GSI JOURNALS
GLOBAL SCIENCE INSTITUTE
e-ISSN 2645-9108
https://dergipark.org.tr/tr/pub/aist
DOI:
GSI Journals Serie C: Advancements in Information Sciences and Technologies
(GSI Dergileri Seri C: Bilgi Bilimleri ve Teknolijilerindeki Gelişmeler)
AIST 2026, 8 (2):53-67-Research Article/Araştırma Makalesi – Received/Geliş T.:13.10.2026– Published/Kabul T.:27.10.2026

Operationalizing Nature Based Intelligence for Ecosystem Early Warning

Alper ÇABUK, Faculty of Architecture and Design, Department of Architecture, Eskişehir Technical University, acabuk@eskisehir.edu.tr, Eskişehir, Türkiye, ORCID: 0000-0002-0684-2247

Alp Deniz ÇABUK, Technical University of Munich (TUM), Munich, Germany, denizcabuk@tum.edu

Volkan TÜRKBAY, Software Developer- Freelancer, volkan@volkanturkbay.com, Eskişehir, Türkiye, ORCID: 0009-0008-9445-4551

Halil DUYMUŞ, Department of Landscape Architecture, Faculty of Architecture, Çukurova University, hduymus@cu.edu.tr, Adana, Türkiye, ORCID: 0000-0002-8113-9674

Abstract

Environmental early-warning systems often treat data quality, anomaly detection, adaptation, and provenance as separate tasks. That separation is problematic in non-stationary ecosystems, where unusual observations may represent either unreliable data or legitimate ecological novelty, and where unrestricted adaptation can normalize persistent degradation. We introduce Nature-Based Intelligence (NBI) as a computational decision-intelligence paradigm grounded in five operational principles: multiscale memory, context-conditioned normality, relational coherence, conservative adaptation, and auditable feedback. We formalize these principles in NBI-EW, an early-warning architecture combining fast and slow state representations, Mahalanobis stress energy, debiased Sinkhorn distributional displacement, optional spatial coherence, gated reference updating, and tamper-evident provenance. The engineering precursor, XAI Kernel v11, passes 48 unit and regression tests; these checks concern software behaviour, and ecological performance remains untested. We specify a leakage-controlled validation protocol using the ICOS Drought-2018 eddy-covariance product to test early detection of drought-associated functional stress at matched false-alarm rates. Temporal and site-transfer performance will be evaluated separately.

Keywords: *Nature-Based Intelligence, Ecological early warning, Drought, Concept drift, Sinkhorn divergence, Data provenance*

1. Introduction

Environmental intelligence is constrained by a problem that sits upstream of prediction. Information systems can store vast quantities of spatial and environmental observations, yet storage does not establish that the incoming observations are trustworthy, ecologically interpretable, or appropriate for updating the system's definition of normality. Geographic information standards explicitly provide principles for describing and evaluating data quality, but they do not prescribe universal minimum acceptable levels because fitness for use remains application dependent (International Organization for Standardization (ISO, 2023). In operational practice, syntactic validity, database ingestion, scientific quality control, anomaly detection, model adaptation, and provenance are frequently implemented as separate stages. That separation is convenient architecturally. It is not necessarily defensible for adaptive environmental monitoring.

The familiar 'garbage in, garbage out' problem is therefore narrower than the problem faced by ecological early-warning systems. A malformed observation should not be allowed to alter a reference state. Neither should a sensor fault. Yet a genuine ecological anomaly must remain observable precisely because it is unusual. The difficult case occurs when a technically valid observation is behaviourally inconsistent with the recent history of a system. The monitoring architecture must then distinguish at least three questions: Is the observation reliable? If reliable, is the environmental state unusual or stressed? If it is unusual, should it be permitted to redefine what the system will subsequently regard as normal? These questions are coupled, but they are not equivalent.

The problem becomes sharper in ecosystems because ecological dynamics are non-stationary, seasonal, historically contingent, and multiscale. Drought impacts can emerge rapidly, as shown by FLUXNET analyses in which gross primary productivity (GPP) frequently declined within weeks of flash-drought onset (Zhang & Yuan, 2020). Their effects can also persist beyond the meteorological event. Experimental work demonstrates ecological memory in recurrent drought responses (Canarini et al., 2021), and cross-biome flux-tower analyses show that drought legacy effects on GPP may persist for months to years and can differ in direction among ecosystems (Yu et al., 2025). A fixed reference distribution is therefore ecologically crude. An unrestricted adaptive reference creates the opposite failure mode: sustained degradation may be absorbed into the baseline and disappear statistically.

Classical ecological early-warning research approaches a related issue through resilience theory and critical slowing down. Variance, autocorrelation, recovery rates, and spatial structure can contain information about loss of resilience near some critical transitions (Scheffer et al., 2009; Dakos et al., 2012; Scheffer, Carpenter, Dakos, & van Nes, 2015). Yet empirical performance is not uniformly reliable. O'Brien et al. (2023), analysing long-term lake records, found limited applicability of both classic and second-generation early-warning signals across real ecological data. This does not invalidate resilience theory. It does, however, caution against treating any single statistical signature as a general-purpose operational warning system.

At the same time, machine-learning research on concept drift has developed extensive methods for detecting distributional change and adapting models in non-stationary streams (Gama, Žliobaitė, Bifet, Pechenizkiy & Bouchachia, 2014; Lu et al., 2019). These methods address a genuine weakness of static models, but adaptation itself is not ecologically neutral. In an environmental monitoring context, a model that adapts successfully to progressive ecosystem degradation may become statistically stable while becoming scientifically less useful. The stability-plasticity problem therefore acquires a domain-specific form: the system must learn legitimate change without automatically legitimizing deterioration.

This paper develops Nature-Based Intelligence (NBI) as a response to that problem. The concept was previously articulated as an ecological planning and decision-support approach that treats natural processes as a source of knowledge rather than merely as objects to be optimized (Çabuk & Çabuk, 2026), a published conference abstract rather than an empirical validation study. The present contribution moves from that conceptual proposition to a formal computational architecture. The novelty claim does not depend on lexical priority for the phrase 'Nature-Based Intelligence'. It concerns an explicit, falsifiable operationalization in which ecological memory,

environmental context, coupled process behaviour, constrained adaptation, and decision provenance are represented as testable mechanisms.

We use ecosystem drought stress as the first empirical target because it exposes all of these tensions simultaneously. Drought produces fast and slow responses, context-specific effects, cross-variable coupling among carbon, water and energy fluxes, and potentially long legacy periods. It also offers high-quality open observations through eddy-covariance infrastructures. The paper therefore has four aims: to define NBI in operational terms; to formalize the NBI-EW early-warning architecture; to show how integrity control and immutable provenance can be coupled to ecological interpretation without confusing data failure with ecological novelty; and to specify an empirical validation protocol that can falsify the proposed advantages before ecological performance claims are made.

2. Nature-Based Intelligence as an Operational Paradigm

Nature-Based Intelligence is defined here as a computational decision-intelligence paradigm in which environmental states are interpreted through context, history, and relationships among coupled processes, while adaptation is constrained so that unreliable observations and persistent deterioration are not automatically assimilated into the definition of normality. The term 'intelligence' is used functionally. It does not imply cognition, consciousness, intention, or optimization by ecosystems. It refers to a design logic in which decision systems learn from observable properties of ecological organization.

NBI must also be distinguished from two neighbouring fields. Nature-based solutions are interventions that protect, sustainably manage, or restore ecosystems to address societal challenges; their value and limits have been extensively discussed in climate and biodiversity research (Seddon et al., 2020). Nature-inspired computation, by contrast, commonly abstracts mechanisms from evolution, swarms, or other natural processes to construct optimization heuristics (Yang, 2020). NBI is neither an intervention category nor an optimization metaphor. It is a framework for environmental sensing, interpretation, adaptation, and decision provenance.

Table 1. Operational principles of Nature-Based Intelligence and their falsification logic.

NBI principle	Ecological rationale	Computational mechanism	Falsification test
Multiscale memory	Current ecosystem behaviour depends on both recent forcing and accumulated state.	Fast response and slow-memory representations are maintained simultaneously.	Removing the slow-memory component should reduce warning performance if memory carries useful information.
Context-conditioned normality	Seasonality, ecosystem type, and local history alter what constitutes normal behaviour.	Reference distributions are conditioned on temporal and ecological context.	A stationary reference should generate more seasonal false alarms or poorer transfer.
Relational coherence	Carbon, water, energy, and meteorological variables are coupled rather than independent channels.	Joint multivariate behaviour is monitored instead of relying on a single variable alone.	Univariate ablation should lose information where cross-process coupling is informative.
Conservative adaptation	Ecological systems change, but persistent deterioration should not be silently normalized.	Reference updates are permitted only for reliable, low-stress observations.	Unrestricted adaptation should produce more missed events during gradual degradation if the gate is useful.
Auditable feedback	Operational decisions must be reconstructable from observations, states, thresholds, and reference changes.	A tamper-evident append-only provenance chain records each transition.	Independent replay should reproduce decisions and detect post-hoc modification of recorded states.

This structure is deliberately falsifiable. Combining established algorithms is not, by itself, a scientific contribution. Exponential memory, Mahalanobis distance, Sinkhorn divergence, Moran's I, quality gates, and cryptographic hashes each have independent histories. The NBI proposition is that a particular architecture linking these mechanisms according to ecological principles provides useful behaviour under environmental non-stationarity. If the corresponding ablations do not degrade performance, the principle is not empirically supported in that application.

2.1 From Post-Hoc Quality Control to Admission-Aware Environmental Intelligence

Geospatial data infrastructures already possess mature quality concepts. ISO 19157-1:2023 defines components, measures, evaluation procedures, and reporting principles for geographic data quality (ISO, 2023). FAIR data principles also treat provenance as integral to reusable scientific data rather than as an optional annotation (Wilkinson et al., 2016). More recent geospatial interoperability work has placed integrity, provenance, and trust at the centre of AI-ready Earth-observation infrastructures (Open Geospatial Consortium [OGC], 2026). NBI therefore does not claim that conventional GIS lacks quality control. The narrower criticism is architectural: formal conformance checking and later scientific interpretation are often separated, whereas adaptive monitoring requires some admission decisions before new observations can influence system memory.

NBI reframes data acceptance as part of the sensing process. Incoming observations are evaluated for integrity before they can alter adaptive reference states. This is not equivalent to rejecting unusual values. The architecture explicitly separates unreliable data from legitimate ecological novelty. An integrity failure may be quarantined from learning while acquisition continues. A severe ecological anomaly, by contrast, should remain under observation and should trigger warning or critical monitoring rather than a data quarantine. The distinction is foundational: novelty is not error, and persistence is not automatically normality.

3. Formal Specification of NBI-EW

NBI-EW is designed as a domain-general architecture with a drought-focused ecological instantiation. Let y(i,t) denote a vector of ecosystem response variables at site i and time t, for example GPP, net ecosystem exchange (NEE), latent heat, or evapotranspiration. Let u(i,t) denote environmental forcing variables such as vapour-pressure deficit (VPD), air temperature, precipitation, and soil moisture. Let q(i,t) denote data-integrity and quality

information. The design separates three layers: observation integrity, ecological interpretation, and adaptive memory.

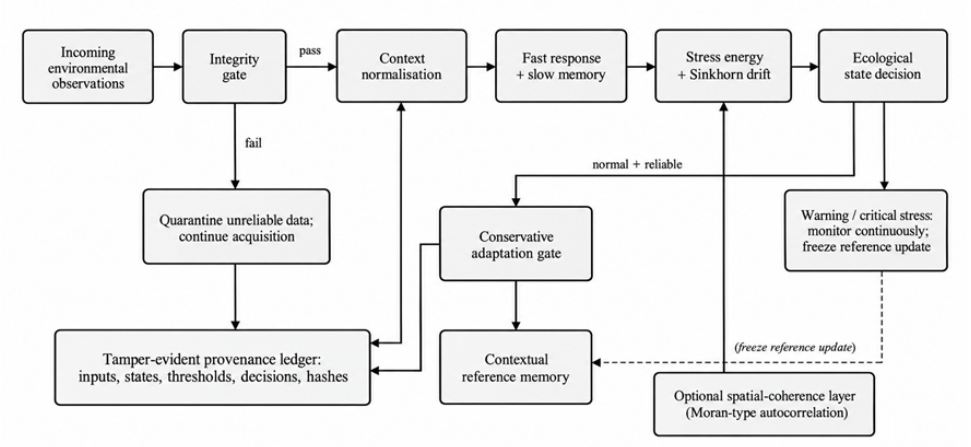

Figure 1. NBI-EW formal architecture. Integrity failure is handled separately from ecological stress; only reliable low-stress states may update contextual reference memory.

3.1 Context-Conditioned Ecological State

A context model is estimated exclusively from the calibration subset defined in Section 5.2. Context c(t) can include day of year, season, ecosystem class, and site identity where site-specific historical data are available. For response vector y(i,t), the context-normalized residual state is

$$r_{i,t} = D_i^{-1}(c_t)\,[\,y_{i,t} - \mu_i(c_t)\,] \quad (1)$$

where $\mu_i(c_t)$ is the expected context-specific state and $D_i(c_t)$ is a diagonal or robust multivariate scale operator learned from calibration data only. Full covariance whitening can be used when sample support is adequate. No test-period outcome information is permitted in μ or D. Signed residuals are retained. This is an explicit departure from the original engineering kernel, whose logarithmic absolute-value transformations were suitable for non-negative error rates but would erase physically meaningful direction in variables such as NEE.

3.2 Fast Response and Slow Ecological Memory

The fast component represents immediate change in the normalized ecological state. A fixed reference interval Δt_0 makes the change dimensionless and comparable with the slow-memory mismatch:

$$f_{i,t} = \frac{\Delta t_0}{\Delta t}[\,r_{i,t} - r_{i,t-1}\,] \quad (2)$$

A slower state estimate is updated using the exact discrete-time solution of a first-order relaxation model under piecewise constant forcing:

$$m_{i,t} = m_{i,t-1} + \alpha_t(\,r_{i,t} - m_{i,t-1}\,),\ \alpha_t = 1 - exp(\,-\Delta t/\,\tau\,) \quad (3)$$

The memory-mismatch vector is

$$p_{i,t} = r_{i,t} - m_{i,t} \quad (4)$$

Here Δt is the observation interval and Δt_0 is a fixed reference interval (one day for the proposed daily hindcast). The timescale τ has the same time unit as Δt. Its admissible range is selected on development data and frozen before test evaluation. Multiple τ values may be used when the ecological hypothesis concerns nested memory scales; each additional scale increases model flexibility and the burden of out-of-sample validation. The transient state m tracks reliable observations, whereas the contextual reference used for scoring is updated only through the adaptation gate in Section 3.6. Integrity failures do not advance m; warning and critical ecological states remain observable without being assimilated into the trusted reference.

3.3 Multiscale Stress Energy

To preserve interpretability while allowing multivariate ecological states, NBI-EW summarizes fast and slow disagreement through three non-negative quantities:

$$a_{fast} = \| f_{i,t} \|_2 , a_{slow} = \| p_{i,t} \|_2 , \kappa = | a_{fast} - a_{slow} | \quad (5)$$

Define $z(i,t) = [a_{fast}, a_{slow}, \kappa]^T$, where κ denotes disagreement between the fast and slow magnitudes and is distinct from context c(t). A regularized Mahalanobis energy is evaluated against the calibration distribution:

$$E_{i,t} = (z_{i,t} - \mu_z)^T (\Sigma_z + \lambda I)^{-1} (z_{i,t} - \mu_z) \quad (6)$$

The current v11 reference kernel uses a related three-dimensional energy score with ridge regularization. In the ecological implementation, warning and critical thresholds will not be interpreted as chi-square p-values unless distributional assumptions are demonstrated. Thresholds are operating-policy parameters estimated or selected on development data and then frozen.

3.4 Behavioural Distribution Displacement

Pointwise energy can identify abrupt departures but may miss coherent shifts in the distribution of behaviour. We therefore define a signed behavioural state vector by stacking context-normalized state, fast response, slow memory, memory mismatch, and selected forcing variables. The operator col stacks its vector arguments into one column:

$$b_{i,t} = col(r_{i,t}, f_{i,t}, m_{i,t}, p_{i,t}, u^*_{i,t}) \quad (7)$$

The forcing vector u* is scaled using calibration-period parameters only. A rolling empirical distribution P_{live} is compared with a context-appropriate reference P_{ref} using debiased Sinkhorn divergence. For entropic regularization ε,

$$S_\varepsilon (P_{ref}, P_{live}) = OT_\varepsilon (P_{ref}, P_{live}) - \tfrac{1}{2} OT_\varepsilon (P_{ref}, P_{ref}) - \tfrac{1}{2} OT_\varepsilon (P_{live}, P_{live}) \quad (8)$$

Sinkhorn divergences retain distributional geometry while correcting the entropic bias of regularized optimal transport (Feydy et al., 2019). Within NBI-EW, Sε is interpreted as behavioural displacement. It is not, by itself, evidence of causal drought impact or an approaching tipping point. The term 'behavioural state space' is used deliberately. The framework does not claim that the observations form a mathematically established manifold.

3.5. Optional Spatial-Coherence Layer

Where sampling density and ecological comparability justify spatial analysis, stress scores can be evaluated over a neighbourhood graph using Moran-type spatial autocorrelation:

$$I = \frac{n}{W} \frac{\sum_i \sum_j w_{ij} (s_i - \bar{s}) (s_j - \bar{s})}{\sum_i (s_i - \bar{s})^2} \quad (9)$$

This component is optional rather than constitutive. The 52 ICOS Drought-2018 stations are geographically dispersed and heterogeneous in biome and climate; geographic proximity does not automatically imply ecological exchangeability. Moran's I is therefore treated as supporting evidence or an ablation component, not as a mandatory NBI principle in the first drought study. Here n is the number of comparable sites, $W = \Sigma_i \Sigma_j w_{ij}$ is the total non-negative spatial weight, s_i is a stress score, and $\bar{s}$ is its mean. The statistic is undefined when W = 0 or the score variance is zero; such cases are recorded as unavailable, not zero coherence. Neighbourhood construction and any permutation scheme are fixed using calibration and development information only.

3.6 Integrity Gate, Ecological State, and Conservative Adaptation

The v11 engineering kernel aggregates detector severity into a single global level and invokes hard-halt behaviour at the critical level. That design is appropriate for a geospatial ingestion pipeline in which critical geometry or

parsing failures should stop state contamination. It is not appropriate to transfer this rule directly to ecosystem stress. NBI-EW therefore separates integrity severity q_t from ecological severity e_t.

Table 2. Separation of data-integrity policy from ecological-stress policy.

Condition	Operational action	Reference-memory update
Integrity pass; ecological normal	Admit observation; continue monitoring	Eligible if adaptation criteria are met
Integrity pass; ecological warning	Admit observation; issue warning; continue monitoring	Frozen
Integrity pass; ecological critical	Admit observation; issue critical alert; continue high-frequency monitoring	Frozen
Integrity suspect	Flag; retain separately for review; continue acquisition	Frozen
Integrity critical	Quarantine from ecological state update; preserve raw evidence; operator review	Frozen

Define A(i,t) as the adaptation gate. In the simplest implementation,

$$A_{i,t} = 1\{ q_{i,t} = pass \} \cdot 1\{ e_{i,t} = normal \} \cdot 1\{ context\ criteria\ satisfied \} \quad (10)$$

Only observations with A(i,t)=1 may update the context-conditioned reference. This rule operationalizes the distinction between legitimate non-stationarity and automatic normalization of degradation. The gate does not prove that accepted observations are ecologically optimal or undisturbed; it merely prevents observations already classified as unreliable or stressed from immediately redefining the reference against which subsequent stress is assessed.

3.7 Tamper-evident Provenance and the Limits of Local WORM Simulation

Each state transition is recorded with its timestamp, input identifiers, quality metrics, contextual reference version, detector outputs, thresholds, decision state, cursor, and an optional snapshot of the adaptive state. A canonical record R_t is chained to the previous record through a cryptographic digest:

$$h_t = SHA3\text{-}512(canonical(R_t) \parallel h_{t-1}) \quad (11)$$

The reference implementation additionally authenticates canonical record payloads using HMAC-SHA3-512 under a secret key and publishes per-record anchors to a no-overwrite store. This changes the governance question from 'who remembers what happened?' to 'can the recorded decision path be independently verified as unchanged?' The mechanism should not be oversold. The present directory-based anchor store is a WORM simulation, not a production immutability guarantee: an administrator controlling the filesystem can delete records. A production deployment would require a separate trust domain with retention-lock semantics or an equivalent append-only service. The provenance model can also be mapped to interoperable provenance vocabularies such as W3C PROV-O (Lebo, Sahoo, & McGuinness, 2013).

4. Reference implementation and software verification

The starting software artifact is XAI Kernel v11 (A. Çabuk & A. D. Çabuk, n.d.), a Python-based geospatial data-quality monitoring kernel. Despite its historical name, it is not a SHAP or feature-attribution system. It ingests paginated WFS features, evaluates geometry validity and parsing failures, detects spatial clustering of invalid geometries, constructs a five-dimensional behaviour vector from error-rate dynamics, computes a calibrated stress-energy score, monitors distributional displacement with debiased Sinkhorn divergence, and records decisions in a HMAC-authenticated chained ledger with separate anchor files.

The v11 behaviour vector is $b_t = [x_t, \Delta x_t, s_{L,t}, dL1_t, dL2_t]$, where x_t is the observed error rate, Δx_t is its increment, $s_{L,t}$ is a slow state estimate, and dL1/dL2 are logarithmic relative-change terms. The kernel calibrates its stress-energy detector from an initial clean reference subset, uses a regularized covariance estimate, and applies an independent absolute quality policy. Spatial error clustering is evaluated using a permutation-based Moran's I statistic. The current Sinkhorn drift threshold and several quality thresholds are policy values requiring domain-specific validation rather than universal constants.

The supplied v11 archive was inspected, and its test suite was independently rerun during manuscript revision. All 48 tests passed, comprising 33 original tests and 15 regression scenarios as classified in the package README. The verification was limited to the supplied engineering implementation and the behaviour exercised by these unit and regression tests. It does not establish ecological validity, real-time operational performance, or live integration reliability with the original external WFS service. Accordingly, NBI-EW inherits software design patterns from v11, not an empirical claim of ecological performance.

Table 3. Required translation from the engineering reference implementation to the ecological NBI-EW kernel.

Component	v11 engineering kernel	NBI-EW ecological kernel
Input signal	Geometry error rate	Multivariate ecosystem responses and environmental forcing
State representation	R^5 behaviour vector from error-rate dynamics	Signed context-normalized ecological state with fast and slow components
Reference	First accepted low-error batches	Historical context-conditioned reference estimated from calibration subset only
Critical response	Global hard halt	Hard quarantine only for integrity failure; ecological critical states remain monitored
Spatial layer	Moran's I on invalid geometries	Optional Moran-type coherence on ecological stress scores
Transformation	Absolute-value logarithmic changes	Sign-preserving scaling appropriate to physical variables
Provenance	Hash-chained HMAC-authenticated ledger + local WORM simulation	Same design pattern, with production retention-lock backend required for immutability
Thresholds	Policy defaults	Development-set calibration at matched false-alarm targets; frozen before test

5. ICOS Drought 2018 Validation Protocol

The first empirical test will use the Drought-2018 ecosystem eddy-covariance product distributed by the Integrated Carbon Observation System (ICOS). The product contains observations from 52 ecosystem stations, with station records extending between 1989 and 2018 depending on site, and was processed through the ONEFlux pipeline in a format compatible with FLUXNET products (Drought 2018 Team & ICOS Ecosystem Thematic Centre, 2020; Pastorello et al., 2020). This combination of long pre-event histories, carbon-water-energy fluxes, meteorological forcing, and the severe 2018 European drought makes the dataset suitable for a retrospective early-warning hindcast.

Table 4. Proposed station and period inclusion rules for the ICOS validation.

Criterion	Prespecified rule
Minimum historical support	At least 4 complete pre-2018 growing seasons at development sites: at least 3 for calibration and at least 1 subsequent season for development. Held-out sites require at least 3 for site-specific calibration.
Response availability	GPP required; NEE and latent heat/evapotranspiration included where quality and completeness criteria are met.
Forcing availability	Air temperature and VPD required for core model; precipitation and soil moisture included when station coverage is adequate.
Data quality	Use published quality flags; exclude periods failing prespecified reliability rules. Missingness thresholds fixed before outcome evaluation.
Leakage control	Context, scaling, covariance, and initial references: calibration only. Hyperparameters, thresholds, and event rules: development only. 2018: frozen held-out evaluation; causal adaptation follows the fixed gate.
Duplicate control	Overlapping station-years from multiple archives are treated as the same observation source, not as independent external validation.

5.1 Outcome definition

The primary outcome is a drought-associated functional-stress event, not ecosystem collapse. The proposed event definition combines a sustained negative context-normalized GPP anomaly with a water-stress criterion based on VPD, soil moisture, or both. The anomaly cutoff, minimum duration, water-stress cutoff, and event-merging rule are selected using the pre-2018 development subset of development stations defined in Section 5.2, supported by ecological rationale. The context-normalization parameters come only from calibration data. Event definitions are finalized before any 2018 test labels are examined and remain common to all methods. This framework paper specifies the selection procedure; it reports neither selected numerical thresholds nor an executed ecological experiment.

The forecasting task is prospective in information structure. At each time t, predictors may use observations available at or before t. The outcome asks whether a new functional-stress episode begins within a 7-day or 14-day horizon. Future observations are allowed only to label the outcome retrospectively. This prevents a common leakage error in which post-event information enters predictor construction.

5.2 Development, Temporal Transfer, and Site Transfer

The primary hindcast reserves 2018 as the temporal test year. Before inspecting 2018 response trajectories or event labels, eligible stations are partitioned into development sites and held-out sites using station metadata, with biome and climate coverage considered where sample size permits. The station list, exclusions, split seed or deterministic assignment rule, and data version are recorded in a frozen split manifest. No station identities or completed split are claimed in this framework paper.

At each eligible development site, the earliest complete pre-2018 growing seasons form D_{cal}, with at least three complete seasons required. The subsequent complete pre-2018 seasons form D_{dev}, with at least one complete season required. Thus eligibility for the primary temporal experiment requires at least four complete pre-2018 growing seasons. D_{cal} estimates seasonal means and scales, forcing normalization, covariance, initial slow memory, and initial Sinkhorn references. D_{dev} selects memory timescales, window lengths, regularization, adaptation policy, warning and critical thresholds, and the common event rule. Development selection uses chronological blocked validation; no random daily mixing is permitted.

The 2018 observations at development sites form the temporal-transfer test. The 2018 observations at held-out sites form the site-transfer test. These test subsets remain separate in reporting. At a held-out site, eligible pre-2018 history may be used solely for unsupervised site-specific context calibration when the deployment scenario assumes an existing site archive. It is not used for outcome-based tuning, threshold selection, or fitting shared hyperparameters. This is transfer to a site with historical calibration data; a separate pooled-context analysis is required to claim transfer to sites without site-specific histories.

The observation subsets are pairwise disjoint. D_{test} is the union of the temporal-transfer and site-transfer test subsets, and the site sets used for development and site-transfer evaluation are disjoint. All fitted context parameters and operating thresholds are frozen before test evaluation. During test-time streaming, only the prespecified adaptation rule may update trusted reference memory, using reliable observations available at or before the current time. No test outcome label may change parameters, gates, event definitions, or model selection.

Chronological boundaries include an embargo at least as long as the maximum predictor lookback plus the 14-day outcome horizon; boundary-crossing predictor or label windows are excluded from scoring and tuning. State initialization uses calibration history only, and subsequent states are advanced causally. Baselines receive identical observations, split boundaries, and availability restrictions. An external post-2018 replication is a distinct future evaluation; overlapping station-years are removed and its dataset is versioned before analysis.

$$D = D_{cal} \cup D_{dev} \cup D_{test}, D_a \cap D_b = \emptyset \text{ for } a \neq b \quad (12)$$

Table 5. Primary hypotheses for empirical validation.

Hypothesis	Falsifiable prediction
H1: Early-warning utility	At a prespecified matched false-alarm rate, NBI-EW improves event recall relative to baseline methods. Median lead time is a secondary comparison.
H2: Multiscale memory	Removing slow memory reduces performance, particularly for gradually developing stress.
H3: Context conditioning	Stationary references increase seasonal false alarms or reduce transfer performance.
H4: Relational coherence	A multivariate state outperforms an otherwise equivalent GPP-only configuration when coupled processes are informative.
H5: Conservative adaptation	Unrestricted reference updating reduces apparent drift but increases missed detections during persistent deterioration.

Baselines will include a context-adjusted GPP anomaly threshold, EWMA and CUSUM change monitoring, classical early-warning indicators such as variance and lag-1 autocorrelation where their assumptions are defensible, and at least one multivariate anomaly detector that receives the same predictor information as NBI-EW. A supervised forecasting model may be included as a predictive ceiling or comparison, but it must be separated conceptually from unsupervised warning methods. The same data-access rules apply to every method.

5.3 Performance Metrics and Uncertainty

Pointwise accuracy is not the primary evaluation criterion because dense ecological time series can make trivial normal predictions appear strong. Performance will be reported at the event level using event recall, missed-event rate, false alarms per site-year, warning lead time, and precision-recall summaries. The primary comparison is event recall at a prespecified matched false-alarm rate; median warning lead time is secondary. The false-alarm operating target and success criterion are fixed on development data before the test is opened. Lead time is summarized alongside recall because estimates conditional on detected events can favour methods that miss difficult events. Uncertainty intervals will be estimated using resampling schemes that preserve site and temporal dependence, such as block bootstrap or site-level bootstrap, rather than treating daily observations as independent.

5.4 Retrospective Processing and Causal Availability

ONEFlux products contain gap-filled and derived variables, including GPP estimates. Pastorello et al. (2020) describe a standardized processing pipeline that is scientifically valuable but not identical to a live sensor stream. Consequently, success on the complete processed archive would demonstrate retrospective early-warning skill, not automatically real-time operational skill. A causal-availability sensitivity analysis will therefore reconstruct predictor variants using only information that could have been available at or before time t. Results from fully processed products and causal variants will be reported separately.

6. Discussion

NBI-EW is motivated by a specific failure mode of adaptive environmental intelligence: the same mechanism that allows a system to cope with non-stationarity can allow it to forget that degradation occurred. This makes environmental adaptation qualitatively different from many conventional concept-drift settings. In commercial prediction, rapid adaptation to a new distribution can be desirable if prediction error falls. In ecological monitoring, a new distribution may be the phenomenon of concern. The reference model must therefore be permitted to change without being permitted to erase evidence indiscriminately.

The distinction between integrity and novelty is equally important. A sensor failure and a drought-induced GPP decline may both appear as departures from historical behaviour, but the operational response should be opposite.

Unreliable data should be prevented from contaminating adaptive memory. Genuine ecological novelty should remain visible and should intensify monitoring. This is where the original geospatial kernel's hard-halt logic requires domain-specific redesign. NBI does not abandon fail-closed principles; it places them at the integrity boundary rather than at the ecological anomaly boundary.

The approach also extends conventional GIS quality reasoning without claiming to replace established standards. ISO 19157-1 provides a formal vocabulary and process for evaluating geographic data quality, FAIR principles require detailed provenance for reusable data, and W3C PROV provides interoperable structures for representing provenance (Lebo et al., 2013; Wilkinson et al., 2016; ISO, 2023). NBI adds an operational policy layer: whether a new observation is allowed to change the internal reference state of an adaptive environmental model. This is a narrower claim than saying that existing GIS is inherently 'garbage in, garbage out', and it is technically more defensible.

A second contribution is methodological restraint. The framework does not assume that every ecosystem stress event is a critical transition. It does not interpret Sinkhorn drift as a tipping point, Moran's I as causality, or a model explanation as ecological mechanism. Classical resilience indicators and modern drift detectors are treated as neighbouring tools whose performance must be compared, not as straw-man alternatives. This matters because empirical ecological early-warning performance is heterogeneous and preprocessing choices can dominate apparent success (O'Brien et al., 2023).

Several limitations are already clear. First, NBI-EW contains design choices that can introduce flexibility, including context windows, memory timescale, behaviour-vector composition, Sinkhorn regularization, and thresholds. A strong evaluation must therefore freeze choices before held-out testing and report sensitivity rather than optimize each site retrospectively. Second, site heterogeneity complicates global thresholding. Third, flux-derived GPP is model-dependent and may include retrospective processing. Fourth, the present WORM implementation is only a local simulation and should not be described as immutable infrastructure. Fifth, the first drought application cannot establish NBI as universally superior across environmental domains. At most, it can support or reject the specified mechanisms in this problem class.

These limitations are not peripheral. They define the conditions under which NBI can become a scientific framework rather than a label. The framework earns explanatory value only if the principles can be removed, challenged, and shown to matter under transfer. A negative ablation result is therefore informative. If context-conditioned references do not reduce false alarms, or if gated adaptation does not protect against normalization of degradation, those principles should be revised rather than preserved rhetorically.

7. Conclusions

Nature-Based Intelligence is proposed here as an operational environmental decision-intelligence paradigm, not as a claim that ecosystems possess cognition and not as a new name for nature-inspired optimization. Its central proposition is architectural: environmental intelligence should represent memory, context, coupling, constrained adaptation, and provenance together because these properties interact under non-stationarity.

The NBI-EW specification converts that proposition into a falsifiable system. Incoming observations first face an integrity gate. Reliable observations are interpreted relative to context and multiscale memory. Distributional displacement is measured without equating change with collapse. Ecological warning freezes adaptation but does not stop observation. Integrity failure can quarantine data from the adaptive state. Every decision is recorded in a tamper-evident provenance chain. NBI-EW is designed so that novelty is not mistaken for error and persistence is not automatically accepted as normality.

The immediate next step is empirical rather than rhetorical. The ICOS Drought-2018 experiment described here will determine whether the proposed architecture provides earlier or more reliable warning of drought-associated functional stress at matched false-alarm rates. Until that evaluation is completed, NBI should be treated as a

formally specified and software-grounded hypothesis. If the predicted gains survive temporal transfer, site transfer, ablation, and causal-availability tests, the framework will have earned a stronger claim.

Declarations

Competing interests: Patent-related interests in the underlying proprietary software are disclosed in the separate title-page declarations supplied to the editors.

Funding and acknowledgements: Details are provided in the separate title-page declarations for anonymous review.

Ethics approval: Not applicable. The proposed empirical study uses open environmental observation data and does not involve human or animal subjects.

Data availability: No new ecological dataset or empirical ecological result is reported in this framework paper. The proposed validation uses the publicly available ICOS Drought-2018 ecosystem eddy-covariance product (Drought 2018 Team & ICOS Ecosystem Thematic Centre, 2020), https://doi.org/10.18160/YVR0-4898. The station-selection and split rules are specified in Section 5.2. No completed station manifest, external replication dataset, or processed ecological validation archive is claimed.

Code availability: XAI Kernel v11 is proprietary and is not publicly deposited or licensed as open source. This manuscript grants no permission to use or redistribute the implementation. Access requests are subject to authorization by the relevant rights holders and applicable terms; access is not guaranteed. The mathematical specification and validation protocol support independent implementation but do not replace access to the exact proprietary source for full computational reproduction. No public ecological NBI-EW release or repository is claimed.

Declaration of AI Assisted Research and Writing

ChatGPT and OpenAI Codex were used to support manuscript organization, language revision, source discovery, document editing, and inspection of the supplied software and test outputs. The software checks reported here were produced by executing the supplied test suite; they were not generated as ecological observations by a language model. AI-generated suggestions are not primary evidence. No AI-generated ecological dataset, performance estimate, or validation outcome is reported. The authors retain responsibility for the scientific claims, mathematical specification, source accuracy, interpretation, and final manuscript. AI tools are not authors and cannot approve the manuscript or assume scientific accountability.

References

Çabuk, A., & Çabuk, A. D. (n.d.). XAI Kernel v11 [Proprietary software; author-supplied, non-public archive]. Unpublished software.

Çabuk, S. N., & Çabuk, A. (2026). Nature based intelligence for sustainable and resilient cities [Conference abstract]. In XI. International Science and Academic Congress, *Congress Abstract Book* (pp. 68–69), 11–13 March 2026, Ereğli/Konya, Türkiye. https://www.insackongre.org/_files/ugd/e42d97_08e2363933a544a684c22e54488d3370.pdf

Canarini, A., Schmidt, H., Fuchslueger, L., Martin, V., Herbold, C. W., Zezula, D., Gündler, P., Hasibeder, R., Jecmenica, M., Bahn, M., & Richter, A. (2021). Ecological memory of recurrent drought modifies soil processes via changes in soil microbial community. *Nature Communications, 12*, 5308. https://doi.org/10.1038/s41467-021-25675-4

Dakos, V., Carpenter, S. R., Brock, W. A., Ellison, A. M., Guttal, V., Ives, A. R., Kéfi, S., Livina, V., Seekell, D. A., van Nes, E. H., & Scheffer, M. (2012). Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. *PLOS ONE, 7*(7), e41010. https://doi.org/10.1371/journal.pone.0041010

Drought 2018 Team, & ICOS Ecosystem Thematic Centre. (2020). Drought-2018 ecosystem eddy covariance flux product for 52 stations in FLUXNET-Archive format [Data set]. ICOS Carbon Portal. https://doi.org/10.18160/YVR0-4898

Feydy, J., Séjourné, T., Vialard, F.-X., Amari, S.-I., Trouvé, A., & Peyré, G. (2019). Interpolating between optimal transport and MMD using Sinkhorn divergences. In Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics (*Proceedings of Machine Learning Research, Vol. 89*, pp. 2681–2690). https://proceedings.mlr.press/v89/feydy19a.html

Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys, 46*(4), Article 44. https://doi.org/10.1145/2523813

International Organization for Standardization. (2023). ISO 19157-1:2023 Geographic information — Data quality — Part 1: General requirements. https://www.iso.org/standard/78900.html

Lebo, T., Sahoo, S., & McGuinness, D. (Eds.). (2013). PROV-O: The PROV Ontology. W3C Recommendation, 30 April 2013. World Wide Web Consortium. https://www.w3.org/TR/prov-o/

Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2019). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering, 31*(12), 2346-2363. https://doi.org/10.1109/TKDE.2018.2876857

O'Brien, D. A., Deb, S., Gal, G., Thackeray, S. J., Dutta, P. S., Matsuzaki, S.-I. S., May, L., & Clements, C. F. (2023). Early warning signals have limited applicability to empirical lake data. *Nature Communications, 14*, 7942. https://doi.org/10.1038/s41467-023-43744-8

Open Geospatial Consortium. (2026). OGC Testbed-21: Data Quality for Integrity, Provenance and Trust Report (OGC Engineering Report 26-005r1). https://docs.ogc.org/per/26-005r1.html

Pastorello, G., Trotta, C., Canfora, E., Chu, H., Christianson, D., Cheah, Y.-W., Poindexter, C., Chen, J., Elbashandy, A., Humphrey, M., Isaac, P., Polidori, D., Reichstein, M., Ribeca, A., van Ingen, C., Vuichard, N., Zhang, L., Amiro, B., Ammann, C., … Papale, D. (2020). The FLUXNET2015 dataset and the ONEFlux processing pipeline for eddy covariance data. *Scientific Data, 7*, 225. https://doi.org/10.1038/s41597-020-0534-3

Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., van Nes, E. H., Rietkerk, M., & Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature, 461*, 53-59. https://doi.org/10.1038/nature08227

Scheffer, M., Carpenter, S. R., Dakos, V., & van Nes, E. H. (2015). Generic indicators of ecological resilience: Inferring the chance of a critical transition. *Annual Review of Ecology, Evolution, and Systematics, 46*, 145-167. https://doi.org/10.1146/annurev-ecolsys-112414-054242

Seddon, N., Chausson, A., Berry, P., Girardin, C. A. J., Smith, A., & Turner, B. (2020). Understanding the value and limits of nature-based solutions to climate change and other global challenges. *Philosophical Transactions of the Royal Society B: Biological Sciences, 375*(1794), 20190120. https://doi.org/10.1098/rstb.2019.0120

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., … Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data, 3*, 160018. https://doi.org/10.1038/sdata.2016.18

Yang, X.-S. (2020). Nature-inspired optimization algorithms: Challenges and open problems. *Journal of Computational Science, 46*, 101104. https://doi.org/10.1016/j.jocs.2020.101104

Yu, X., Orth, R., Reichstein, M., Reimers, C., Gomarasca, U., Migliavacca, M., Papale, D., Bahn, M., & Bastos, A. (2025). Widespread but divergent drought legacy effects on gross primary productivity across biomes. *Global Change Biology, 31*(10), e70541. https://doi.org/10.1111/gcb.70541

Zhang, M., & Yuan, X. (2020). Rapid reduction in ecosystem productivity caused by flash droughts based on decade-long FLUXNET observations. *Hydrology and Earth System Sciences, 24*, 5579-5593. https://doi.org/10.5194/hess-24-5579-2020