**Entomology and Applied Science Letters**
**Volume 13, Issue 1, Page No: 86-98**

Available Online at: www.easletters.com

**ISSN No: 2349-2864**

# A Digital Twin of Mosquito-Borne Disease Control Linking Vector Life Cycles, Insecticide Resistance, Human Mobility, Weather, and Intervention Feedback

**Andrew Miller[1]*, Jessica Davis[1], Daniel Carter[2], Siti Rahman[3]**

*[1]Department of Computational Epidemiology and Digital Twin Engineering, College of Engineering, Georgia Institute of Technology, Atlanta, United States.*
*[2]Department of Spatial Epidemiology and Urban Health, Faculty of Public Health, National University of Singapore, Singapore.*
*[3]Department of Vector Control and Intervention Modelling, Faculty of Health Sciences, University of the Witwatersrand, Johannesburg, South Africa.*

## ABSTRACT

Mosquito-control decisions are made within biological and operational systems that change across life stages, locations, weather conditions, movement patterns, intervention histories, and surveillance cycles. Existing predictive models can represent selected parts of this complexity, but they are commonly separated from one another, updated irregularly, or interpreted beyond the conditions under which their outputs are valid. This architecture article addresses the resulting gap by proposing an evidence-grounded mosquito-control digital twin that links vector life-cycle states, environmental and habitat inputs, human mobility, insecticide resistance, intervention response, surveillance assimilation, state updating, simulation feedback, validation, and governance. The approach integrates mechanistic biological representation with multimodal observations and data-driven updating while preserving the distinction between observed quantities, latent state estimates, inferred relations, and proposed architectural components. The strongest defensible synthesis is that mosquito control requires a dynamic, spatially connected, and uncertainty-aware representation because environmental forcing, population processes, movement, resistance, and intervention effects vary across species, populations, places, time scales, and measurement systems. However, increased simulation detail does not establish real-world control effectiveness; state estimation does not provide complete observability; assimilation may reproduce surveillance bias; and an architectural specification does not constitute an operationally validated system. The principal implication is that development should proceed through modular validation, explicit parameter and data provenance, prospective testing, human authorization, monitoring, and rollback rather than through unqualified automation. The proposed structure is therefore a bounded scholarly architecture for organizing evidence, hypotheses, updating processes, and decision evaluation, not a deployment-ready mosquito-control platform.

**Keywords:** Mosquito-control digital twin, Predictive entomology, Artificial intelligence, Sensor surveillance, Remote sensing, Human mobility.

**HOW TO CITE THIS ARTICLE:** Miller A, Davis J, Carter D, Rahman S. A Digital Twin of Mosquito-Borne Disease Control Linking Vector Life Cycles, Insecticide Resistance, Human Mobility, Weather, and Intervention Feedback. Entomol Appl Sci Lett. 2026;13(1):86-98. https://doi.org/10.51847/SNJdGqW4n8

**Corresponding author:** Andrew Miller
**E-mail** ✉ andrew.miller@gatech.edu
**Received:** 04/12/2026
**Accepted:** 11/03/2026

## INTRODUCTION

Mosquito-control programmes must interpret observations generated by traps, environmental sensors, remote sensing, insecticide-resistance assays, epidemiological surveillance, and intervention records while the biological system represented by those observations continues to change. A digital twin offers a possible organizing construct, but the term should not be applied to

any model, dashboard, or periodically rerun simulation. A mosquito-control digital twin must be more than a stand-alone simulator: it requires persistent data-mediated linkage to the represented system, while ecological use also requires process knowledge to be combined with place-specific data-driven learning [1, 2]. This definition places synchronization, biological representation, uncertainty, and decision feedback at the centre of the architecture.

The scientific rationale is supported by forecasting evidence showing that vector and human states can be represented jointly but not with uniform accuracy. Mechanistic ensemble forecasting has produced linked estimates of mosquito infection and human disease outcomes, although skill varied among locations and forecast targets [3]. Such variation is not a peripheral technical problem. It indicates that observation systems, ecological structure, parameter values, and transmission conditions differ across control jurisdictions. A useful twin must therefore retain local uncertainty and model discrepancy rather than converting a forecast into an apparently complete description of the system.

The decision environment is also nonstationary. Climate change can redistribute Aedes-borne transmission suitability across regions, seasons, and vector species rather than producing a uniform increase in risk [4]. Local mosquito abundance, pathogen transmission, and intervention demand can consequently move in different directions even under a common climate scenario. Weather is only one part of that change: habitat availability, vector movement, human mobility, prior control pressure, resistance, reporting delays, and operational coverage can modify the observed trajectory. A static risk surface or fixed-parameter population model cannot represent all these transitions as an updating control problem.

The unresolved gap is therefore architectural rather than merely predictive. Existing evidence supports individual components—thermal response functions, population models, mobility networks, resistance maps, surveillance filters, and intervention models—but does not by itself define how they should be connected, updated, interpreted, validated, and governed as one mosquito-control digital twin. This article develops an explicitly non-validated architecture that separates empirical inputs from latent states, mechanistic relations from data-driven corrections, and simulated intervention effects from demonstrated field effectiveness. Its central argument is that the twin should function as a modular, uncertainty-aware representation for hypothesis testing and bounded decision support, not as an autonomous or intrinsically trustworthy representation of mosquito-control reality.

*Why mosquito control requires dynamic system representation*

Mosquito populations and transmission potential emerge from transitions whose rates change with environmental conditions and biological history. Dynamic representation is necessary because thermal and hydroclimatic effects are nonlinear and interact with biological traits and prior system state [5–7]. Temperature can simultaneously affect development, mortality, biting, vector competence, and pathogen incubation, while drought or rainfall may modify habitat availability, host contact, and transmission differently among regions. These pathways cannot be reduced safely to a single linear weather coefficient. Their representation must instead preserve bounded responses, delayed effects, species and population differences, and the possibility that the same environmental observation acts through different mechanisms in different systems.

Model structure also changes what can be inferred. Mechanistic mosquito-borne disease models vary in their aims, compartments, parameter sources, spatial assumptions, and treatment of vector and host processes [8]. A model designed to reproduce seasonal incidence may not contain the stage structure needed to evaluate larval control, and a model designed for regional invasion may not resolve household exposure or intervention delivery. Parameter transfer can create an appearance of biological precision while importing assumptions from another population, climate regime, vector species, or surveillance system. Dynamic system representation therefore requires purpose-specific components, explicit parameter provenance, and alternative structural hypotheses. Increasing complexity is defensible only when the added component corresponds to an identifiable mechanism, observable

consequence, decision requirement, or validation test.

Vector traits themselves may not remain constant under changing environmental conditions. Phenotypic plasticity can alter predicted transmission, but its magnitude and direction may depend on population history, exposure conditions, and the traits included in the model [9]. Apparent trait change may alternatively reflect adaptation, population replacement, altered species composition, or biased observation. The proposed twin should therefore permit selected parameters to vary through constrained updating or competing scenarios without treating every discrepancy as biological adaptation. This preserves four essential boundaries: simulation fidelity is not equivalent to real-world control effectiveness; state estimation is not equivalent to complete observability; data assimilation is not equivalent to unbiased updating; and a digital-twin architecture is not equivalent to a validated operational system.

The evidence dimensions and interpretive boundaries for mosquito control requires dynamic system representation are summarized in **Table 1**.

**Table 1.** Why Mosquito Control Requires Dynamic System Representation: Data Inputs, Analytical Tasks, Biological Representations, Validation, Explainability, Generalization, and Decision Boundaries

| Analytical component or task | Input data | Model or representation | Expected output | Validation requirement | Explainability or interpretation need | Generalization risk | Decision boundary | Representative supporting reference(s) |
|---|---|---|---|---|---|---|---|---|
| **Define the digital-twin relation** | Surveillance streams, intervention records, environmental observations, model versions | Persistently linked physical–virtual representation | Versioned estimate of the represented control system | Confirm update cadence, data linkage, provenance, and intended use | Distinguish observed data, inferred states, simulations, and recommendations | Engineering definitions may not transfer directly to ecological systems | A connected simulation is not automatically a digital twin | [1] |
| **Integrate process and data representations** | Biological process knowledge, field observations, environmental data, data-driven predictions | Hybrid mechanistic and data-driven representation | Place-specific state estimates and scenarios | Compare hybrid, mechanistic-only, and data-driven alternatives under withheld observations | Show which outputs arise from biological assumptions and which arise from statistical correction | Local ecological relations may change among places and seasons | Data integration is not complete biological representation | [2] |
| **Estimate linked vector and human states** | Mosquito surveillance, human cases, transmission states | Mechanistic ensemble assimilation model | Mosquito infection and human-disease forecasts | Evaluate calibration and accuracy by location, season, and forecast target | Retain interpretable compartments, observation models, and uncertainty intervals | Surveillance systems and transmission ecology differ among jurisdictions | Forecast skill is not evidence of intervention effectiveness | [3] |
| **Represent nonstationary environmental forcing** | Temperature, rainfall, drought, immunity, and temporal context | Nonlinear trait-based and epidemic representations | Environment-conditioned transition rates and risk trajectories | Test across climates, vector–pathogen systems, and temporal scales | Identify which life-history or transmission pathway each environmental variable modifies | Thermal and hydroclimatic relations may not transfer across species or regions | Climatic suitability is not realized incidence or controllability | [5] |
| **Specify model purpose and parameter provenance** | Published structures, parameter estimates, surveillance sources, and modelling aims | Modular mechanistic representation with documented assumptions | Purpose-specific state and prediction outputs | Compare structural alternatives and test transferred parameters locally | Expose compartments, parameter sources, omissions, and identifiability limits | Imported parameters can create false precision | Greater model detail does not ensure greater decision value | [8] |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **Represent context-dependent trait change** | Environmental history, population traits, and transmission parameters | Constrained plasticity, adaptation, or alternative-parameter scenarios | Conditional trait and transmission trajectories | Discriminate plasticity from adaptation, population turnover, and measurement error using longitudinal evidence | Label trait change as observed, estimated, or hypothesized | Population-specific responses may not generalize | Updated parameters are not proof of a biological mechanism | [9] |

*Vector life-cycle and population components*

The population core of the proposed twin should represent mosquito development and survival as a set of linked but partly latent states rather than as one undifferentiated abundance index. Eggs, larvae, pupae, newly emerged adults, host-seeking adults, infected adults, and infectious adults are connected through transitions that respond differently to habitat, temperature, density, movement, and intervention exposure. Population representation must also connect local establishment with dispersal because climate and movement pathways jointly shape the invasion potential of *Aedes aegypti* [10]. These processes imply that a local decline in observed adult traps may reflect mortality, delayed emergence, movement, sampling variability, or redistribution rather than a single population-wide response.

Thermal constraints should be assigned to specific life-cycle and pathogen components. Comparative evidence across temperate mosquito-borne viruses shows that broadly similar transmission peaks can arise from different combinations of survival, biting, development, competence, and incubation responses [11]. Temperature can therefore delimit broad transmission conditions without completely explaining local dynamics, as shown in analyses of Ross River virus transmission [12]. In architectural terms, thermal input should modify named transition functions rather than act as a generic risk multiplier. Where trait evidence is sparse, the model should propagate uncertainty and identify the transferred population or experimental context from which the parameter was obtained.

The assumption of permanently fixed species parameters is also difficult to defend. Mosquito responses to warming may reflect acclimation, phenotypic plasticity, evolutionary change, demographic filtering, or altered species composition, and the evidence available to separate these processes remains uneven [13]. At broader scales, climate- and population-dependent diffusion can reproduce and forecast regional spread of *Aedes albopictus*, supporting the inclusion of colonization and spatial dispersal processes alongside local stage dynamics [14]. Nevertheless, diffusion is a simplified representation of transport and establishment, not a direct estimate of local abundance or intervention susceptibility. The life-cycle module should consequently provide uncertain state trajectories that can be tested against stage-specific observations, not claim complete population observability.

*Weather, habitat, and human-mobility inputs*

Weather, environmental barriers, and human mobility jointly condition when and where transmission can emerge [15–17]. Comparative climate evidence shows that the direction, lag, and strength of weather–disease relationships vary among pathogens and locations, while mobility networks can couple otherwise separated transmission units and environmental barriers can constrain the routes along which emergence occurs. These inputs should therefore interact with biological states rather than enter as independent, universally weighted predictors. Temperature and hydrological observations require local temporal transformations; movement requires time-varying spatial connectivity; and environmental suitability should constrain whether introduced vectors or infections can establish and persist.

Habitat representation can be strengthened through Earth observation, but remote-sensing products remain indirect measurements. Environmental and land-surface variables have been used to predict temporal distributions of *Aedes aegypti*, demonstrating the value of multimodal habitat information [18]. The same result also defines an interpretive limit: a remotely sensed association does not establish that a particular breeding site exists, that larvae occupy it, or that an intervention can reach it.

Sensor resolution, cloud cover, changing land classifications, sparse trapping, and mismatch between pixel scale and mosquito-experienced habitat can all produce distribution shift. The twin should therefore store raw products, derived features, transformation methods, acquisition dates, and uncertainty separately from the inferred habitat state.

Local field evidence further indicates that mosquito abundance can be associated with built-environment and sociodemographic conditions, although these associations may be shaped by trapping design, unmeasured environmental factors, infrastructure, or differential surveillance [19]. Social context may help identify heterogeneous exposure and service conditions, but it should not be encoded as an intrinsic biological cause or used to justify unexamined targeting. The proposed input layer should consequently distinguish weather observations, habitat observations, habitat proxies, human-mobility estimates, demographic context, and model-derived suitability. Their spatial and temporal supports must be aligned with the state transition they are intended to inform, and their contribution should be retained only when validation shows that it improves a prespecified biological, forecasting, or decision target without concealing bias.

*Insecticide resistance and intervention response*

Insecticide resistance should not enter the digital twin as a binary or species-wide attribute. Global evidence for major *Aedes* vectors shows substantial geographical variation in resistance phenotypes, mechanisms, insecticide classes, and assay coverage [20]. Every resistance observation should therefore retain its mosquito population, collection date, insecticide, dose, exposure protocol, mortality endpoint, species composition, and uncertainty. Otherwise, differences among assays or populations may be mistaken for temporal biological change. Resistance surfaces are consequently model-based summaries of incomplete observations, not complete measurements of susceptibility throughout a control jurisdiction.

Spatiotemporal mapping can identify broad resistance trends, but regional estimates depend on interpolation among unevenly distributed bioassays [21]. The twin should display observation density and posterior uncertainty alongside any resistance layer and should permit local assays to supersede regional priors when methods are sufficiently comparable. It must also distinguish phenotypic resistance from the genetic or metabolic mechanisms that may generate it. Similar mortality outcomes can arise from different mechanisms, while operational underperformance may arise from inadequate coverage, product decay, behavioral avoidance, environmental exposure, or application failure rather than resistance alone.

Intervention response should therefore be represented as a pathway linking susceptibility, realized contact, product efficacy, behavioral response, operational coverage, and transmission consequences. Evolutionary modelling shows that the relative durability of insecticide mixtures and sequences depends on efficacy, exposure, dominance, and genetic assumptions rather than on a universally superior strategy [22]. Standard bioassays likewise do not translate directly into practical field resistance [23]. Entomological effects can inform epidemiological projections, but their implications depend on baseline transmission, vector behavior, and intervention properties [24]. Simulation of a promising strategy is thus a conditional counterfactual, not evidence of real-world control effectiveness.

The evidence dimensions and interpretive boundaries for insecticide resistance and intervention response are summarized in **Table 2**.



**Table 2.** Insecticide Resistance and Intervention Response: Data Inputs, Analytical Tasks, Biological Representations, Validation, Explainability, Generalization, and Decision Boundaries

| Analytical component or task | Input data | Model or representation | Expected output | Validation requirement | Explainability or interpretation need | Generalization risk | Decision boundary | Representative supporting reference(s) |
|---|---|---|---|---|---|---|---|---|
| **Characterize resistance** | Standardized bioassays, species, population, | Population- and product-specific resistance state | Local susceptibility estimate with uncertainty | Repeat comparable assays and | Retain assay protocol and mechanism metadata | Assay and population differences limit transfer | Technical resistance is not | [20] |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | insecticide, dose, collection date | | | test withheld observations | | | operational failure | |
| **Map resistance through space and time** | Georeferenced bioassays and intervention history | Spatiotemporal resistance surface | Regional resistance distribution and trend | Local holdout validation and uncertainty calibration | Display observation density and interpolation uncertainty | Sparse sampling can create false spatial precision | Estimated surfaces are not complete observability | [21] |
| **Simulate resistance evolution** | Genotype, dominance, efficacy, exposure, coverage | Population-genetic intervention model | Conditional resistance trajectories | Sensitivity analysis using locally plausible parameters | Identify assumptions driving strategy comparisons | Genetic and exposure conditions vary among settings | Simulated durability is not field effectiveness | [22] |
| **Separate technical and practical resistance** | Bioassays, mosquito physiological state, product and environmental conditions | Linked phenotype–exposure–performance representation | Context-specific practical response estimate | Compare assay results with field product performance | Explain discordance among assay, contact, and outcome | Age, feeding status, substrate, and environment modify response | Bioassay mortality is not practical resistance | [23] |
| **Project intervention consequences** | Mortality, feeding inhibition, behavior, coverage, baseline transmission | Entomological-to-epidemiological pathway | Conditional transmission-effect projection | Compare predictions with independent operational outcomes | Trace how each entomological effect changes transmission | Vector behavior and programme delivery may differ | Biological efficacy is not operational effectiveness | [24] |

*Surveillance assimilation and model updating*

Surveillance assimilation should update a probability distribution over plausible biological states rather than replace the model state with each new observation. Temperature-forced West Nile virus forecasting demonstrates that an additional data stream can improve selected targets while offering little or inconsistent benefit for others [25]. Input inclusion should therefore be justified by prespecified forecast or decision targets, comparator models, and ablation tests. A richer data environment does not necessarily reduce uncertainty when measurements are redundant, delayed, spatially mismatched, or weakly connected to the state being estimated.

Observation time must also be separated from event time. Mosquito and human surveillance reports can arrive after substantial and jurisdiction-specific delays, and those delays can affect forecast accuracy [26]. The assimilation layer should represent reporting latency, revisions, missingness, trap effort, diagnostic uncertainty, and changes in surveillance practice. Without such observation models, the update procedure may interpret administrative variation as biological change. Data assimilation is therefore not equivalent to unbiased updating: a mathematically coherent filter can reproduce systematic underreporting or sampling bias with increasing confidence.

Sequential filtering can update latent mosquito, pathogen, and human states, but these quantities remain conditional estimates. A real-time West Nile virus system used incoming observations to update states and parameters during an outbreak [27]. Metapopulation filtering has also propagated dengue information among connected cities, supporting spatially coupled updating when movement connects local epidemics [28]. These approaches justify a state-updating component in the proposed twin, but they do not establish complete observability. Parameter identifiability, filter stability, network specification, and surveillance representativeness must be tested separately.

Validation should examine calibration, interval coverage, peak timing, peak intensity, cumulative outcomes, and failure across outbreak scales rather than rely on one aggregate accuracy measure. An ensemble dengue forecasting system performed favourably against a comparator in its evaluated setting, but its evidence remained local and retrospective [29]. The proposed twin should consequently record model versions, priors, observations, posterior states, rejected updates, and forecast errors. Updating should be reversible when data or models are corrected, and persistent discrepancy

should trigger structural review rather than unlimited parameter adjustment.

*Proposed mosquito-control digital twin*

The proposed mosquito-control digital twin is a modular, continuously updateable representation of a defined control jurisdiction. Its core contains stage-structured mosquito states, vector infection states, spatial transmission units, and observation models. Weather and habitat modify specified biological transitions; human mobility couples local units; resistance modifies intervention susceptibility; and intervention records represent product, coverage, timing, delivery, and realized exposure. Digital-twin modelling scholarship identifies uncertainty, model discrepancy, computational constraints, and continual data integration as central design challenges [30]. Layered twin architectures further support separating the represented physical system, virtual state, data exchange, and decision services so that each can be examined independently [31].

**Figure 1** presents the digital-twin architecture within the analytical logic developed in this section.

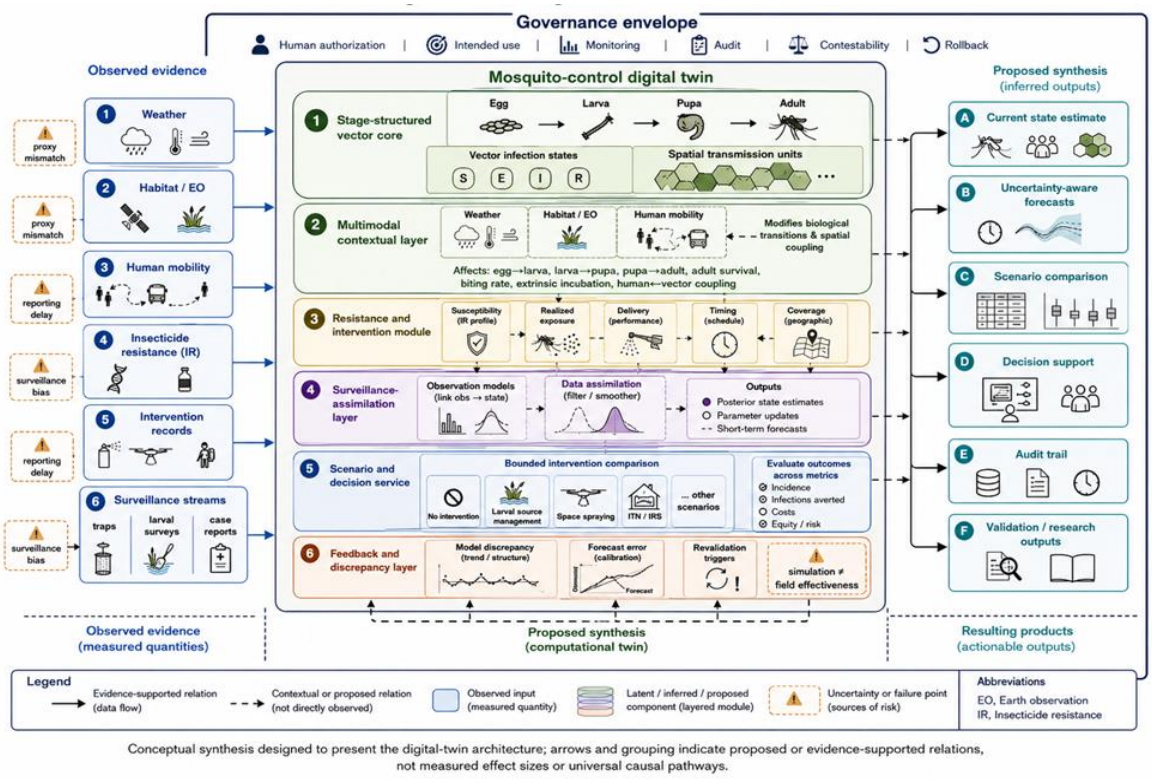


**Figure 1.** The digital-twin architecture

*Alt text*

A structured conceptual diagram that presents the digital-twin architecture, with labelled components, directional relations, contextual modifiers, uncertainty points, and a clear boundary between observed evidence and proposed synthesis.

The feedback cycle begins when surveillance updates the current state distribution and model discrepancy. Scenario simulation then compares bounded intervention alternatives, after which a human decision process may authorize, modify, postpone, or reject an action. Deployment records and subsequent observations return to the twin, but they should not trigger an automatic causal interpretation or policy update. Health digital-twin evidence shows that many systems described as twins remain prototypes with limited prospective or external validation [32]. **Figure 2** therefore represents feedback as a monitored learning process rather than an autonomous control loop.

**Figure 2** shows feedback cycles linking surveillance, simulation, intervention, and model updating within the analytical logic developed in this section.

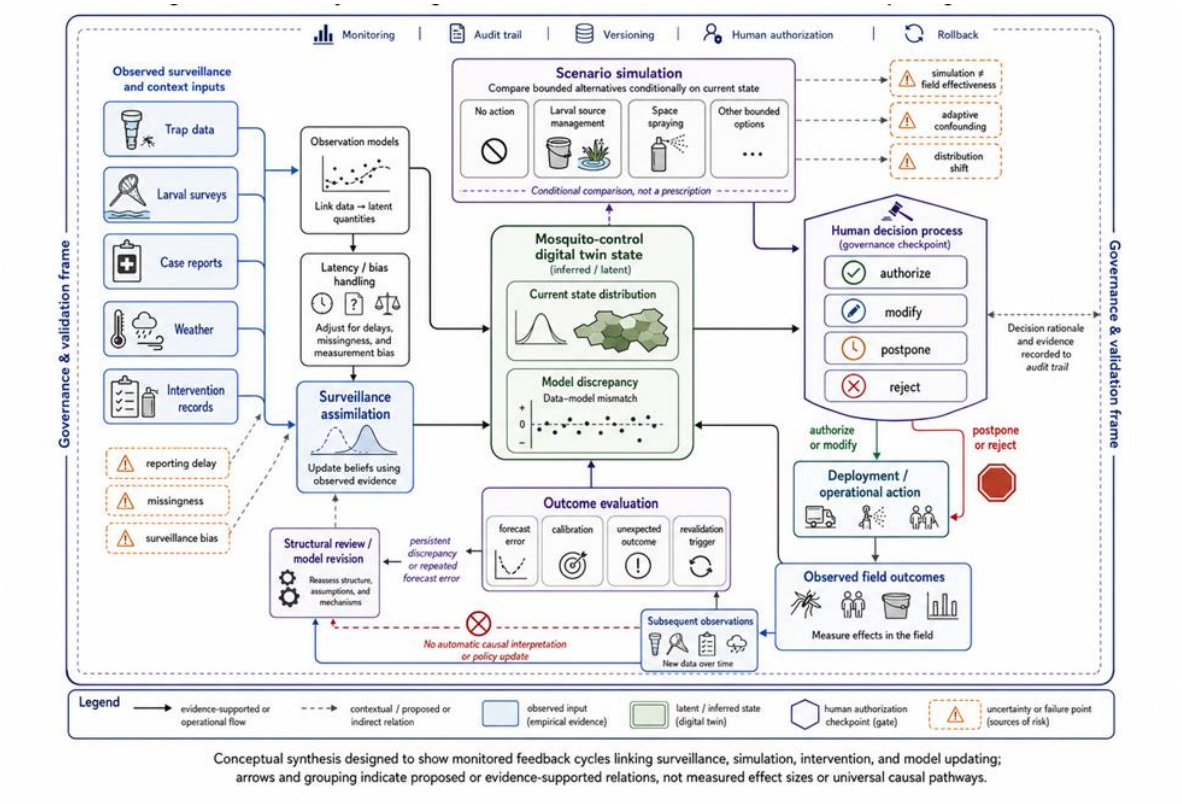


**Figure 2.** Feedback cycles linking surveillance, simulation, intervention, and model updating

*Alt text*

A structured conceptual diagram that shows feedback cycles linking surveillance, simulation, intervention, and model updating, with labelled components, directional relations, contextual modifiers, uncertainty points, and a clear boundary between observed evidence and proposed synthesis.

Multimodal representation should preserve the provenance, temporal support, spatial support, transformation, uncertainty, and validation status of each modality. Reviews of health digital twins emphasize the value of integrating heterogeneous physiological, imaging, sensor, and historical information while also exposing the difficulty of validating their connections [33]. The same principle applies here: weather, Earth observation, mobility, resistance assays, traps, cases, and intervention records are not interchangeable measurements. Their integration supports a proposed synthesis only when each stream is linked to a defined state, mechanism, observation model, or decision requirement.

The proposed components, evidence bases, boundary conditions, failure modes, and validation requirements are organized in **Table 3**.



**Table 3.** Proposed Mosquito-Control Digital Twin: Components, Evidence Basis, Relations, Boundary Conditions, Failure Modes, and Validation Requirements

| Proposed component | Purpose | Evidence basis | Relation or mechanism | Input or precondition | Expected output | Boundary condition or failure mode | Validation requirement | Representative supporting reference(s) |
|---|---|---|---|---|---|---|---|---|
| **Stage-structured vector core** | Represent development, survival, infection, and infectious-vector availability | Thermal and population evidence | Environment-conditioned transitions among latent stages | Species- and population-specific parameter sets | Uncertain stage and infection trajectories | Overparameterization or transferred traits create false precision | Stage-specific calibration and simplification tests | [5] |
| **Multimodal contextual layer** | Represent weather, habitat, mobility, and | Environmental and spatial evidence | Typed inputs modify named states | Provenance, scale, timing, and | Context-conditioned state estimate | Proxy variables may be mistaken for direct biological measurements | Ablation, drift testing, and independent | [15] |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | surveillance context | | or spatial connections | uncertainty metadata | | | field comparison | |
| **Resistance and intervention module** | Separate susceptibility, exposure, delivery, and outcome | Resistance and intervention evidence | Intervention effects propagate through contact, behavior, mortality, and transmission | Assays, product properties, coverage, timing, and delivery records | Conditional intervention-response distribution | Technical resistance or efficacy may be mistaken for operational impact | Local assay and operational outcome validation | [23] |
| **Surveillance-assimilation layer** | Update latent states and parameters | Ensemble forecasting evidence | Observation-specific sequential updating with latency and error models | Time-stamped observations and explicit priors | Posterior states, parameters, and forecasts | Biased surveillance can produce biased updating | Calibration, posterior checks, and reporting-delay audits | [26] |
| **Scenario and decision service** | Compare bounded intervention alternatives | Digital-twin modelling principles | Simulate alternatives using current uncertain state | Defined decision question and admissible actions | Conditional scenario comparison | Simulation fidelity is not control effectiveness | Prospective comparison with documented decisions and outcomes | [30] |
| **Feedback and discrepancy layer** | Learn from surveillance and intervention outcomes | Twin lifecycle and governance principles | Log prediction, decision, action, outcome, and model revision | Versioned audit trail and update authority | Model discrepancy and revalidation trigger | Adaptive targeting can confound feedback | Counterfactual or quasi-experimental evaluation | [32] |
| **Governance envelope** | Bound use and responsibility | Responsible high-stakes system principles | Human authorization, monitoring, audit, contestability, and rollback | Defined intended use, roles, thresholds, and incident process | Traceable decision support | Architecture may be mistaken for deployment approval | Independent review and prospective operational validation | [34] |



*Validation, governance, and decision implications*

Validation must proceed from data and component models to interfaces, forecasts, recommendations, and real-world outcomes. Artificial-intelligence experience in high-stakes settings shows that technical performance does not establish workflow-integrated benefit [35]. For the proposed twin, simulation fidelity should therefore be evaluated separately from entomological accuracy, forecast calibration, decision usefulness, operational feasibility, and control effectiveness. Progress would be demonstrated by prospective evaluation against explicit comparators across different seasons, vector populations, surveillance systems, and programme conditions—not by improved fit to the data used for development.

Transferability must also be tested rather than presumed. Models can fail when populations, prevalence, measurements, interventions, or data-generating processes differ from their development settings [36]. Explanation interfaces cannot solve this problem because plausible post-hoc explanations may remain unreliable and do not establish that a model is biologically correct, causally valid, fair, or safe [37]. Explanation should instead support interrogation: users should be able to identify the data, assumptions, mechanisms, uncertainty, and boundary conditions responsible for a recommendation. External validation, local recalibration, sensitivity analysis, and explicit abstention should take precedence over claims of universal generalizability.

Governance should define intended use, authorized users, decision authority, monitoring, audit, incident response, correction, contestability, and rollback. Responsible high-stakes machine learning requires lifecycle oversight rather than one-time technical approval [34]. A mosquito-control twin should therefore preserve immutable versions of data, models, assumptions, scenarios, recommendations, and human decisions. Automated intervention authority is not justified by the architecture proposed here. Its

appropriate role is bounded decision support under accountable human and institutional control, with revalidation after material changes in data, vector ecology, products, surveillance practice, or programme objectives.

The evidence gaps, their consequences, and the corresponding priorities are summarized in **Table 4**.

**Table 4.** Validation, Governance, and Decision Implications: Evidence Gaps, Consequences, Priority Actions, Validation Needs, and Decision Relevance

| Evidence or implementation gap | Why it matters | Priority action | Required study or capability | Validation indicator | Context or equity consideration | Residual limitation | Decision relevance | Representative supporting reference(s) |
|---|---|---|---|---|---|---|---|---|
| **Limited prospective and external validation** | Retrospective local performance may not transfer | Test the architecture across sites, seasons, and surveillance regimes | Prospective multi-jurisdiction evaluation with comparators | Calibration, coverage, and target-specific performance outside development data | Include settings with different resources and surveillance density | No finite test establishes universal transferability | Determines whether outputs may inform local planning | [35] |
| **Incomplete observability of mosquito states** | Traps and cases measure only parts of the system | Validate observation models and latent-state uncertainty | Stage-specific surveillance and posterior predictive assessment | Reliable uncertainty coverage for withheld observations | Avoid deprioritizing areas with sparse surveillance | Latent states remain model dependent | Bounds confidence in current-state estimates | [27] |
| **Biased or delayed surveillance assimilation** | Updating can reproduce systematic measurement error | Model latency, missingness, revisions, and sampling effort | Reporting-delay audits and synthetic-bias stress tests | Stable calibration under plausible observation bias | Assess unequal reporting and device coverage | Unknown biases may remain | Determines whether updating should be accepted or suspended | [26] |
| **Weak linkage between resistance assays and field response** | Technical resistance may be mistaken for programme failure | Collect assay, delivery, exposure, and outcome evidence together | Integrated entomological and operational evaluation | Agreement or explained discordance among assay and field outcomes | Ensure local populations and intervention conditions are represented | Epidemiological impact remains context dependent | Supports product selection and resistance management | [23] |
| **Distribution shift in environmental and mobility inputs** | Relationships may change after climate, land-use, or behavioral change | Implement drift detection and scheduled revalidation | Versioned data pipelines and out-of-distribution testing | Detected drift linked to recalibration or abstention | Mobility and sensor data may underrepresent disadvantaged groups | Drift detection does not identify every causal change | Prevents silent reuse of obsolete relations | [36] |
| **Explanation without demonstrated validity** | Plausible displays can produce unjustified trust | Link explanations to provenance, uncertainty, and validation evidence | User-centred assurance and error-detection studies | Users identify limitations and avoid unsupported causal conclusions | Interfaces must suit different expertise and authority | Explanation cannot certify correctness | Supports critical review rather than automation bias | [37] |
| **Undefined authority and rollback** | Errors can become operational actions without accountability | Establish governance before deployment testing | Audit trail, authorization rules, incident response, and rollback exercises | Decisions and model changes remain traceable and reversible | Include affected communities and operational staff in oversight | Governance reduces but does not eliminate harm | Defines whether and how recommendations can be used | [34] |

## CONCLUSION

A mosquito-control digital twin is scientifically defensible only as a dynamic, modular, and uncertainty-aware representation that connects vector life cycles, environmental and habitat conditions, human mobility, resistance, intervention response, surveillance assimilation, state updating, and feedback. The evidence supports many of these components individually and shows why fixed, isolated, or context-free representations are inadequate. It does not demonstrate that their integration will automatically improve mosquito control. The proposed architecture must therefore preserve distinctions between observation and inference, simulation and effectiveness, updating and unbiased learning, and architectural completeness and operational validation. The highest priority is prospective, multi-context evaluation under explicit governance, with component-level validation, data and parameter provenance, human authorization, monitoring, abstention, and rollback. Until such evidence exists, the architecture should be treated as a structured research and decision-support hypothesis rather than a deployment-ready control system.

**ACKNOWLEDGMENTS:** None

**CONFLICT OF INTEREST:** None

**FINANCIAL SUPPORT:** None

**ETHICS STATEMENT:** None

## REFERENCES

1. Barricelli BR, Casiraghi E, Fogli D. A survey on digital twin: Definitions, characteristics, applications, and design implications. IEEE Access. 2019;7:167653-71. doi:10.1109/ACCESS.2019.2953499
2. Blair GS. Digital twins of the natural environment. Patterns (N Y). 2021;2(10):100359. doi:10.1016/j.patter.2021.100359
3. DeFelice NB, Little E, Campbell SR, Shaman J. Ensemble forecast of human West Nile virus cases and mosquito infection rates. Nat Commun. 2017;8(1):14592. doi:10.1038/ncomms14592
4. Ryan SJ, Carlson CJ, Mordecai EA, Johnson LR. Global expansion and redistribution of Aedes-borne virus transmission risk with climate change. PLoS Negl Trop Dis. 2019;13(3). doi:10.1371/journal.pntd.0007213
5. Mordecai EA, Caldwell JM, Grossman MK, Lippi CA, Johnson LR, Neira M, et al. Thermal biology of mosquito-borne disease. Ecol Lett. 2019;22(10):1690-708. doi:10.1111/ele.13335
6. Tesla B, Demakovsky LR, Mordecai EA, Ryan SJ, Bonds MH, Ngonghala CN, et al. Temperature drives Zika virus transmission: Evidence from empirical and mathematical models. Proc Biol Sci. 2018;285(1884):20180795. doi:10.1098/rspb.2018.0795
7. Paull SH, Horton DE, Ashfaq M, Rastogi D, Kramer LD, Diffenbaugh NS, et al. Drought and immunity determine the intensity of West Nile virus epidemics and climate change impacts. Proc Biol Sci. 2017;284(1848):20162078. doi:10.1098/rspb.2016.2078
8. de Wit MM, Martins AD, Delecroix C, Heesterbeek H, ten Bosch QA. Mechanistic models for West Nile virus transmission: A systematic review of features, aims and parametrization. Proc Biol Sci. 2024;291(2018):20232432. doi:10.1098/rspb.2023.2432
9. Brass DP, Cobbold CA, Purse BV, Ewing DA, Callaghan A, White SM. Role of vector phenotypic plasticity in disease transmission as illustrated by the spread of dengue virus by Aedes albopictus. Nat Commun. 2024;15(1):7823. doi:10.1038/s41467-024-52144-5
10. Iwamura T, Guzman-Holst A, Murray KA. Accelerating invasion potential of disease vector Aedes aegypti under climate change. Nat Commun. 2020;11(1):2130. doi:10.1038/s41467-020-16010-4
11. Shocket MS, Verwillow AB, Numazu MG, Slamani H, Cohen JM, El Moustaid F, et al. Transmission of West Nile and five other temperate mosquito-borne viruses peaks at temperatures between 23°C and 26°C. Elife. 2020;9. doi:10.7554/eLife.58511
12. Shocket MS, Ryan SJ, Mordecai EA. Temperature explains broad patterns of

Ross River virus transmission. Elife. 2018;7. doi:10.7554/eLife.37762
13. Couper LI, Farner JE, Caldwell JM, Childs ML, Harris MJ, Kirk DG, et al. How will mosquitoes adapt to climate warming? Elife. 2021;10. doi:10.7554/eLife.69630
14. Barman S, Semenza JC, Singh P, Sjödin H, Rocklöv J, Wallin J. A climate and population dependent diffusion model forecasts the spread of Aedes albopictus mosquitoes in Europe. Commun Earth Environ. 2025;6(1):276. doi:10.1038/s43247-025-02199-z
15. Caldwell JM, LaBeaud AD, Lambin EF, Stewart-Ibarra AM, Ndenga BA, Mutuku FM, et al. Climate predicts geographic and temporal variation in mosquito-borne disease dynamics on two continents. Nat Commun. 2021;12(1):1233. doi:10.1038/s41467-021-21496-7
16. Massaro E, Kondor D, Ratti C. Assessing the interplay between human mobility and mosquito borne diseases in urban environments. Sci Rep. 2019;9(1):16911. doi:10.1038/s41598-019-53127-z
17. Harish V, Colón-González FJ, Moreira FRR, Gibb R, Kraemer MUG, Davis M, et al. Human movement and environmental barriers shape the emergence of dengue. Nat Commun. 2024;15(1):4205. doi:10.1038/s41467-024-48465-0
18. Mudele O, Bayer FM, Zanandrez LFR, Eiras AE, Gamba P. Modeling the temporal population distribution of Ae. aegypti mosquito using big Earth observation data. IEEE Access. 2020;8:14182-94. doi:10.1109/ACCESS.2020.2966080
19. Moise IK, Riegel C, Muturi EJ. Environmental and social-demographic predictors of the southern house mosquito Culex quinquefasciatus in New Orleans, Louisiana. Parasit Vectors. 2018;11(1):249. doi:10.1186/s13071-018-2833-5
20. Moyes CL, Vontas J, Martins AJ, Ng LC, Koou SY, Dusfour I, et al. Contemporary status of insecticide resistance in the major Aedes vectors of arboviruses infecting humans. PLoS Negl Trop Dis. 2017;11(7). doi:10.1371/journal.pntd.0005625
21. Hancock PA, Hendriks CJM, Tangena JA, Gibson H, Hemingway J, Coleman M, et al. Mapping trends in insecticide resistance phenotypes in African malaria vectors. PLoS Biol. 2020;18(6). doi:10.1371/journal.pbio.3000633
22. Levick B, South A, Hastings IM. A two-locus model of the evolution of insecticide resistance to inform and optimise public health insecticide deployment strategies. PLoS Comput Biol. 2017;13(1). doi:10.1371/journal.pcbi.1005327
23. Namias A, Jobe NB, Paaijmans KP, Huijben S. The need for practical insecticide-resistance guidelines to effectively inform mosquito-borne disease control programs. Elife. 2021;10. doi:10.7554/eLife.65655
24. Sherrard-Smith E, Ngufor C, Sanou A, Guelbeogo MW, N'Guessan R, Elobolobo E, et al. Inferring the epidemiological benefit of indoor vector control interventions against malaria from mosquito data. Nat Commun. 2022;13(1):3862. doi:10.1038/s41467-022-30700-1
25. DeFelice NB, Schneider ZD, Little E, Barker C, Caillouet KA, Campbell SR, et al. Use of temperature to improve West Nile virus forecasts. PLoS Comput Biol. 2018;14(3). doi:10.1371/journal.pcbi.1006047
26. DeFelice NB, Birger R, DeFelice N, Gagner A, Campbell SR, Romano C, et al. Modeling and surveillance of reporting delays of mosquitoes and humans infected with West Nile virus and associations with accuracy of West Nile virus forecasts. JAMA Netw Open. 2019;2(4). doi:10.1001/jamanetworkopen.2019.3175
27. Yi C, Cohnstaedt LW, Scoglio CM. A real-time forecasting and estimating system of West Nile virus: A case study of the 2023 WNV outbreak in Colorado, USA. R Soc Open Sci. 2024;11(12):240513. doi:10.1098/rsos.240513
28. Zeng Q, Yu X, Ni H, Xiao L, Xu T, Wu H, et al. Dengue transmission dynamics prediction by combining metapopulation networks and Kalman filter algorithm. PLoS Negl Trop Dis. 2023;17(6). doi:10.1371/journal.pntd.0011418
29. Chen Y, Liu T, Yu X, Zeng Q, Cai Z, Wu H, et al. An ensemble forecast system for tracking dynamics of dengue outbreaks and its validation in China. PLoS Comput Biol. 2022;18(6). doi:10.1371/journal.pcbi.1010218

30. Rasheed A, San O, Kvamsdal T. Digital twin: Values, challenges and enablers from a modeling perspective. IEEE Access. 2020;8:21980-2012. doi:10.1109/ACCESS.2020.2970143
31. Liu M, Fang S, Dong H, Xu C. Review of digital twin about concepts, technologies, and industrial applications. J Manuf Syst. 2021;58(Pt B):346-61. doi:10.1016/j.jmsy.2020.06.017
32. Katsoulakis E, Wang Q, Wu H, Shahriyari L, Fletcher R, Liu J, et al. Digital twins for health: A scoping review. NPJ Digit Med. 2024;7(1):77. doi:10.1038/s41746-024-01073-0
33. Coorey G, Figtree GA, Fletcher DF, Snelson VJ, Vernon ST, Winlaw D, et al. The health digital twin to tackle cardiovascular disease—a review of an emerging interdisciplinary field. NPJ Digit Med. 2022;5(1):126. doi:10.1038/s41746-022-00640-7
34. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: A roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337-40. doi:10.1038/s41591-019-0548-6
35. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17(1):195. doi:10.1186/s12916-019-1426-2
36. Futoma J, Simons M, Panch T, Doshi-Velez F, Celi LA. The myth of generalisability in clinical research and machine learning in health care. Lancet Digit Health. 2020;2(9). doi:10.1016/S2589-7500(20)30186-2
37. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9