

Auditing marine survey evidence for submerged prehistoric landscapes on the Spanish Gulf of Cádiz shelf

Arman Askari Nazari¹

armarmask@alum.us.es

Universidad de Sevilla

Pablo Fraile-Jurado¹

Universidad de Sevilla

Juan Carlos Mejías-García²

Universidad de Sevilla

Research Article

Keywords: submerged prehistoric landscapes, marine geophysics, legacy-data reuse, observation support, evidence quality, Gulf of Cádiz

Posted Date: September 29th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11028557/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Marine archives are often treated as if their existence proves that the seabed has been archaeologically observed. We tested that assumption for the Spanish Gulf of Cádiz shelf by auditing 21 source families identifiable by 21 August 2026 against a predefined target: preserved Late Pleistocene to middle Holocene terrestrial or nearshore deposits capable of retaining occupation-related evidence. Two worked applications anchor the audit: parametric sub-bottom profiling at El Saladillo documents a buried Last Glacial Maximum palaeochannel, whereas anomaly-to-diver linkage in Cádiz Bay verifies exposed objects without testing buried prehistoric contexts. Each family was evaluated for identifiable metadata, testable observation geometry, achieved configuration, target fit, controlled result, reuse and support for negative inference. Seventeen families passed metadata, eight retained testable geometry under the strict cumulative rule and two also retained interpretable configuration; none fully matched the primary target. All 21 failed the final requirement for an adequate negative inference. This does not show that archaeology is absent or that the shelf has never been investigated; it shows that no recovered source presently supports a regional areal ruling-out of the target. Twenty method-specific observations and 13 capability envelopes instead provided bounded local evidence. The main contribution is a target-specific audit that separates catalogue availability, observation opportunity and controlled archaeological outcome. The Gulf archive supports local palaeolandscape interpretation and exposed-object verification, but not defensible regional exclusion of preserved prehistoric contexts.

1. Introduction

Repeated Late Quaternary sea-level lowstands exposed extensive areas of the continental shelves as terrestrial and nearshore landscapes. These areas were not peripheral to human settlement: they contained coastlines, river valleys, wetlands, resource zones and routes that could have structured mobility and occupation. Their subsequent inundation removed substantial parts of the archaeological record from conventional land-based research while preserving other parts beneath marine sediments. Postglacial transgression, however, did not affect all settings uniformly. Erosion, reworking, burial and later seabed disturbance produced a spatially uneven archive in which landforms, palaeoenvironmental deposits and archaeological materials may survive at different scales and with different degrees of integrity (Bailey and Flemming 2008; Flemming et al. 2014; Bailey et al. 2020). The scientific problem is therefore not simply to reconstruct where former land existed, but to determine where interpretable evidence may have survived and whether available investigations were capable of observing it.

Research has made major advances at the regional, landscape and site scales. Regional assessments combine sea-level histories, bathymetry, geomorphology, terrestrial archaeology and cultural evidence to identify broad areas for further investigation (McCarthy et al. 2022). Industry three-dimensional seismic data have been reused to resolve drowned barriers, lagoons, shorelines and drainage systems at scales that substantially improve palaeolandscape interpretation (O'Leary et al. 2020), while petroleum-exploration archives in the Gulf of Mexico show both the archaeological value of commercial data and the institutional work required to identify and reuse them (Jones 2023). Less conventional relationships

between instruments and archaeological targets have also been demonstrated: in the Netherlands, magnetometer patterns were linked to buried channel and levee systems through sub-bottom profiling, coring and comparison with onshore stratigraphy (van den Brenk et al. 2025). Development-led investigations can likewise generate valuable observations, as shown by the integration of geophysics, cores and dredging monitoring in southern California (Hildebrand and York 2022). Together, these studies establish that existing scientific, hydrographic and commercial data may hold considerable archaeological value. They also show that this value is conditional on the target, environment, acquisition design, processing and ground truth.

That conditionality creates an unresolved evidence-denominator problem. A catalogue entry, mapped raster, transit line, survey swath, seismic profile, core and diver inspection do not provide equivalent opportunities to detect a specified target. Even surveys described by the same method name can differ materially in frequency, bandwidth, line spacing, positioning, penetration, achieved resolution, seabed conditions, processing and quality control. Multibeam and side-scan data may document exposed morphology or large seabed objects while remaining unable to exclude buried palaeosurfaces. Sub-bottom data may resolve stratigraphic architecture yet provide no direct evidence of occupation or chronology. Cores sample a minute physical volume and can establish depositional setting without providing a realistic test for discrete artefacts between stations. Visual inspection can verify exposed targets but normally cannot exclude buried evidence. Archaeological survey quality assurance therefore requires explicit objectives, calibrated methods, documented coverage and reviewable error or detection conditions rather than nominal method labels (Banning et al. 2017). Current marine guidance likewise treats survey design as aim- and environment-specific (Historic England 2025). Experimental work reinforces the point: a CHIRP-based claim for direct lithic detection in Apalachee Bay could not be independently replicated even though the survey produced useful palaeolandscape information (Fitch and Cook Hale 2025).

Legacy-data reuse adds a second difficulty because the available evidence is not a representative sample of the seabed. Surveys are commissioned for navigation, engineering, aggregate extraction, dredging, route selection, environmental assessment, unexploded-ordnance investigation, geological research or heritage management. Their locations, configurations and archived products reflect those purposes. Access is also fragmented among scientific repositories, port authorities, heritage bodies, contractors and project archives. An English national audit identified 67 projects and 204 geophysical or geotechnical survey events, while also documenting uneven reuse, variable technical reporting and dependence on particular development and research programmes (Wessex Archaeology 2013). The principle of collecting data once and using them many times is therefore valuable but not automatic: reuse requires recoverable geometry, configuration metadata, processing provenance, target-relevant quality review and permission to examine or redistribute the necessary products (Historic England 2025). Apparent gaps in archaeological knowledge may consequently reflect where investigations were commissioned, what was retained and what can be accessed, rather than the distribution of preserved archaeological contexts.

Inference must also preserve the meaning of survey outcomes. A detected modern object, a palaeochannel interpreted from acoustic profiles, a target not relocated by divers, an inaccessible dataset and a target not detected under adequate conditions are different evidence states. Collapsing them into a binary detection/absence variable creates false negatives and conceals the conditions under which non-detection could be informative. The same principle applies to contextual data used to explain survey placement or modern seabed conditions. A supplied measurement, an explicit zero and an omitted record are not interchangeable. Observation support, measurement support and the archaeological result must therefore remain separate throughout the analytical chain. This separation is necessary both for scientific validity and for proportionate heritage decisions, because uncertainty about observability should lead to targeted information requirements rather than an unsupported conclusion that archaeology is absent.

The Spanish Gulf of Cádiz shelf provides a strong setting in which to address this problem. Sea-level reconstruction indicates that large parts of the Iberian shelf were repeatedly exposed during the Last Glacial Cycle, creating changing coastal and fluvial landscapes with potential relevance to prehistoric settlement (Fraile-Jurado and Mejías-García 2022; Fraile-Jurado et al. 2025; Mejías-García and Fraile-Jurado 2026). The region also contains dense and heterogeneous marine activity, scientific cruises, hydrographic mapping, port and dredging investigations, geotechnical sampling and heritage interventions. Yet the existence of these holdings does not establish how much of the shelf has been observed under conditions suitable for preserved palaeolandscape contexts or occupation-capable deposits. This evidential problem matters within a Spanish maritime-planning system in which cultural heritage is recognised but remains unevenly integrated across plans, evidence bases and responsible institutions (Canoura 2026). Regional work has reconstructed relative archaeological attention from exposure, hydrographic and lithological proxies and has distinguished sensitivity from institutional evidence conditions. The present study examines the missing evidence-production stage: whether existing investigations supplied a documented opportunity to observe the target and what conclusions their outcomes legitimately permit.

This article develops and applies a target-specific audit for deciding when a marine archive constitutes archaeological observation rather than merely nominal data availability. It asks one central question: to what extent do recovered investigations on the Spanish Gulf of Cádiz shelf document a realistic opportunity to observe preserved Late Pleistocene to middle Holocene contexts capable of retaining occupation-related evidence? A provenance-controlled inventory is evaluated against explicit geometry, achieved configuration, target-fit, outcome and reuse requirements, and admissible results are allowed to revise only the evidence state supported by the observation. The contribution is methodological: a reproducible way to construct and audit the observation denominator for heterogeneous legacy records. The Gulf of Cádiz case demonstrates the framework and identifies the smallest archive-recovery or new-survey actions likely to change the present evidence state. It is not a new archaeological-potential surface, a complete census of every survey undertaken or a model of archaeological probability.

2. Materials and methods

2.1 Study design and spatial domain

The study used a regional evidence audit combined with a bounded record-linkage pilot. The regional component covered the Spanish Gulf of Cádiz shelf represented by an established 100 m analytical grid. This domain provides the strongest overlap among the reconstructed exposure framework, the archaeological-potential index (API), the institutional evidence architecture and the survey records recovered for this study. Portuguese and Moroccan campaigns were retained in the discovery inventory when they supplied regional context, but they were not pooled into the Spanish analysis because access rules, heritage designations and reporting systems are not equivalent across jurisdictions.

Spatial processing used ETRS89 / UTM zone 30N (EPSG:25830). The authoritative mask contains 858,501 valid cells on an approximately 100 m grid and represents 8,582.846 km². The API and its exposure, hydrography and lithology components were treated as contextual hypotheses, not as observations of archaeology. Cádiz Bay formed a separate pilot because project-led geophysical anomalies could be linked to later diver inspections through identifiers and coordinates. The pilot was used to test record linkage, outcome semantics and inference limits; it was not treated as a representative sample of anomaly or archaeological prevalence on the regional shelf.

2.2 Archaeological target and nested evidence levels

The primary target was defined before any survey was classified: a spatially bounded, preserved Late Pleistocene to middle Holocene terrestrial or nearshore depositional context with a credible capacity to retain occupation-related archaeological evidence and whose confirmation, exclusion, condition or geometry could change a planning or project decision. The broad temporal envelope was approximately 120–6 ka BP, but records were evaluated against period- and environment-specific expressions rather than treated as one undifferentiated behavioural interval. A candidate target required a plausible terrestrial, estuarine, lagoonal, riverine or nearshore setting; an interpretable preservation pathway; a feasible chronological test; sufficient geometry to relate it to an observation; and a result capable of changing investigation, routing, mitigation, monitoring or preservation requirements.

Evidence was organised into five nested but non-equivalent levels (Table 1). T0 records a documented opportunity to observe; T1 and T2 distinguish preserved palaeolandscape context from an occupation-capable context; T3 records direct occupation-related evidence; and T4 adds the characterisation required for management. The principal audit concerned T1 and T2. Historic wrecks, isolated modern objects, undated anomalies, deep geological structures without a demonstrated habitable-shelf relationship and high-API cells were not accepted as substitutes for the primary target.

Table 1
Target and evidence levels used in the audit

Level	Evidential meaning	Minimum interpretation
T0	Documented observation opportunity	Recoverable method, geometry, effort and usable data; not archaeological evidence by itself
T1	Preserved palaeolandscape context	Buried surface, channel fill, peat, soil, lagoonal body or another preservation trap
T2	Occupation-capable context	T1 context with compatible chronology, environment, accessibility and preservation
T3	Direct occupation-related evidence	Material, features, residues or patterned deposits
T4	Characterised heritage occurrence	T3 occurrence assessed for character, integrity, extent, importance and vulnerability

The levels are nested but not interchangeable. The principal regional audit concerns T1 and T2; progression is never inferred from level labels alone.

2.3 Evidence architecture and inventory construction

Evidence discovery was iterative rather than prospectively registered. It combined the local project library, scientific repositories, hydrographic and geological catalogues, port and heritage archives, cruise reports, published literature, recovered navigation and geophysical files, and direct responses from custodians. Consequently, the original document-level hit count and every verbatim query cannot be reconstructed, and no systematic-review document denominator is claimed. Reproducible accounting begins with the 21 candidate source families formed and frozen on 21 August 2026. Online Resource 3 records this retrospective boundary and supplies the protocol for future updates.

Discovery used three routes: scientific, hydrographic or geological catalogues and repositories; heritage, port or project archives; and an inherited legacy compilation. Searches combined a geographic block (Gulf or Golfo de Cádiz, Cádiz Bay, adjacent shelf and named ports), a method block (multibeam, side-scan, sub-bottom, seismic, magnetometry, cores, boreholes, dredging, diving or remotely operated vehicle) and a campaign or archive block (cruise, vessel, project, catalogue and custodian names). Citations, vessel identifiers, project references and custodian advice expanded the candidate set. These blocks are a prospective update protocol, not a claim that identical strings were used in every original search.

A source was retained when it documented a potentially relevant campaign, dataset, intervention or archive holding, even if primary files were restricted or target fit was poor. Duplicate catalogue and report entries were consolidated when they described the same acquisition family; materially different methods, dates, configurations or outcomes remained separate at the observation stage. No family was excluded after the 21-family denominator was formed. Fifteen custodian or access routes were recorded. Readiness gates were assigned at this source-family level. Five priority families were then

decomposed into 20 controlled observation units for outcome coding, while six partly overlapping families supplied 13 method-target-environment envelopes for capability calibration. These linked populations answer different questions and are not additive.

Two registers were maintained. Register A contained contextual and model evidence: the API and components, topobathymetry, emerged-time and palaeocoast products, coarse lithology and modelled drainage, heritage designations, aggregate-extraction and fishing context, and earlier model outputs. These data could define strata, explain survey placement or support planning, but could not supply the observation denominator. Register B contained observation evidence and recorded the project and commissioning purpose; acquisition date and custodian; native or reconstructed geometry; sensor and configuration; processing and positioning; environmental and operational limitations; observable target level; result and interpretation; confidence; access; licence; reuse conditions; and missingness. Only Register B was eligible to support claims about whether a target had been observed.

The source-family inventory used one row per coherent campaign, project or archive family. An observation unit was a method-, target-, configuration- and outcome-specific record derived from a family; an envelope was a separate calibration of what a method could or could not show under stated environmental and operational conditions. This hierarchy prevented, for example, a cruise track, multibeam mosaic, sub-bottom profile and core station from being treated as one interchangeable survey. Missing values were coded explicitly as not reported, unknown, restricted, not applicable, not assessed or derivable. No technical value was inferred from a method name or nominal specification unless it could be linked to the acquisition under review. Figure 1 summarises how the target ontology, the two registers, analytical grain, readiness gates and secondary opportunity classes relate. The complete source-family accounting and reusable discovery log are provided in Online Resource 3.

2.4 Sensor-target compatibility and archaeological quality control

Compatibility was assessed for each method-target configuration rather than for a sensor category in the abstract. Multibeam and side-scan sonar were considered capable of documenting T0 support, exposed seabed morphology and sufficiently large surface objects, but not of excluding buried palaeosurfaces or occupation horizons. Magnetometry could support the detection or discrimination of ferrous and other magnetically contrasting targets, but could not exclude non-magnetic prehistoric contexts. Shallow sub-bottom or parametric profiling could support buried T1 architecture and, where bandwidth, penetration, line geometry and ground truth were adequate, aspects of T2. Conventional or deep seismic reflection was retained for regional architecture only when the preserved shallow bandwidth and processing allowed the target interval to be evaluated. Cores and boreholes could establish local stratigraphy, formation, preservation, palaeoenvironment and chronology, but their minute physical support did not permit continuous absence between stations. Diver, remotely operated vehicle (ROV) and video observations were suitable for exposed T3 verification within a documented search geometry, not for excluding buried evidence.

The minimum configuration fields depended on the method. Acoustic records required instrument and source, frequency or bandwidth, pulse, range or swath, line and cross-line spacing, orientation, positioning, motion and tidal correction, water depth, achieved resolution, penetration, processing and masking conditions. Towed systems additionally required altitude and layback information. Cores and boreholes required diameter, recovered length, depth datum, recovery, lithology, sampling interval, analyses, chronological control and disturbance. Visual inspection required searched geometry, navigation, visibility, current, altitude or distance, search pattern, recording method and burial limitation. Guidance ranges, including approximately 50–100 m main-line spacing for regional shallow palaeolandscape work and 10–30 m for site-specific work, were used only as transparent comparators conditioned on target scale, burial, geology and achieved performance; they were not universal pass/fail thresholds (Historic England 2025).

Archaeological quality control was evaluated separately from hydrographic or contractor quality assurance, following the principle that survey reliability depends on objective-specific design, coverage, calibration and reviewable error conditions (Banning et al. 2017). A configuration was not treated as archaeologically adequate unless the record supported review of navigation, gaps, interference, weather, instrument settings, seabed or geological masking, achieved penetration and target-scale resolution. Cross-method corroboration by independent acoustics, cores, samples, diving or archaeological observations was recorded but did not erase the limitations of the original observation.

2.5 Observation geometry and evidence-readiness gates

Native tracklines, swaths, footprints and inspected points were preserved wherever available. Swaths were derived only when acquisition parameters justified them. Systematic acquisition lines, targeted lines and transit tracks were kept separate; the bounding extent of a composite raster was not substituted for its native source-survey footprint. Point inspections remained points unless a searched radius or polygon was documented. Records with different sensors, dates, configurations or outcomes were not dissolved into a binary coverage layer, and repeat observations were retained because they could change confidence or reveal temporal change. We use “testable geometry” for a footprint that can be related to an observation; “strict” refers only to the cumulative rule that retains full passes at every preceding gate.

Every source family was then evaluated through seven ordered readiness gates: identifiable metadata; testable geometry; interpretable observation-scale configuration; fit to the primary T1/T2 target; a controlled target-level result; workable access and reuse; and support for an adequate negative inference. Each gate received pass, partial or fail with a written rationale, evidence locator, decisive gap and recovery action. A strict cumulative series retained only records that fully passed the current and every preceding gate. A second pass-or-partial series retained incomplete but non-failed records to show potentially recoverable evidence. Comparison of the two series tested whether admitting documented incompleteness changed the terminal inference without converting partial evidence into a pass. The terminal result was also recomputed for every permutation of the seven gates. The series were

descriptive audit summaries, not weighted quality scores; no strong attribute could compensate for failure elsewhere. Source-level decisions are exposed in Online Resource 2.

Field-level validation required unique inventory, observation and envelope identifiers; complete links from the 21 inventory families to the readiness matrix; valid controlled-state values; and parent inventory keys for every observation project and envelope. Cumulative counts were rederived from the row-level matrix rather than entered manually, and spatial-model category totals were required to exhaust the authoritative mask. Gate assignment was completed by the study team as a documented interpretive audit, not as independent blinded coding; no inter-rater coefficient is therefore reported. Online Resource 2 exposes the source-level decisions and rationales, Online Resource 3 prespecifies independent recoding and adjudication, and Online Resource 4 contains a blinded 21-family coding workbook and the machine-readable audit data. Independent recoding must be completed before any agreement statistic or validated-gate claim is added.

For families that retained a controlled result and a workable reuse pathway under the pass-or-partial assessment, method-target detection envelopes were constructed. Each envelope documented observable target expression and scale, burial or exposure state, spatial support, acquisition and environmental conditions, validation, decisive failure conditions and prohibited inference. Envelopes were assigned one of four roles: local positive calibration, bounded project result, conditional pending primary evidence, or context only. Effective archaeological coverage was permitted only where geometry, configuration, environment, target fit, controlled result, archaeological review and reuse all passed.

2.6 Controlled observation outcomes

Every observation received exactly one primary result state: target detected; target not detected under adequate conditions; target not resolved or inconclusive; source target not relocated; data unavailable or unusable; or observation not designed for the target. A separate interpretation field classified the evidenced object as prehistoric, historic, modern, natural or unresolved and recorded whether it was a site or material occurrence, potential landform, stratigraphic horizon, palaeoenvironmental context or non-target feature. This separation prevented a physically detected modern object from being labelled an archaeological negative and prevented an inaccessible source from being absorbed into an undifferentiated inconclusive category.

An adequate non-detection required the target, footprint, configuration, environmental conditions, processing, repeatability and archaeological review to demonstrate a realistic opportunity for observation. “No reported archaeology,” a target not relocated, an unsuitable sensor, an unavailable profile, a transit line or a core without artefacts did not meet that definition. Confidence and the reason negative inference was or was not permitted were recorded independently from the primary result state.

For each admissible outcome, UPDATE was defined as a bounded transition in Register B rather than as a posterior-probability calculation. A supported detection could advance the evidenced target level or revise the interpreted object class within its documented geometry. An adequate non-detection could

reduce or redirect a subsequent information requirement only within the target, footprint, horizon and achieved conditions that had passed the applicable gates. Inconclusive, not-relocated, unavailable, unusable or not-designed-for-target outcomes did not reduce archaeological uncertainty; they instead updated the record of observation support, the decisive gap or the next recovery requirement. The revised outcome, confidence, permitted inference and decisive gap were then carried forward for institutional review. The outcome-to-update and review rules are summarised in Online Resource 1. No archaeological prior, survey likelihood, posterior probability, loss function or expected value of sample information was estimated.

2.7 Cádiz Bay anomaly-to-verification pilot

The Cádiz pilot linked ten underwater inspections undertaken in 2024 to source geophysical anomalies recorded during the 2019 approach-channel investigation. Linkage used anomaly identifiers where available, target labels, dimensions, geographic coordinates, projected coordinates, survey sector and the detailed observation narrative. World Geodetic System 1984 (WGS84) latitude and longitude were preserved as the preferred independent positional reference, and ETRS89 / UTM zone 30N coordinates were reconciled against them. Apparent transcription errors in projected coordinates were corrected only where the decimal-degree positions, summary table and sequence supplied an unambiguous check; each correction remained flagged in the register.

The diver reports were recoded using the controlled vocabulary rather than the source report's positive/negative wording. Directly observed objects were classified by interpreted type and age; targets that could not be relocated remained not-relocated because burial, movement, positioning error, visibility and unreported search effort could not be excluded. Search radius, bottom time and detailed visibility were retained as missing when not reported. The pilot was summarised as linked detected non-target and not-relocated observations. Sensitivity, specificity, false-negative rate and anomaly prevalence were not estimated because the complete 2019 target population, target-selection rule and comparable search-effort denominator had not been recovered.

2.8 Secondary planning and reporting-support stress tests

A secondary planning model tested whether the audit could guide evidence recovery without turning contextual data into archaeological evidence. On the authoritative 100 m grid, the inherited API was grouped into higher-attention classes 3–4 and lower-attention classes 0–2. Selected bathymetry from the Spanish Navy Hydrographic Institute (IHM) MBAR24 product was classified as no valid selected cell, regional 100 m context only, or fine 16 m context. Their exhaustive, non-compensatory cross-classification produced six opportunity classes. Existing observations and calibration points were overlays and did not lower a class or imply effective coverage.

Eight-neighbour components of at least 1 km² were retained in the three higher-attention classes as reproducible scoping units. Three illustrative design experiments used the largest components to compare reconnaissance and validation demands under absent, regional and fine bathymetric context. Their lines and stations are reproducibility geometries, not operational survey plans; every element

remains contingent on primary-profile interpretation, safety, heritage, permitting, substrate, vessel and instrument review.

Bottom-contact fishing supplied a separate reporting-support stress test, not evidence of archaeological presence, preservation or absence. The native evidential unit was the C-square. A 2021–2024 row with positive surface or subsurface swept-area ratio (SAR) was reported activity; an explicit zero required both populated values to equal zero; and a missing row remained missing rather than zero. Rasterised 100 m cells were used only for overlap. Historical products tested omission semantics, and the held Scenario 3 comparison candidate was rejected as a low-activity or causal control when current non-reporting could not be distinguished from missing support.

Evidence gaps were ranked by decision relevance, target observability, likely uncertainty reduction, feasibility and access. Archive recovery was preferred where navigation, sensor-on intervals, profiles, event tables, core logs or processing history could move a source through a decisive gate. New acquisition was considered only where recoverable records could not supply the primary target. These analyses are secondary applications of the audit, not additional models of archaeological probability.

2.9 Analysis, inference limits and reproducibility

The minimum analysis consisted of counts, areas, line kilometres, metadata-completeness rates, readiness-gate attrition and cross-tabulations by sensor, purpose, access, target level and controlled outcome. Where valid geometry existed, nominal coverage, documented T0 support, target-appropriate T1 support, T2 support, inconclusive support and repeat coverage were calculated separately. Observation-support availability, rather than archaeological presence, was the response for any survey-selection comparison. Configuration-specific detection performance would be estimated only if a sufficiently large set of comparable labelled targets and effort denominators became available.

Prohibited analyses included coding unsurveyed cells, unsuitable surveys, omitted records or “no reported archaeology” as absences; treating composite bathymetry as source-survey effort; using 100 m overlay cells as independent observations; or estimating a fishing effect from the Scenario 3 block and its held candidate. All outputs therefore preserve uncertainty and distinguish source family, observation unit, envelope, project, C-square, trackline, footprint, profile, core and inspection units.

The reproducibility structure separated protocol and controlled vocabularies, inventory and contact provenance, contextual layers and their native evidential grain, non-sensitive observation geometries, processing and validation scripts, figures and tables, and metadata-only descriptions of restricted sources. Coordinate reference systems, transformations, source hashes and derivation rules were recorded with each spatial product. Restricted or heritage-sensitive coordinates were not redistributed; the package retained a discovery and request route instead. Online Resources 1–3 provide the update crosswalk, source-readiness matrix, and discovery and validation protocol; Online Resource 4 provides the eligible machine-readable data, recoding workbook, scripts, validation output, manifests and

checksums. This design allows the audit to be updated as custodians release new metadata or files without silently changing earlier evidence states.

OpenAI ChatGPT was used for language editing, code-development assistance, consistency checking and reproducibility documentation. All analytical definitions, source classifications, calculations and interpretations were reviewed and accepted by the authors, who retain full responsibility for the work.

3. Results

3.1 Inventory composition and evidential granularity

Reproducible accounting began at 21 coherent source families (GOC-001–GOC-021) frozen on 21 August 2026; earlier document-level hits were not reconstructable because discovery had been iterative. Eleven families came through scientific, hydrographic or geological catalogues and repositories, nine through heritage, port or project archives, and one through an inherited legacy compilation. No family was excluded after this denominator was formed. The holdings spanned multibeam and single-beam bathymetry, side-scan sonar, shallow and multichannel seismic reflection, parametric sub-bottom profiling, magnetometry, cores, grabs, borehole locations, diving, dredging monitoring and derived terrain products. No single custodian exposed all fields required for target-specific evaluation. The full accounting and prospective update protocol are reported in Online Resource 3.

The source-family count substantially understated analytical heterogeneity (Fig. 2a). Five priority families were decomposed into 20 method-, target- and outcome-specific observations: two Spanish Navy Hydrographic Institute (IHM) MBAR24 records, four MOWER records, one El Saladillo record, three Camposoto records and ten Cádiz Bay anomaly-verification records. A partly overlapping group of six families supplied 13 method-target-environment envelopes. These were linked but different populations: observations tested result semantics, whereas envelopes calibrated capability or bounded contextual roles. The counts are therefore not additive. The legacy geographic information system (GIS) family (GOC-021), for example, contained 34 borehole points and restored layers but no drill logs, acquisition configuration, chronology or controlled target outcomes, so it remained inventory evidence rather than an observation.

Access status was heterogeneous rather than simply open or closed. Complete or partial files were locally available for several scientific repositories and derived products, while MOWER and related campaigns of the Marine Technology Unit of the Spanish National Research Council (UTM-CSIC) exposed public navigation and event metadata but retained raw primary holdings behind principal-investigator authorization. The Institute of Marine Sciences of the Spanish National Research Council (ICM-CSIC) confirmed that records must be identified and requested through its own collection platform rather than through a bespoke Gulf-wide extract. The Port of Algeciras directed unresolved technical requests to its formal electronic procedure. These responses established credible recovery pathways but did not improve observation readiness by themselves.

3.2 Evidence-readiness attrition

Nominal availability contracted sharply through the ordered readiness gates (Table 2; Fig. 2b). Seventeen of the 21 audited families fully passed metadata, eight retained testable observation geometry under the strict cumulative rule and two also retained interpretable achieved configuration. Neither fully matched the primary T1/T2 target, so the strict series reached zero at target fit. When partial evidence was retained, 21 families remained at metadata, 16 at geometry, 15 at configuration and nine at primary-target fit; six also retained a controlled result and workable reuse route, but none supported an adequate negative inference. Boundary treatment therefore changed the recoverable intermediate set but not the terminal conclusion. Recalculation across all 5,040 gate orders also ended at zero for both strict and pass-or-partial rules because every audited family failed the adequate-negative gate. Intermediate attrition remains sensitive to definitions and analytical grain; the terminal result does not. Source-level states, rationales and recovery actions are reported in Online Resource 2.

Non-cumulative states explained the attrition: geometry failed for five families, configuration for five, primary-target fit for seven and controlled result for 12. A reusable derived bathymetric surface could have a defined extent while lacking source-survey settings and a controlled archaeological result; a published local positive interpretation could remain useful even where exact navigation and raw profiles were unavailable. The non-compensatory design preserved both facts without turning either source into effective regional coverage.

Under the study's stringent definition, no audited source family qualified as effective regional T1/T2 coverage capable of supporting an areal negative inference, and no primary-target negative cells were produced. This finding does not mean that the shelf is archaeologically unobserved in every sense. It means that no retained family combined testable observation geometry, interpretable achieved configuration, primary-target fit, controlled result, archaeological review and workable reuse at the level required to generalise a non-detection beyond its documented support.

Table 2
Cumulative evidence retained through the ordered readiness gates (n = 21 audited source families)

Readiness gate	Strict pass	Pass or partial
Metadata identifiable	17	21
Geometry testable	8	16
Configuration interpretable	2	15
Primary-target fit	0	9
Controlled result	0	6
Reuse workable	0	6
Adequate negative inference	0	0

Strict pass retains only families that fully passed the current and all preceding gates. Pass or partial excludes failures but does not imply full readiness.

3.3 Controlled outcomes in the observation register

The 20 observation records comprised 11 detected outcomes, four not-relocated outcomes, three data-unavailable or unusable outcomes, one not-designed-for-target outcome and one inconclusive outcome. There were no adequate non-detections. The 11 detections did not represent 11 prehistoric discoveries. They included derived regional morphology, a documented ship-navigation opportunity, a locally interpreted palaeochannel, exposed historical or modern archaeological material at Camposoto (Rodríguez Mariscal and Martí Solano 2020) and six contemporary non-target objects in Cádiz Bay. Separating the observation state from the interpreted object class prevented these records from being pooled as equivalent archaeological presences.

Under the bounded update logic, the palaeochannel detection advanced local T1 evidence, while the remaining detections revised contextual, historical or modern non-target classifications within their supports. The four not-relocated outcomes and the unavailable, unusable, inconclusive or unsuitable observations did not lower primary-target uncertainty; they revised the recorded observation limitation or next evidence requirement. No observation produced a posterior archaeological probability, converted a planning cell to absence or independently determined a management action.

The MOWER family illustrated the difference between ship presence and achieved observation support. Its public one-minute navigation contained 47,326 valid positions, of which 1,349 positions and 221.733 km of track lay inside the analytical mask; recorded depths along the in-mask route ranged from 18.66 to 624.32 m (UTM-CSIC 2014). This was a recoverable T0 opportunity. Sensor-on intervals, line roles and swath or profile geometry were not public, however, so the associated multibeam, Parasound and core records remained data-unavailable or unusable for T1/T2 assessment. The 221.733 km route was therefore retained as a navigation proxy, not a target-capable track length.

The IHM MBAR24 family supplied a contrasting example. Six verified derivative GeoTIFFs covered 6,781.061 km², or 79.007% of the analytical mask, including 2,033.447 km² of the higher-attention API classes (Instituto Hidrográfico de la Marina 2024). The product supported regional seabed morphology at selected 16 m or 100 m grid scales, but source survey dates, systems, density, uncertainty and per-cell lineage were not recoverable from the mosaics. The uncovered 20.993% of mask cells could not be called unsurveyed, and the covered cells could not be called archaeologically observed. MBAR24 therefore passed as contextual morphology and planning support, not as a survey-effort denominator.

3.4 Target-specific detection envelopes

Thirteen method-target envelopes were retained from six evidence families. Four were local positive calibrations, one was a bounded project result, three were conditional pending primary evidence and five were context only. None permitted a primary negative inference. The four local positive calibrations

comprised side-scan sonar, magnetometry and diver verification of exposed or proud T3 material at the Campamento dry-dock project, together with parametric sub-bottom imaging of a buried T1 palaeochannel at El Saladillo. These calibrations demonstrated that particular target expressions could be observed under local conditions; they did not establish transferable sensitivity across the shelf.

El Saladillo supplied the strongest local T1 result. Approximately ten 200 m east–west reciprocal transects and several transects longer than 300 m imaged a previously undocumented Last Glacial Maximum palaeochannel and basal units using an Innomar SES-2000 system (Solana Rubio et al. 2026). Reported performance exceeded 10 m penetration in muddy seabed at approximately 10–20 m water depth, with best acquisition near 2–2.5 knots and degradation under greater wave motion, vessel vibration or speeds above approximately 4 knots. The observation establishes that buried late-Pleistocene channel architecture was locally detectable under these shallow muddy conditions. It does not demonstrate occupation, date a habitable surface or revise regional palaeogeography because exact navigation, line spacing, pulse settings, raw profiles, cores and chronological ties were unavailable. The next archaeological test is therefore whether channel margins, abandoned fills or adjacent surfaces preserve dated subaerial horizons or occupation-capable deposits, requiring profile recovery, targeted coring, chronological control and palaeoenvironmental or microarchaeological proxies.

The Campamento programme supplied a complementary T3 calibration. Side-scan sonar, magnetometry and targeted diver inspection detected and classified modern discarded objects inside a small preventive-archaeology project (Rojas Roca 2014). Missing target dimensions, raw imagery, magnetic amplitudes, tow geometry, search radii and achieved environmental conditions limited transferability. A one-month dredging-monitoring statement that no archaeological remains were found was retained as the sole negative-style bounded project result, but absent dredged-volume, depth, continuity and observation logs prevented it from becoming a primary-target non-detection.

The remaining envelopes clarified why method presence was insufficient. Processed European Marine Observation and Data Network (EMODnet), FRAME and INSIGHT bathymetric cells supplied local T1 morphology context but not native survey effort (EMODnet Bathymetry Consortium 2024). A Parasound configuration remained conditional because exact runlines, profiles and achieved penetration were inaccessible. The RC2912 archive supplied 15.86 km of in-mask transit centreline, of which 9.49 km contained corrected bathymetry at 80–123 m water depth, but neither its full swath nor systematic target search could be reconstructed (Marine Geoscience Data System 1989). Five International Ocean Discovery Program (IODP) Expedition 339 Gulf sites provided well-controlled formation and chronological context at approximately 559–993 m water depth, outside the primary shallow shelf domain. MBAR24 remained a regional contextual surface. These records increased geological knowledge without increasing effective T1/T2 archaeological coverage.

3.5 Cádiz Bay anomaly-to-verification linkage

All ten 2024 diver inspections were linked to their source anomaly identifiers or target records (Autoridad Portuaria de la Bahía de Cádiz 2024). Six targets were physically observed and classified as

contemporary non-target material: ropes, chains and a clamp; chain links; a concrete block with nets; a metal gangway fragment; a concrete mooring block; and an iron wreck at least 35 m long. Four targets—ONI-26, Anomaly-30, ONI-38 and ONI-113—were not relocated. The source suggestion that some might have been buried was retained as an explanation, not resolved as an observation.

Controlled recoding changed the meaning of several source labels. A report-level “negative” attached to an observed concrete or mooring object became a detected modern non-target outcome, while failure to relocate a target remained not-relocated rather than not-detected. Three projected-coordinate transcription problems were also identified and documented: a missing zero in the ONI-111 easting, an extra digit in the ONI-38 northing and an extra zero in the ONI-113 northing, with the ONI-113 easting discrepancy retained in the quality note. WGS84 coordinates and internally consistent table values supplied the independent checks.

The pilot did not yield a false-negative rate or detection probability. The complete 2019 anomaly population, the rule used to choose the ten 2024 targets, search radii, bottom times and comparable visibility or effort were not available. The six detections demonstrated that selected large exposed modern targets could be verified. The four not-relocated records demonstrated uncertainty in relocation under the reported conditions. Neither group tested preserved, buried occupation-related contexts.

3.6 Regional opportunity classes and contextual support

The authoritative mask contained 858,501 cells and 8,582.846 km². Higher-attention API classes occupied 2,684.403 km² and lower-attention classes 5,898.443 km². Cross-classification with selected MBAR24 context produced six exhaustive categories (Table 3; Fig. 3). Among higher-attention cells, 650.956 km² lacked a selected valid MBAR24 cell (H0/P1), 673.230 km² had regional 100 m context only (H1/P2) and 1,360.217 km² had fine 16 m context (H2/P3). The corresponding lower-attention categories occupied 1,150.830, 1,528.555 and 3,219.058 km². The three higher-attention classes together represented 31.276% of the mask.

The class distribution described planning context, not graded archaeological evidence. Fine MBAR24 coverage did not reduce an H2 cell's need for a target-capable buried-context survey, and lack of a selected MBAR24 cell did not prove lack of previous hydrographic work. Observation overlays were therefore displayed without lowering the base class. The GOC-003 screening parsed 22,237,862 rows from the processed composite digital terrain model (cDTM) but found only 6,709 source-labelled cell centres within the mask, representing an estimated 18.103 km² and 2,084 unique API pixels (0.242749% of the mask). Because those values were processed source-cell centres rather than native swaths, they were retained as sparse contextual support. The RC2912 transit line and the local Campamento and El Saladillo calibration references were treated in the same non-subtractive way.

Eight-neighbour analysis identified 71 H0/P1 components, 114 H1/P2 components and 24 H2/P3 components. Eleven, 12 and 12 components, respectively, were at least 1 km². The largest components occupied 193.581 km² in P1, 558.179 km² in P2 and 980.613 km² in P3. These three components formed

the reproducible envelopes for the provisional field scenarios; their size was not interpreted as archaeological site extent.

Table 3
Exhaustive survey-opportunity classes on the 8,582.846 km² analytical mask

Class	API attention	Selected MBAR24 context	Area (km ²)	Mask (%)
H0 / P1	Higher	No valid selected cell	650.956	7.584
H1 / P2	Higher	Regional 100 m only	673.230	7.844
H2 / P3	Higher	Fine 16 m	1,360.217	15.848
L0 / S1	Lower	No valid selected cell	1,150.830	13.408
L1 / S2	Lower	Regional 100 m only	1,528.555	17.809
L2 / S3	Lower	Fine 16 m	3,219.058	37.506

Classes combine inherited API attention with contextual bathymetric support. They are not archaeological probabilities, detections, non-detections or effective survey coverage.

3.7 Secondary planning and reporting-support stress tests

As a secondary application, three illustrative design experiments translated the largest higher-attention components into nested reconnaissance and validation geometries (Table 4; Fig. 4). Scenario 1 used the approximately 194 km² P1 component because selected bathymetric context was absent, with an approximately 12 km² validation footprint. Scenarios 2 and 3 used morphology-stratified reconnaissance footprints of approximately 80 and 60 km² and validation footprints of approximately 16 km². Geometry checks confirmed nesting and clipping, but these checks do not constitute achieved survey coverage; exact reproducibility metrics remain in the companion package.

The constraint screen confirmed that the experiments are not operational plans. Depth remained unresolved for Scenario 1; contextual median depths were approximately 33 m in Scenario 2 and 24 m in Scenario 3. Traffic, utilities, wrecks, unexploded ordnance, live restrictions, permits, substrate, target-horizon depth and instrument trials remain unresolved. All lines and stations therefore require redesign after primary-data review and project-specific safety, heritage and permitting assessment.

The fishing stress test likewise changed interpretation rather than archaeological evidence. The Scenario 3 block was almost entirely covered by two positive 2021–2024 C-square records, while the held equal-area candidate had no current row but was positive in archived products. Only four of 86,448 complete European records contained explicit zero surface and subsurface swept-area ratios, none locally. The candidate was therefore rejected as a low-activity, undisturbed or absence control and retained only as reporting-support uncertainty; C-squares, not rasterised cells, remained the evidential units.

Table 4
Illustrative design experiments for three higher-attention contexts

Scenario	Available context	Reconnaissance (km ²)	Validation (km ²)	Decisive pre-field uncertainty
1 (S-P1)	No selected bathymetric cell	≈ 194	≈ 12	Depth, morphology and operational constraints
2 (S-P2)	Regional 100 m bathymetry	≈ 80	≈ 16	Primary shallow profiles and target definition
3 (S-P3)	Fine 16 m bathymetry	≈ 60	≈ 16	Target horizon, substrate and live constraints

Rounded areas compare research-design demands. Exact reproducibility geometries are retained in the companion package; they are not navigation- or contractor-ready plans.

3.8 Empirical gap hierarchy

The audit resolved the regional evidence state into three operational tiers. First, contextual products with validated geometry—especially MBAR24, processed EMODnet cells, RC2912 and the inherited API components—could support stratification, morphological hypotheses and transparent selection but not target outcomes. Second, locally positive or bounded observations—El Saladillo, Campamento, Camposoto and the Cádiz pilot—could calibrate particular target expressions and result semantics but not regional detection performance. Third, candidate primary sources—especially MOWER multibeam, Parasound and core holdings; INSIGHT_LEG2 profiles; El Saladillo profiles and navigation; and port project archives—could potentially change readiness if their geometry, settings, primary data and controlled results were recovered.

The decisive gaps were therefore narrower than the nominal inventory suggested. Recovery of sensor-on navigation, runline roles and sample-event tables would determine whether MOWER intersected the shallow target domain under active acquisition. Access to exact INSIGHT_LEG2 and El Saladillo profiles, settings and interpretations would permit configuration- and target-specific T1 review. Source-lineage metadata for MBAR24 would improve hydrographic context without changing its archaeological role unless native acquisition and target review were also recovered. The complete 2019 Cádiz anomaly population and selection rule would strengthen the verification denominator but would still concern T3 anomaly verification rather than preserved T1/T2 contexts. Until those gates are crossed or new field data are acquired, the supported regional result is an evidence-opportunity map and a set of bounded local calibrations—not an archaeological absence surface.

4. Discussion

4.1 Nominal data abundance is not an observation denominator

The central result is a divergence between the apparent richness of the recovered archive and its present ability to support target-specific inference. Among the 21 source families, 17 could be identified, but strict cumulative retention fell to eight when geometry was required, to two when achieved configuration was required and to zero at primary-target fit. Allowing partial states retained more recoverable evidence but did not alter the endpoint. Because all 21 failed the adequate-negative gate, both rules also ended at zero under every gate ordering. This robustness does not make intermediate counts universal: they remain conditional on the recovered inventory, field definitions and source-family grain. It establishes only the narrower conclusion that none of the recovered sources presently supports a regional primary-target non-detection.

This distinction develops a problem already visible in national and project-level reviews. The English submerged-palaeolandscape audit found substantial legacy holdings but uneven technical reporting, reuse and project concentration (Wessex Archaeology 2013), and petroleum-exploration archives show how commercial acquisition can be archaeologically valuable while remaining institutionally difficult to discover and repurpose (Jones 2023). Archaeological quality-assurance research and marine-geophysics guidance both make reliability conditional on objectives, coverage, calibration, environment and achieved performance (Banning et al. 2017; Historic England 2025). The present non-compensatory chain makes the consequence explicit: access cannot repair unsuitable target fit, geometry cannot repair missing settings, and a strong configuration cannot create a controlled result that was never recorded.

The terminal zero should therefore be read only within the study's stringent evidential definition: no audited family qualified as effective regional T1/T2 coverage capable of supporting an areal negative inference. It is not a statement that no hydrographic, geological or archaeological work crossed the shelf, and it does not imply that future primary-data review will fail. Several families remained one or two recoverable gates short of a different classification. The result is useful because it prevents uncertainty about documentation from becoming certainty about archaeological absence.

4.2 Detection envelopes reveal target-specific rather than generic capability

The detection envelopes show why a generic “surveyed” category is analytically unstable. El Saladillo provides a local positive calibration for a buried T1 landform because a parametric sub-bottom system imaged a Last Glacial Maximum palaeochannel under documented shallow, muddy and operationally favourable conditions (Solana Rubio et al. 2026). Campamento, Camposoto and the Cádiz pilot instead calibrate exposed or proud T3 material, target relocation and modern-object discrimination. IHM and EMODnet bathymetry support seabed morphology at their output scales; MOWER and RC2912 identify navigation or line opportunities; IODP cores and seismic data provide deep regional formation and chronological context. These are complementary forms of evidence, but they are not interchangeable observations of preserved occupation-capable deposits.

The same principle is visible in comparable research. Industry three-dimensional seismic data can resolve drowned landscape architecture where processing and scale preserve the relevant landforms

(O'Leary et al. 2020). Magnetometer patterns can become palaeolandscape evidence when sub-bottom profiles, cores and onshore analogues establish the causal interpretation (van den Brenk et al. 2025). Development-led programmes can connect geophysics, cores and monitoring to a defined archaeological question (Hildebrand and York 2022). Conversely, a plausible sensor-target claim may fail independent replication even when useful landscape information is produced (Fitch and Cook Hale 2025). These examples and the Gulf of Cádiz envelopes support the same inference: capability belongs to a method-target-environment-configuration relationship, not to an instrument name.

This has two consequences for reuse. First, local positive calibration is valuable even when it cannot be extrapolated regionally. El Saladillo supplies realistic acquisition and failure conditions for future shallow-shelf work, while Campamento and the Cádiz pilot improve result semantics and ground-truth design for exposed targets. Second, contextual sources should not be devalued because they cannot support primary outcomes. MBAR24, EMODnet, IODP and validated transit records constrain morphology, depth, formation processes and survey selection; their role is to sharpen hypotheses and acquisition design, not to replace archaeological observation.

El Saladillo provides a worked example of the full evidence transition. Published SES-2000 profiles imaged a buried channel form with more than 10 m reported penetration and interpreted it as a Last Glacial Maximum palaeochannel. This establishes locally detectable late-Pleistocene channel architecture under shallow muddy conditions and advances the evidenced state to T1 within the published support. It does not establish occupation, demonstrate T2 preservation or revise regional palaeogeography because digital navigation, exact settings, raw profiles, cores and chronology remain unavailable. Archaeologically, the result narrows the next question: do channel margins, abandoned fills or adjacent surfaces preserve dated subaerial horizons or occupation-capable deposits? Answering that question requires recovery of navigation and representative profiles followed by targeted coring, chronological control and palaeoenvironmental or microarchaeological tests. The surrounding shelf cannot be treated as observed or absent.

4.3 Non-detection is an achieved property of an observation

The controlled outcome vocabulary exposed a second source of potential bias. Eleven of the 20 observation units were coded as detected, but this group combined regional morphology, a navigation opportunity, a palaeochannel, historical material and modern objects. Four Cádiz targets were not relocated, three MOWER configurations were unavailable or unusable, one record was not designed for its stated target and one was inconclusive. None was an adequate primary-target non-detection. A binary presence/absence transformation would therefore have combined unlike detections and manufactured negatives from missing or unsuitable observations.

The Cádiz pilot demonstrates the practical effect. Six inspections observed modern objects, even where source wording described a result as negative in heritage terms. Four source targets were not relocated, possibly because of burial, movement, positioning, visibility or search limitations. The first group

supports object classification; the second supports uncertainty about relocation. Neither establishes that buried prehistoric evidence was absent. Sensitivity or false-negative rates would require the complete anomaly population, selection rule and comparable search-effort denominator. Without those elements, statistical performance measures would be more precise in appearance than in meaning.

Treating adequate non-detection as an achieved observation property reverses a common default. A negative is not what remains after detections are removed. It requires a predefined target, a documented footprint, sufficient achieved resolution and penetration, suitable environmental conditions, quality-controlled processing, repeatability and archaeological review. This standard is demanding, but it is also scalable: a well-documented negative may remain bounded to a line, swath, horizon, core interval, searched radius or monitored work volume. Bounded negative evidence can be useful without being generalized to the surrounding shelf.

4.4 Survey selection and reporting support are part of the evidence

The inventory was produced by scientific, hydrographic, engineering, dredging, port-management and heritage decisions rather than by a probability sample of the seabed. Survey purpose shaped location, sensor choice, spacing, depth and archive contents. MOWER's 221.733 km in-mask route, for example, documents vessel presence but not active-system intervals. MBAR24 offers extensive fused context while concealing contributing survey effort. Port projects provide close archaeological supervision within small development footprints, whereas deep scientific programmes provide strong stratigraphic evidence outside the primary shallow domain. These contrasts are not nuisances around an otherwise representative dataset; they are the mechanism that created the evidence distribution.

As a secondary stress test, the fishing-data audit showed the same support problem in a different dataset. The Scenario 3 block was mainly covered by supplied positive C-square records, whereas the held comparison candidate lacked a current row; historical positivity showed that omission was not persistent low activity, and no locally usable explicit-zero unit existed. Rejecting the candidate prevented missing reporting support from becoming a false environmental control or rasterised overlay cells from becoming pseudo-replicates.

More generally, omitted records, inaccessible profiles and undocumented sensor intervals should remain visible analytical states. If the regional inventory becomes large enough for modelling, the initial response should be observation-support availability at the independent project or configuration level. Commissioning purpose, custodian, depth domain, infrastructure context and reporting support may then be evaluated as selection variables. Archaeological occurrence or preservation should not become the response until comparable target-capable observations and valid negative opportunities exist. The current evidence cannot estimate a causal fishing effect on preservation, occurrence or detectability.

4.5 Linking archaeological attention, observation support and heritage review

The audit links three analytical layers that are often conflated in maritime heritage assessment. Spatial archaeological attention is informed by subaerial exposure, hydrographic and lithological proxies. Observation support records whether an investigation could detect a defined target under documented conditions. Institutional review considers the resulting evidence, uncertainty and proportionality of the response. The present study addresses the evidence-production link among these layers: whether an investigation created a documented opportunity to observe the specified target and whether its result can alter the state of knowledge.

The distinction matters because none of the three layers can replace the others. A high archaeological-potential value identifies where further attention may be proportionate; it does not predict a site or prove preservation. Contextual bathymetry refines planning but does not demonstrate target-capable acquisition. A controlled result changes the evidence state only within its support and does not determine legal significance or management action. This separation is especially relevant in Spain, where cultural heritage is formally present in maritime spatial planning but evidence, procedures and responsibilities remain unevenly integrated (Canoura 2026). Institutional review remains the stage at which competent authorities consider the evidence, uncertainty, applicable policy and proportionality; the framework is evidential support, not legal advice.

The six opportunity classes preserve these distinctions spatially. P1-P3 identify higher-attention contexts with absent, regional or fine selected MBAR24 support. The classes do not decline from “unknown” to “surveyed,” because bathymetric context and archaeological observation are different dimensions. Existing lines and local calibrations are overlays rather than deductions from opportunity. This architecture allows new evidence to update the register and readiness gates without silently changing the meaning of the inherited API.

The framework therefore implements UPDATE in a deliberately bounded sense: investigate or recover evidence; classify the observation after the support gates; revise the controlled evidence state within its documented support; and return that revised state to institutional review. It does not convert the inherited API into a prior probability, calculate a posterior archaeological probability or optimise a planning action through a loss function. The bounded update is the auditable precursor that would be required before such a probabilistic extension could be justified.

4.6 A staged strategy for reducing decision-relevant uncertainty

The empirical gap hierarchy supports archive recovery before broad new acquisition where a focused request can cross a decisive gate. For MOWER, sensor-on navigation, line roles, event tables, core locations and acquisition logs would determine whether any shallow target-capable subset exists. For INSIGHT_LEG2 and El Saladillo, exact profiles, navigation, settings, processing and interpretations would permit configuration-level T1 review. For port projects, full anomaly populations, selection rules, raw geophysics and monitoring denominators would strengthen bounded T3 and works-related inference.

MBAR24 source-lineage metadata would improve hydrographic provenance, although native target review would still be required before its role could change.

Where recovery cannot supply the primary target, the secondary scenario analysis illustrates research-design consequences rather than prescribing fieldwork. Missing bathymetric context expands reconnaissance, whereas regional or fine morphology can stratify smaller footprints without removing the need for shallow sub-bottom profiling and targeted ground truth. Exact lines and stations make the comparison reproducible, but traffic, utilities, wrecks, unexploded ordnance, environmental requirements, permits, substrate, target depth, vessel limits and system trials can change every operational parameter. Field design must follow primary-profile interpretation and project-specific safety, heritage and permitting review.

4.7 Implications for archives, reporting and heritage decisions

The largest gains may come from improving the structure of survey records rather than increasing catalogue counts. A reusable minimum record should link commissioning purpose, target definition, native geometry, sensor-on intervals, achieved settings, environmental conditions, processing history, quality review and a controlled outcome register. Reports should preserve detected, adequate non-detection, inconclusive, not-relocated, unavailable and not-designed-for-target states. Where sensitive coordinates or third-party rights prevent open release, generalized footprints, stable identifiers, metadata extracts and supervised-access routes can still make evidence discoverable and auditable.

For heritage management, the framework supports proportionate information requirements. High sensitivity with weak observation support justifies targeted archive recovery or staged investigation rather than an assumption of either presence or absence. Strong contextual morphology may narrow the required field area without reducing the need for buried-target evidence. A bounded adequate non-detection may justify reducing or redirecting requirements within its support, while an inaccessible or unsuitable survey should leave uncertainty explicit. The approach therefore treats uncertainty as a condition to be managed and reduced, not as a score to be averaged away.

The operative sequence is INVESTIGATE → OBSERVE → UPDATE → DECIDE. INVESTIGATE covers focused archive recovery or a defined target-capable survey; OBSERVE requires a controlled outcome supported by the applicable readiness gates; UPDATE revises the bounded evidence state and next information requirement; and DECIDE remains the responsibility of the competent authority. Outcomes that are inconclusive, not relocated, unusable or not designed for the target may update knowledge of observation support or recovery needs, but they do not lower archaeological belief by default.

4.8 Limitations and future validation

The inventory is a dated and incomplete reconstruction of a fragmented archive. Discovery was iterative, so the raw document-level hit count and all original query strings cannot be reconstructed; the reproducible denominator begins at the 21 source families frozen on 21 August 2026. Restricted primary

datasets remained unavailable at the freeze. Subsequent custodian replies clarified catalogue and access routes but did not recover primary observations; they therefore did not alter the frozen classifications or counts. The analytical domain was limited to the Spanish Gulf of Cádiz shelf so that coordinate systems, inherited models and governance remained coherent. Portuguese and Moroccan sources were retained for context but not pooled. These choices constrain completeness and may overrepresent open metadata or responsive custodians. Online Resource 3 makes this boundary explicit and defines a prospective update protocol.

The observation register remains small and purposefully heterogeneous. Local positives do not estimate regional sensitivity, and the absence of an adequate non-detection prevents prevalence or absence modelling. The inherited potential and contextual layers retain their source assumptions and scales; scenario geometries are experiments, not field-tested designs. Gate classifications are documented study-team judgements rather than independent blinded measurements, so no inter-rater coefficient was estimated. Auditability rests on explicit definitions, locators, rationales, decisive gaps, executable checks and the strict/pass-or-partial comparison in Online Resource 2. Online Resource 3 prespecifies independent recoding, and Online Resource 4 includes the blinded 21-family workbook, but recoding has not yet been completed and no validation claim is made. No inference is made about cells lacking target-capable observation, and UPDATE remains an evidence-state transition rather than a Bayesian pre-posterior analysis.

The companion workbook operationalises that future test. One or more independent coders review the blinded source packets and assign pass, partial or fail at each gate without access to the study classifications; their decisions are frozen before discussion or adjudication. Pre-adjudication reporting should present the category distribution, contingency table, raw percentage agreement and linearly weighted Cohen's kappa with confidence intervals for each gate. A pooled statistic across all 147 family-by-gate judgements should remain secondary because it could conceal gate-specific disagreement and ignores clustering within families. Where a gate has little category variation, the contingency table and raw agreement remain primary because kappa may be unstable. Adjudicated states and reasons should then be versioned separately from the independent ratings; with more than two coders, Krippendorff's alpha may supplement the gate-specific results. Until this sequence is completed, the workbook is a validation instrument rather than evidence of reliability.

Future validation should proceed in three steps. First, the living inventory should be updated as custodians release navigation, profiles, event tables, logs and reuse decisions, with every change versioned through the readiness gates. Second, a small shallow-shelf pilot should predefine T1/T2 target expressions, achievable acquisition settings, quality checks, core-validation criteria and controlled outcomes before mobilization. Third, transferability should be tested across replicated depositional and operational settings rather than inferred from one harbour or one configuration. Only comparable labelled targets and adequate negative opportunities would justify detection-performance estimates. The protocol is intended to be transferable; the present envelope values and opportunity classes are not.

5. Conclusions

The recovered Gulf of Cádiz archive is substantial, but nominal abundance does not constitute an observation denominator for preserved prehistoric contexts. None of the 21 audited source families passed the full chain required for a regional negative inference. Extensive bathymetry, scientific tracks, deep cores and local interventions remain useful, but their evidential roles are not interchangeable.

The principal contribution is a target-specific audit that separates discovery, context, observation opportunity, achieved capability and controlled outcome. It prevents composite rasters from becoming survey effort, ship tracks from becoming active acquisition, modern-object detections from becoming prehistoric evidence and missing or unsuitable observations from becoming negatives. The archaeological result is bounded but positive: the archive supports local palaeolandscape interpretation at El Saladillo and exposed-object verification in project settings, while not supporting regional exclusion of preserved T1/T2 contexts.

The immediate priority is recovery of sensor-on geometry, primary profiles, acquisition settings, event and core tables, processing histories and complete outcome registers, followed where necessary by shallow profiling, targeted coring and chronological or palaeoenvironmental validation. Archaeological sensitivity can identify where attention is proportionate; documented observation support determines what was actually tested; admissible outcomes update only the bounded evidence state; and competent authorities decide what follows.

Declarations

Competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethics statement

The study involved no human participants or animals; ethical approval and consent were not applicable.

Declarations of generative AI and AI-assisted technologies

During preparation of this work, the authors used OpenAI ChatGPT to support language editing, code development, consistency checking and reproducibility documentation. The authors reviewed and verified the resulting text, code, calculations and interpretations and take full responsibility for the content of the publication.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial or not-for-profit sectors.

Author Contribution

A.A.N. conceived and designed the study, developed the methodology, assembled and curated the evidence, conducted the investigation, performed the formal analysis and validation, developed the software and reproducibility materials, prepared all figures, tables and supplementary materials, administered the project, and wrote, reviewed and edited the manuscript. P.F.-J. and J.C.M.-G. provided supervision. All authors reviewed and approved the final manuscript and accept accountability for the work.

Acknowledgement

The authors thank staff of the marine-data, geological, hydrographic, port and heritage institutions who clarified catalogue, access and reuse routes during construction of the survey inventory. This was an archive-based study: no new field access, intervention at a protected location or data collection was undertaken, and no additional field or location permission was required. Public, authorised or metadata-only materials were consulted under their recorded provider conditions; restricted primary data and sensitive coordinates are not redistributed. No individual respondent is responsible for the interpretations presented here.

Data Availability

Online Resource 1 contains the bounded outcome-to-update crosswalk, Online Resource 2 the source-level readiness matrix, and Online Resource 3 the discovery, accounting and validation protocol. Online Resource 4 contains the author-created source-family accounting, observation and envelope registers, blinded recoding workbook, selected opportunity and scenario summaries, validation scripts, manifests and checksums. Restricted primary geophysical data, institutional records, sensitive coordinates and third-party products are not redistributed; identifiers, custodians, access routes and reuse conditions are recorded instead.

References

1. Autoridad Portuaria de la Bahía de Cádiz (2024) Memoria preliminar y final de la actividad arqueológica preventiva de prospección arqueológica subacuática: Dragado de profundización del canal de navegación del Puerto de Cádiz. Cádiz: Autoridad Portuaria de la Bahía de Cádiz.

- <https://www.puertocadiz.com/download/memoria-preliminar-y-final-de-la-actividad-arqueologica-preventiva-de-prospeccion-arqueologica-sub-acuatica-dragado-de-canal-de-acceso/>
2. Bailey G, Flemming NC (2008) Archaeology of the continental shelf: Marine resources, submerged landscapes and underwater archaeology. *Q Sci Rev* 27(23–24):2153–2165.
<https://doi.org/10.1016/j.quascirev.2008.08.012>
 3. Bailey G, Galanidou N, Jöns H, Peeters H, Mennenga M (eds) (2020) *The Archaeology of Europe's Drowned Landscapes*. Springer, Cham. <https://doi.org/10.1007/978-3-030-37367-2>
 4. Banning EB, Hawkins AL, Stewart ST, Hitchings P, Edwards S (2017) Quality assurance in archaeological survey. *J Archaeol Method Theory* 24(2):466–488. <https://doi.org/10.1007/s10816-016-9274-2>
 5. Canoura LG (2026) The role and gaps of cultural heritage within Spain's maritime spatial planning. *Discover Oceans* 3:41. <https://doi.org/10.1007/s44289-026-00156-y>
 6. EMODnet Bathymetry Consortium (2024) EMODnet Digital Terrain Model 2024, version 12, mean-sea-level product. European Marine Observation and Data Network.
<https://emodnet.ec.europa.eu/en/bathymetry>
 7. Fitch S, Cook Hale J (2025) A needle in a haystack: Landscape survey and archaeological detection experiments in Apalachee Bay. *Archaeol Prospection* 32(2):263–280.
<https://doi.org/10.1002/arp.1959>
 8. Flemming NC, Çağatay MN, Chiocci FL, Galanidou N, Jöns H, Lericolais G, Missiaen T, Moore F, Rosentau A, Sakellariou D, Skar B, Stevenson A, Weerts H (2014) *Land Beneath the Waves: Submerged Landscapes and Sea Level Change*. European Marine Board Position Paper 21. European Marine Board, Ostend
 9. Fraile-Jurado P, Mejías-García JC (2022) Método para el cálculo, análisis y representación espacial de la variable tiempo sumergido bajo el nivel del mar durante la última glaciación en la plataforma continental del Golfo de Cádiz (España y Portugal). *Revista de Geografía Norte Grande* 81:183–205.
<https://doi.org/10.4067/S0718-34022022000100183>
 10. Fraile-Jurado P, Mejías-García JC, Roldán-Muñoz ME, Borja-Barrera C (2025) Mapping subaerial exposure timing on the Iberian Peninsula continental shelf during the Last Glacial Cycle. *Boletín de la Asociación de Geógrafos Españoles* 106. <https://doi.org/10.21138/bage.3425>
 11. Hildebrand JA, York AL (2022) Modeling and monitoring submerged prehistoric sites during offshore sand dredging and implications for the study of Early Holocene coastal occupation of Southern California. *Geoarchaeology* 37:768–782. <https://doi.org/10.1002/gea.21914>
 12. Historic England (2025) *Marine Geophysics: Data Acquisition, Processing and Interpretation Guidance Notes*, 2nd edn. Historic England, Swindon
 13. Instituto Hidrográfico de la Marina (2024) MBAR24: Modelos batimétricos de alta resolución. Centro de Descargas del Centro Nacional de Información Geográfica.
<https://cdihm.cnig.es/CentroDescargasIHM/home>

14. Jones D (2023) The offshore oil exploration industry and archaeology in the Gulf of Mexico. *J Maritime Archaeol* 18:351–369. <https://doi.org/10.1007/s11457-023-09376-1>
15. Marine Geoscience Data System (1989) RC2912 processed underway geophysical data. Interdisciplinary Earth Data Alliance. <https://doi.org/10.1594/IEDA/315260>
16. McCarthy J, Wiseman C, Woo K, Steinberg D, O'Leary M, Wesley D, Brady LM, Ulm S, Benjamin J (2022) Beneath the Top End: A regional assessment of submerged archaeological potential in the Northern Territory, Australia. *Australian Archaeol* 88(1):65–83. <https://doi.org/10.1080/03122417.2021.1960248>
17. Mejías-García JC, Fraile-Jurado P (2026) Spatial analysis of changing sea levels and its implications for understanding prehistoric human settlement patterns in the southwestern Iberian Peninsula. *J Island Coastal Archaeol* 21(1):104–129. <https://doi.org/10.1080/15564894.2025.2512155>
18. O'Leary MJ, Paumard V, Ward I (2020) Exploring sea country through high-resolution 3D seismic imaging of Australia's NW Shelf: Resolving early coastal landscapes and preservation of underwater cultural heritage. *Q Sci Rev* 239:106353. <https://doi.org/10.1016/j.quascirev.2020.106353>
19. Rodríguez Mariscal NE, Martí J (2020) Solano. Prospecciones arqueológicas en el área de Camposoto: Nuevas aportaciones sobre el navío francés Fougueux (1805), San Fernando, Cádiz. *Anuario Arqueológico de Andalucía 2010*. Junta de Andalucía. <http://hdl.handle.net/20.500.11947/21881>
20. Rojas Roca C (2014) Intervención arqueológica preventiva para el dragado de mejora de calado y limpieza en el área de la bocana del dique seco de Campamento, San Roque. *Anuario Arqueológico de Andalucía 2014*: 614–621. Junta de Andalucía. <http://hdl.handle.net/20.500.11947/20065>
21. Solana Rubio SE, Cerezo Andreo F, Ramallo Asensio SF, López-Castejón F, Bernal-Casasola D, Cau Ontiveros MÁ, Sabio R, González, Gamo E (2026) Pazos. Sub-bottom profiler in underwater archaeology: Comparative analysis for non-intrusive surveying and documentation of underwater cultural heritage in Spain. *Journal of Marine Science and Engineering* 14 (10): 943. <https://doi.org/10.3390/jmse14100943>
22. UTM-CSIC (2014) MOWER research cruise, R/V Sarmiento de Gamboa, campaign 29SG20140901. Unidad de Tecnología Marina. <https://doi.org/10.20351/29SG20140901>. Consejo Superior de Investigaciones Científicas
23. van den Brenk S, Huisman H, Willemse NW, Smit B, van Os BJH (2025) Magnetometer mapping of drowned prehistoric landscapes for archaeological heritage management in the Netherlands. *Archaeol Prospection* 32(1):3–18. <https://doi.org/10.1002/arp.1925>
24. Wessex Archaeology (2013) Audit of Current State of Knowledge of Submerged Palaeolandscapes and Sites. English Heritage Project 6231, report 84570.03. Wessex Archaeology, Salisbury

Figures






Figure 1

Target-specific evidence-audit pipeline. The T0-T4 ontology defines the evidential target; Register A supplies contextual planning support, while Register B alone enters the observation chain through preserved analytical grain, seven ordered readiness gates, controlled outcome coding and bounded update before institutional review

a Linked analytical populations



b Cumulative readiness across the 21 source families



Partial evidence identifies recovery potential; it is not equivalent to a strict pass

Figure 2

Linked analytical populations and evidence-readiness attrition. a The source-family audit branches into overlapping observation and detection-envelope populations that are not additive. b Strict and pass-or-partial cumulative retention across seven gates; both series reach zero for adequate negative inference





Figure 3

Survey opportunity and validation context on the authoritative archaeological-potential grid. Existing contextual and calibration overlays do not reduce the base opportunity class



Interpretation boundary

Planning geometry only. P2 and P3 selection balances depth × local-relief strata; P1 is geometry-centred because selected bathymetry is absent. The lines are clipped to the planning cells but require water-depth, hazard, route, instrument-range and sea-trial revision before field use. No footprint, line or station represents archaeological probability, effective survey coverage, a detection, a non-detection, or absence.

Figure 4

Illustrative design experiments for three higher-attention planning contexts. Exact geometries provide reproducibility, not navigation- or contractor-ready plans; every line and station requires primary-data, safety, heritage, operational and permitting review

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [ESM1UpdateDecideCrosswalk.pdf](#)
- [ESM2SourceReadinessMatrix.pdf](#)
- [ESM3DiscoveryandValidationProtocol.pdf](#)
- [ESM4ReproducibilityPackage.zip](#)