noe, K., Olson, D., Parker Jones, O., Running Wolf, C., Silva, M., Fragnito, S. and Whaanga, H. (2020) Indigenous Protocol and Artificial Intelligence Position Paper. Honolulu: Initiative for Indigenous Futures and CIFAR. https://doi.org/10.11573/spectrum.library.concordia.ca.00986506

Mittelstadt, B. (2019) 'Principles alone cannot guarantee ethical AI', Nature Machine Intelligence, 1, pp. 501-507. https://doi.org/10.1038/s42256-019-0114-4

OECD (2023) Regulatory Sandboxes in Artificial Intelligence. OECD Digital Economy Papers No. 356. Paris: OECD Publishing. https://doi.org/10.1787/8f80a0e6-en

OECD (2025) Regulatory Sandbox Toolkit. Paris: OECD Publishing. https://doi.org/10.1787/de36fa62-en

Sloane, M., Moss, E., Awomolo, O. and Forlano, L. (2022) 'Participation is not a design fix for machine learning', Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization. https://doi.org/10.1145/3551624.3555285

Smart Nation Singapore (2023) Singapore National AI Strategy 2.0: AI for the Public Good, for Singapore and the World. Singapore: Government of Singapore. https://go.gov.sg/nais2023

Smart Nation Singapore (2026) Update to National AI Strategy. Singapore: Government of Singapore. https://www.smartnation.gov.sg/initiatives/national-ai-strategy/

State of Hawai'i State Data Office (2025) GenAI Assistant Technologies Usage Guidelines. Document CDO-007. Honolulu: State of Hawai'i. https://data.hawaii.gov/wp-content/uploads/2025/09/CDO007-GenAI-Assistant-Technologies-Usage-Guidelines-NEW.pdf

UNDP (2026) 'Mauritius launches National AI Strategy and FAIR Guidelines supported by UNDP', 10 April. https://www.undp.org/mauritius-seychelles/press-releases/mauritius-launches-national-ai-strategy-and-fair-guidelines-supported-undp

UNESCO (2021) Recommendation on the Ethics of Artificial Intelligence. Paris: UNESCO. https://unesdoc.unesco.org/ark:/48223/pf0000381137

UNESCO (2024) Needs Assessment for Artificial Intelligence, Digital Transformation and Open Data for Small Island Developing States. Paris: UNESCO. https://doi.org/10.58338/DEMF2854

UNESCO and Innovation for Policy Foundation (2022) Multistakeholder AI Development: 10 Building Blocks for Inclusive Policy Design. Paris and Ebene: UNESCO and i4Policy. https://unesdoc.unesco.org/ark:/48223/pf0000382570

SOURCE NOTE

Sources and links were checked on 21 September 2026. Peer-reviewed research, legislation and official institutional materials were preferred. Government strategies and programme pages are used as evidence of stated policy and institutional design, not as proof of successful implementation. Because several examples are recent, operational outcomes should be reassessed as evaluations and audit evidence become available.