Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

CAN EXPLANATIONS SURVIVE BIOLOGICAL COMPLEXITY? A SYSTEMATIC MAP OF EXPLAINABLE ARTIFICIAL INTELLIGENCE IN OMICS-DRIVEN DRUG DISCOVERY

Koffi Mensah[1]*, Hawa Ouedraogo[2], Ismael Diallo[3], Stefan Becker[4], Mahamadou Traoré[5]

1. ***Department of Explainable AI in Omics-Driven Discovery, Faculty of Pharmacy, University of Ghana, Accra, Ghana.***
2. ***Department of Biological Complexity and Explanation Validity, Faculty of Pharmacy, University of Ouagadougou, Ouagadougou, Burkina Faso.***
3. ***Department of Systematic Mapping of XAI in Drug Discovery, Faculty of Pharmaceutical Sciences, University of Niger, Niamey, Niger.***
4. ***Department of Explanation Survival and Biological Context, Faculty of Veterinary Medicine, University of Zurich, Zurich, Switzerland.***
5. ***Department of Omics-Driven Drug Discovery and Explainability, Faculty of Pharmacy, University of Abomey-Calavi, Cotonou, Benin.***

ARTICLE INFO

Received:
01 March 2025
Received in revised form:
15 June 2025
Accepted:
17 June 2025
Available online:
28 June 2025

Keywords: Explainable artificial intelligence, Omics data, Drug discovery, Biological interpretability, Explanation faithfulness, Context dependence

ABSTRACT

Explainable artificial intelligence is increasingly used to interpret genomic, transcriptomic, and multi-omics models in drug discovery, yet the biological meaning of those explanations remains uncertain. High-dimensional dependence, pathway redundancy, incomplete prior knowledge, cohort and platform shifts, and context-specific regulation can make an explanation appear coherent while remaining unstable, model-contingent, or biologically circular. This systematic mapping review examines how explanation methods are defined, applied, and evaluated across omics-driven pharmaceutical research. A protocol-led mapping strategy was used to organize the selected peer-reviewed literature by omics layer, drug-discovery task, model architecture, explanation family, biological object, validation practice, expert involvement, and downstream decision relevance. The synthesis indicates that technical method availability is more mature than evidence of biological credibility. Feature attribution, perturbation analysis, rule extraction, attention-based interpretation, and biologically structured neural networks can reveal model-relevant patterns, but benchmark performance, pathway familiarity, and expert agreement do not alone establish mechanism, causality, or translational value. The review proposes an evidence hierarchy in which computational faithfulness, stability, context transportability, non-circular biological coherence, experimental confirmation, and decision utility are treated as distinct requirements. It also identifies major neglected areas, including cross-platform explanation replication, prospective experimental use, negative evidence, uncertainty communication, and controlled studies of expert decision impact. The central contribution is an evaluative map of where explanations are technically capable, scientifically plausible, weakly validated, or not yet ready to influence pharmaceutical decisions. The proposed standards are review-derived rather than empirically validated and require adaptation to specific biological systems, development stages, and decision consequences.

To Cite This Article: Mensah K, Ouedraogo H, Diallo I, Becker S, Traoré M. Can Explanations Survive Biological Complexity? A Systematic Map of Explainable Artificial Intelligence in Omics-Driven Drug Discovery. Pharmacophore. 2025;16(3):86-95. https://doi.org/10.51847/4hFpoC4b6j

Introduction

Artificial intelligence is increasingly used to extract predictive structure from genomic sequences, expression profiles, molecular networks, multi-omics measurements, and pharmacological response data. These applications have intensified interest in explanations that can expose which variants, genes, transcripts, pathways, cell states, or cross-modal relationships influence a model prediction. Recent scholarship positions explanation as a possible interface between opaque computational

Corresponding Author: Koffi Mensah; Department of Explainable AI in Omics-Driven Discovery, Faculty of Pharmacy, University of Ghana, Accra, Ghana. E-mail: koffi.mensah@ug.edu.gh.

models and scientific or high-stakes interpretation, while warning that domain adaptation and validity testing remain essential [1–3]. In drug discovery, this interface is especially consequential because an apparently meaningful explanation may influence target prioritization, biomarker selection, experimental design, patient stratification, or interpretation of drug-response heterogeneity. The scientific value of such an explanation therefore depends on more than whether it is visually persuasive or intelligible to a user.

The term *explanation* nevertheless covers fundamentally different objects. It may describe a feature-attribution map, a local approximation of model behaviour, a rule extracted from a predictive system, a counterfactual change, a perturbational sensitivity profile, an attention pattern, or the internal organization of a biologically structured model. These objects differ in whether they are intrinsic or post hoc, local or global, model-specific or model-agnostic, and descriptive or intervention-oriented. For consequential applications, post-hoc explanation should not automatically be treated as equivalent to using a model whose reasoning is directly inspectable [4]. Nor does an interpretable architecture eliminate the possibility that the model has learned confounded, incomplete, or context-limited relationships. A model can be transparent about an invalid association, just as a black-box model can sometimes encode a useful signal that remains difficult to interpret.

Omics data magnify this problem because biological variables are neither independent nor stable across contexts. Genes participate in overlapping pathways; transcript abundance depends on cell composition, disease stage, treatment exposure, time, and measurement platform; and multi-omics integration introduces mismatched scales, missing modalities, batch structure, and unequal information density. Explanations may therefore change because a method is unstable, because the model is poorly specified, or because the underlying biology legitimately differs between populations, tissues, laboratories, or perturbational states. Conversely, an explanation may appear stable because the model repeatedly reproduces the same confounding structure or because a prior-knowledge architecture constrains it toward familiar pathways. Stability, familiarity, and biological plausibility must consequently be evaluated as distinct properties rather than collapsed into a single claim of interpretability.

This systematic mapping review asks whether explanations can remain scientifically defensible as they encounter biological complexity. Its purpose is not to estimate a pooled effect or declare one explanation method superior across all omics tasks. Instead, it maps explanation families, application contexts, evaluation practices, maturity levels, and neglected validation requirements across omics-driven drug discovery. The central contribution is an evaluative synthesis that separates demonstrated computational capability from biological plausibility, model faithfulness, experimental confirmation, downstream utility, and routine pharmaceutical readiness. The review further proposes a biologically oriented credibility standard and an omics explanation fragility map. Both are conceptual products of the synthesis rather than empirically validated qualification systems.

Systematic-Mapping Questions and Protocol

The review was structured as a systematic evidence map because the relevant literature is methodologically heterogeneous and is not organized around a common intervention, comparator, or outcome that would support quantitative pooling. The protocol predefined seven linked questions: how the map should be bounded; how evidence should be searched, screened, coded, and charted; which explanation methods are used across genomics, transcriptomics, and multi-omics; how biological complexity and context affect explanation stability; how faithfulness, expert interpretation, and downstream utility are evaluated; which areas are comparatively mature, weak, or neglected; and which standards are required for biologically credible evaluation. Transparent reporting was treated as a condition for review auditability, including explicit objectives, eligibility concepts, information sources, selection logic, synthesis methods, and limitations [5]. The research questions were therefore formulated to identify distributions and evidence relationships rather than calculate pooled estimates.

Protocol decisions distinguished the *prediction object*, the *explanation object*, and the *biological claim*. The prediction object could be a drug response, disease class, molecular phenotype, regulatory activity, target association, or another pharmaceutical outcome. The explanation object denoted what the method actually exposed, such as an input feature, sequence element, gene module, pathway node, latent factor, cell state, or cross-modal relationship. The biological claim described the interpretation attached to that output, ranging from model relevance and association to putative mechanism or causal influence. Eligibility and coding rules required these layers to remain separate because a method can faithfully explain a model without validating the biological proposition inferred from that explanation. Protocol-level guidance was used to make the eligibility logic, selection decisions, synthesis choices, and departures from planned procedures explicit rather than leaving them implicit in the narrative [6].

The protocol also distinguished evidence demonstrating technical availability from evidence supporting scientific or translational credibility. Technical demonstrations included the production of coherent feature rankings, local explanations, perturbation profiles, rules, attention maps, or biologically labelled internal representations. Stronger evidence required tests of model faithfulness, retraining stability, resistance to leakage, replication across cohorts or platforms, independent biological validation, expert task performance, or prospective decision impact. Search reproducibility was supported by predefined concept blocks, documented database interfaces, complete Boolean strategies, eligibility limits, supplementary citation searches, and DOI-level bibliographic verification [7]. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop systematic-mapping questions and protocol within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Systematic-mapping questions and protocol

Evidence domain	Eligible evidence or method	Reported capability or claim	Strength of support	Reproducibility or bias concern	Unresolved gap
Systematic-review governance	Protocols, reporting guidelines, eligibility frameworks, search documentation, selection records, and structured synthesis methods	Enables a transparent and auditable account of how the evidence map was constructed	Strong methodological support for reporting and traceability	Complete reporting does not by itself eliminate publication bias, search omissions, or subjective coding	Limited consensus on methodological standards tailored specifically to systematic maps of omics explainability
Omics and pharmaceutical scope	Genomic, transcriptomic, single-cell, pharmacogenomic, and multi-omics studies linked to drug response, target discovery, molecular phenotyping, or related development decisions	Defines the biological and pharmaceutical contexts in which explanations are being applied	Direct support where omics data and pharmaceutical tasks are central to the article	Broad use of "drug discovery" may combine fundamentally different tasks, evidence levels, and decision consequences	Insufficient agreement on when an omics prediction has a sufficiently direct pharmaceutical connection
Explanation taxonomy	Attribution, perturbation, surrogate, rule-based, attention, counterfactual, sparse, knowledge-primed, and biologically structured approaches	Reveals model-relevant features, relationships, internal units, or hypothetical changes	Strong support for method availability; variable support for biological interpretation	Terminological inconsistency may cause intrinsically interpretable models and post-hoc explainers to be treated as equivalent	Lack of a shared taxonomy linking explanation form to biological object and intended scientific use
Biological credibility	Stability analyses, external validation, pathway interpretation, replication, perturbational evidence, and independent experimental confirmation	Suggests that an explanation may remain meaningful beyond the original model and dataset	Uneven; strongest for bounded technical or experimental demonstrations	Correlated features, circular use of prior knowledge, batch effects, population structure, and context shifts can generate misleading coherence	Few studies evaluate several biological credibility requirements together
Human and expert evaluation	Expert review, user studies, task-based assessment, trust measurement, error detection, and decision comparisons	Examines whether an explanation is understandable or useful to an intended user	Developing evidence, with limited direct pharmaceutical evaluation	Familiarity and plausibility may increase confidence without improving decision quality	Sparse controlled evidence involving pharmacologists, molecular biologists, medicinal chemists, or translational teams
Downstream utility and readiness	Explanation-guided target prioritization, biomarker selection, experimental design, assay selection, or drug-response investigation	Proposes that explanations can influence scientific or development decisions	Predominantly plausible or retrospective support	Benchmark success may be mistaken for experimental, translational, or operational usefulness	Limited prospective evidence that explanations improve decisions, reduce errors, or allocate experimental resources more effectively

Search, Screening, Coding, and Evidence-Charting Methods

The search strategy combined concepts for explainable or interpretable artificial intelligence with terms for genomics, transcriptomics, proteomics, metabolomics, single-cell analysis, pharmacogenomics, and multi-omics, together with terms representing drug discovery, drug response, target identification, molecular mechanism, therapeutic development, and pharmaceutical decision-making. Supplementary concepts captured explanation faithfulness, robustness, stability, context dependence, external validation, expert evaluation, reproducibility, and downstream utility. Searches were designed for major biomedical and multidisciplinary bibliographic databases, with citation chasing used to identify methodologically or biologically relevant journal articles not recovered through the principal strings. Database names, interfaces, search expressions, limits, and search updates were retained in the review record. The mapping design followed scoping-review principles by charting the extent, characteristics, and gaps of a heterogeneous evidence base rather than estimating pooled intervention effects [8].

Eligibility screening was conducted against predefined article-level criteria. Eligible records had to be peer-reviewed journal articles with sufficient methodological or conceptual detail to support at least one mapping question. Studies were retained when omics data, explanation methods, biological interpretation, validation design, expert evaluation, reproducibility, or pharmaceutical translation formed a substantive part of the article rather than an incidental mention. Exclusion categories covered non-journal materials, records without verifiable bibliographic identity, studies outside the intended scientific scope, narrow benchmarks lacking interpretable validity evidence, and articles whose claims could not be linked to a precise manuscript use. Full-text decisions were assigned one principal exclusion reason to avoid retrospective rationalization. Methodological appraisal focused on critical weaknesses—including inappropriate validation, unclear data provenance, leakage risk, circular biological confirmation, unsupported causal language, and absence of relevant external testing—rather than reducing heterogeneous articles to a single quality score [9].

Evidence charting used a structured coding dictionary that separated article type, omics layer, biological scale, pharmaceutical task, prediction target, model family, explanation family, intrinsic or post-hoc status, local or global scope, biological object, dependence on prior knowledge, faithfulness testing, stability assessment, external validation, experimental confirmation, expert involvement, downstream decision task, uncertainty reporting, and reproducibility resources. "Not clearly reported"

was recorded when an article did not provide sufficient information for a category. Each publication was assigned to an evidence-maturity interpretation based on the strongest form of support it actually contained, while weaker and stronger claims within the same article were charted separately. Because the evidence was too heterogeneous for meaningful pooling, synthesis proceeded through structured comparison of convergent findings, divergent findings, contextual explanations, reproducibility concerns, and boundaries on what could be concluded [10]. Distributional statements were reserved for patterns supported by the verified corpus, and no absent numerical detail was reconstructed from narrative descriptions.

Connecting Release Profiles to Local and Systemic Exposure

The transition from depot release to pharmacologically relevant exposure begins in the tissue surrounding the formulation rather than in the systemic circulation. Once a drug or prodrug leaves a depot, it may remain associated with carrier material, bind to extracellular components, precipitate, dissolve incompletely, enter lymphatic structures, undergo enzymatic transformation, or diffuse toward local blood vessels. Computational and experimental analyses of long-acting antiretroviral formulations have shown that lymphatic sequestration and the balance between carrier-associated and dissociated drug species can contribute materially to prolonged pharmacokinetics [11]. These processes demonstrate why depot output cannot be interpreted as a direct input into a conventional plasma pharmacokinetic model without an intervening transport description. In the proposed Release–Exposure–Response Translation Framework, depot exit is therefore followed by a distinct local disposition layer that represents drug state, location, transport route, and transformation before either systemic entry or access to a local pharmacological target.

Local disposition is especially important after subcutaneous administration, where tissue architecture, extracellular matrix properties, injection volume, formulation viscosity, binding, vascular density, and lymphatic drainage may shape absorption. Multiscale pharmacokinetic modeling has been used to connect local transport and absorption processes with systemic exposure following subcutaneous administration of biotherapeutics [12]. Such models illustrate the need to bridge spatial and temporal scales: microscopic diffusion and binding can influence macroscopic plasma profiles, while systemic distribution and clearance can determine whether local absorption remains rate-limiting. RERTF does not prescribe one mathematical representation for this bridge. Compartmental, physiologically based, partial differential equation, agent-based, and hybrid models may each be suitable under different conditions. The critical requirement is that the model explicitly states which species and transport processes it represents, which mechanisms are inferred rather than measured, and which outputs have been independently evaluated.

Multiphysics simulations further indicate that local diffusion, convection, tissue deformation, and vascular or lymphatic absorption can be coupled with whole-body pharmacokinetics to explain systemic exposure after injection [13]. This coupling is important because a release-rate parameter may otherwise absorb several unobserved mechanisms and appear more transferable than it is. A fitted slow-release constant could represent true matrix-controlled liberation, poor dissolution, tissue sequestration, limited vascular access, reversible binding, or a combination of these processes. Mechanistic pharmacokinetic models of long-acting injectable formulations accordingly seek to connect drug physicochemical properties and depot release rates with desired systemic profiles [14]. Their usefulness depends on parameter identifiability and on whether the modeled release term corresponds to an experimentally defensible process. A model may reproduce plasma data while assigning the wrong mechanism if compensating parameters are available. Predictive agreement is therefore necessary for many applications but cannot independently establish mechanistic correctness.

Physiologically based pharmacokinetic modeling of intramuscular antiretroviral drugs demonstrates the additional importance of muscle physiology, local blood flow, formulation behavior, drug properties, and systemic disposition [15]. The relevant exposure may include parent drug, active metabolite, prodrug, carrier-associated species, or local tissue concentrations, depending on the therapeutic mechanism. For locally acting systems, plasma concentrations may provide an incomplete or even misleading representation of pharmacological opportunity. For systemic therapies, local accumulation may still determine tolerability, absorption variability, and residual drug after treatment discontinuation. RERTF therefore requires the exposure layer to specify the chemical species, anatomical compartment, temporal interval, and therapeutic purpose under consideration. The framework does not assume that local and systemic exposure must always be measured directly, but any surrogate or computational estimate must be bounded by its assumptions and evaluated against evidence appropriate to the intended claim.

Responsive and Long-Acting Systems under Changing Physiological Conditions

Responsive delivery systems add a further condition to the release–exposure relationship: drug liberation or carrier transformation depends on a stimulus that may vary across space, time, patients, and disease states. Stimuli-responsive nanocarriers can be designed to react to pH, redox conditions, enzymes, temperature, light, magnetic fields, ultrasound, or combinations of endogenous and exogenous triggers [16]. Material responsiveness is commonly demonstrated by comparing release under two or more controlled laboratory conditions. Such evidence establishes that a formulation can respond to the tested stimulus range, but it does not establish that the trigger is sufficiently available, specific, persistent, or spatially localized in vivo. RERTF therefore treats responsiveness as a relationship among the material, the stimulus trajectory, the biological environment, and the intended pharmacological effect. It is not regarded as a fixed property of the formulation independent of context.

Glucose-responsive insulin delivery illustrates this distinction. A useful system must not merely release more insulin at higher glucose concentrations; it must sense relevant changes, respond within an appropriate time, avoid excessive release, limit

delayed action after glucose normalization, preserve insulin activity, and function repeatedly within a physiologically variable environment [17]. These requirements introduce response lag, gain, threshold, reversibility, exhaustion, and hysteresis into pharmaceutical evaluation. The same nominal glucose concentration may not produce the same delivery response if the carrier has aged, partially depleted, swollen, become protein-coated, or undergone repeated activation. Therapeutic control therefore depends on a dynamic feedback relationship rather than a static release ratio measured under fixed conditions. Any computational model of responsive delivery should distinguish stimulus prediction from drug-release prediction and both from prediction of the resulting biological response.

Route-specific physiology can further determine whether an apparently responsive material produces useful exposure. Nasal in situ gelling systems, for example, depend on gelation, mucus interaction, mucociliary clearance, residence time, epithelial transport, enzymatic stability, and anatomical access to relevant absorption pathways [18]. Enzyme-responsive systems face a related problem because disease-associated protease activity may be heterogeneous within and between tissues and may change during treatment [19]. A carrier can therefore be chemically responsive yet pharmacologically unreliable if the trigger is inaccessible, transient, insufficiently specific, or present in non-target tissues. The proposed framework separates trigger detection, carrier transformation, drug liberation, transport, and biological response as distinct transitions. Evidence at one transition should not be used automatically to claim success at the next.

Dual redox-responsive systems demonstrate that multiple triggers can improve selectivity in theory while also introducing additional failure modes [20]. Sequential or combined activation may depend on the ordering, magnitude, and duration of extracellular and intracellular conditions. Competing triggers may produce partial activation, premature release, or failure to activate. For long-acting systems, the physiological environment can also change over weeks or months because of inflammation, disease progression, treatment response, tissue remodeling, or repeated administration. RERTF consequently conceptualizes responsive and long-acting delivery as a dynamic state space rather than a single input–output curve. The formulation may move among dormant, hydrated, activated, depleted, degraded, encapsulated, and structurally altered states, each with different transport and response properties. This theoretical representation does not imply autonomous therapeutic regulation or validated control performance; it identifies the variables that future evaluation must resolve.

Evaluation Criteria beyond Cumulative Release Curves

Evaluation should begin by asking whether the release method measures the intended formulation process. For nanocarriers and related systems, apparent release can be influenced by separation technique, membrane transport, maintenance of sink conditions, sampling procedures, drug instability, and incomplete recovery [21]. Similar concerns apply to accelerated testing, where elevated temperature, altered pH, surfactants, organic cosolvents, or mechanical agitation may change the governing mechanism rather than merely increase its rate. A suitable method should therefore be reproducible, mass-balanced, discriminatory, and mechanistically interpretable within a defined context of use. Biorelevance is not an intrinsic property of an apparatus; it is a claim that must be justified by showing that the measured process corresponds sufficiently to the in vivo mechanism relevant to the proposed decision.

Evaluation of complex long-acting injectables also requires comparison of formulation composition, manufacturing history, particle or depot structure, release mechanism, pharmacokinetics, and local tolerability [22]. Release-curve similarity may contribute to product understanding, but it cannot independently resolve whether two systems form comparable depots, generate the same active species, produce similar local exposure, or interact similarly with tissue. The proposed framework therefore distinguishes four levels of evaluation: analytical performance, concerning whether a measurement is technically valid; mechanistic performance, concerning whether the relevant processes are correctly represented; translational performance, concerning whether release predicts local or systemic exposure; and therapeutic performance, concerning whether exposure and biological response support the intended benefit–risk profile. Evidence should be interpreted only at the level it directly supports.

Advanced analysis of PLGA microspheres demonstrates the importance of polymer properties, excipient characterization, microstructure, residual components, drug distribution, and finished-product attributes in explaining long-acting behavior [23]. Quality-by-design approaches similarly organize critical material attributes, critical process parameters, and critical quality attributes as interacting determinants of product performance [24]. RERTF extends this reasoning beyond product quality by adding local transport, active-moiety exposure, biological feedback, and therapeutic context. The framework does not replace quality-by-design, pharmacokinetic modeling, or release-method development. It supplies an interpretive bridge among them, preventing an upstream quality or release metric from being treated as a complete therapeutic endpoint. **Figure 1** presents the release-to-exposure translation bridge, showing how the article's key components and boundaries are connected within evaluation criteria beyond cumulative release curves.

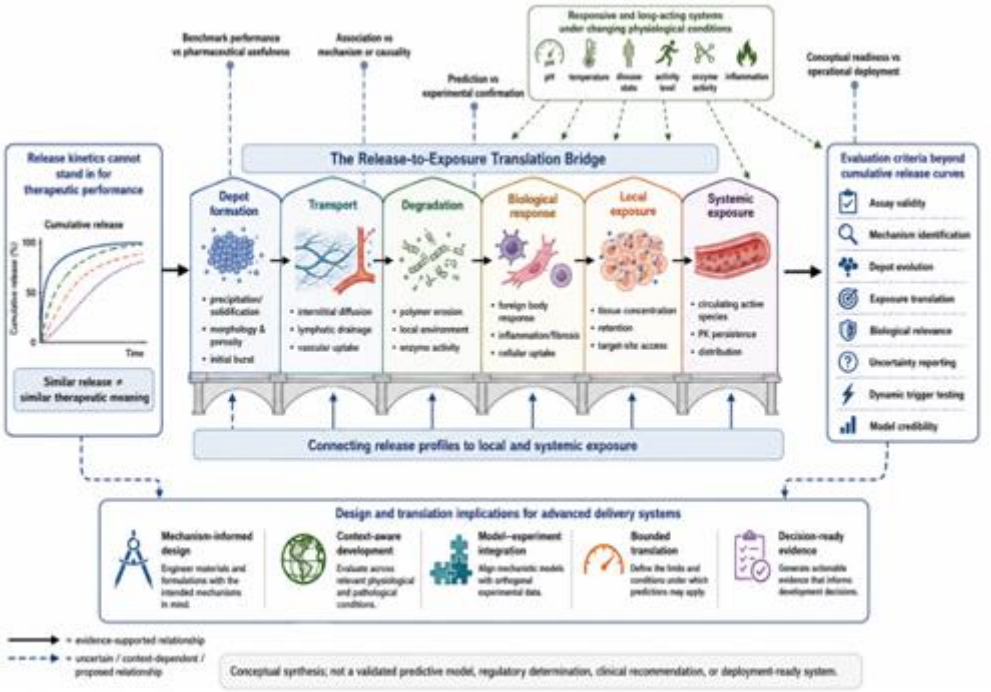

Figure 1. Release-to-Exposure Translation Bridge.

The figure is an original conceptual synthesis that organizes the article's central contribution across release kinetics cannot stand in for therapeutic performance, depot formation, transport, degradation, and biological response, depot formation, biological response, connecting release profiles to local and systemic exposure, responsive and long-acting systems under changing physiological conditions. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

No universal dissolution method can be assumed suitable for every long-acting injectable because different products depend on different combinations of dissolution, diffusion, erosion, degradation, and depot formation [25]. Evaluation must consequently be modular and claim-specific. At minimum, a strong assessment should identify the operative release mechanism, characterize depot evolution, demonstrate assay validity, examine relevant local transport, establish active-species exposure, evaluate biological response where claimed, and report uncertainty. Responsive systems additionally require dynamic stimulus testing, response lag, reversibility, repeated activation, false-trigger behavior, and loss-of-function assessment. Computational models should report parameter provenance, sensitivity, structural uncertainty, calibration, external evaluation, and the precise decision for which their predictions are intended. These dimensions should not be combined prematurely into one universal performance score, because aggregation can hide a critical failure at an individual layer.

Design and Translation Implications for Advanced Delivery Systems

The first design implication is that formulation objectives should be derived from the intended therapeutic function rather than from a preferred release-curve geometry. Implantable sustained-delivery devices require coordinated consideration of material selection, geometry, drug loading, mechanical integrity, manufacturing, insertion, retrieval, degradation, and patient use [26]. An extended duration may reduce dosing frequency, but it can also prolong adverse exposure, complicate dose adjustment, delay discontinuation, or increase the importance of implantation and removal procedures. RERTF therefore recommends defining a target exposure–response profile before selecting a target release profile. The desired release mechanism should then be treated as one design variable within a larger product architecture rather than as the final objective.

Prodrug-based long-acting systems make the distinction between release and active exposure particularly clear. Liberation of a prodrug from a formulation is followed by chemical or enzymatic conversion to the pharmacologically active species, and the conversion step may become the dominant determinant of exposure [27]. A formulation could therefore release prodrug reproducibly while producing variable active-drug concentrations because of differences in tissue enzymes, local chemistry, transport, or systemic metabolism. Design models must specify whether their output concerns formulation release, prodrug appearance, active-moiety formation, or biological response. Combining these quantities under a generic label such as "drug release" creates ambiguity and can misdirect optimization.

Long-acting microneedle systems likewise require integration of formulation behavior with device mechanics, skin insertion, dissolution or depot formation, tissue transport, user administration, and acceptability [28]. More broadly, clinical translation of long-acting formulations depends on manufacturability, stability, dose feasibility, route suitability, safety, pharmacokinetics, pharmacodynamics, adherence, and the comparative value of prolonged exposure [29]. Artificial-intelligence and computational-design workflows should therefore use multi-domain objective functions and explicit

constraints. A candidate should not be prioritized solely because a model predicts prolonged release or low release-profile error. It should also satisfy bounded requirements concerning material feasibility, active-moiety exposure, tolerability, reversibility, administration burden, uncertainty, and the intended patient population. Model outputs may support prioritization, but they do not replace experimental confirmation or clinical evaluation.

Patient-centric development adds another layer to therapeutic performance. Injection volume, needle characteristics, procedure complexity, implant visibility, retrieval options, residual drug, duration, administration setting, and patient preference can influence whether a technically successful system is appropriate in practice [30]. These factors are not external conveniences added after formulation optimization; they can define the acceptable design space from the beginning. The proposed framework accordingly treats translation as a sequence of bounded claims: conceptual plausibility, analytical characterization, mechanism-supported behavior, exposure translation, biological response, manufacturable product performance, and context-appropriate clinical use. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop design and translation implications for advanced delivery systems within the article's central argument.

Table 2. Evaluation, implementation, and research priorities arising from Release Profiles Do Not Equal Therapeutic Performance in Controlled, Responsive, and Long-Acting

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement
Therapeutic target definition	Intended indication, target tissue, therapeutic window, duration, safety needs, reversibility, administration setting	Defines the pharmaceutical question that the delivery system must answer	Context-specific target exposure and response objectives	Clinical and biological objectives may be incompletely defined early in development	Iterative refinement using pharmacological, clinical, and patient evidence
Formulation and depot design	Drug properties, polymer or lipid characteristics, geometry, loading, manufacturing parameters	Creates the physical system governing drug liberation	Reproducible and mechanistically interpretable depot state	Post-administration structure may differ from manufactured structure	Orthogonal structural characterization before and after administration
Release-mechanism evaluation	Dissolution, diffusion, swelling, erosion, degradation, cleavage, triggered transformation	Identifies how and when drug leaves the formulation	Mechanism-supported release description	Several mechanisms may yield similar cumulative curves	Mechanism-sensitive assays, imaging, mass balance, and perturbation studies
Local disposition assessment	Tissue architecture, binding, solubility, lymphatic transport, vascularity, local metabolism	Connects depot exit to local and systemic availability	Spatially and chemically defined local disposition	Direct local measurements may be difficult or invasive	Paired experimental measurements and qualified mechanistic models
Systemic exposure translation	Absorption, distribution, metabolism, elimination, active-moiety conversion	Predicts or explains circulating exposure	Context-appropriate pharmacokinetic profile	Parameters may be non-identifiable or platform-specific	External predictive evaluation and prospective pharmacokinetic testing
Biological-response assessment	Target engagement, biomarkers, tissue response, immune effects, toxicity, adaptation	Determines whether exposure produces intended and unintended effects	Exposure–response and safety interpretation	Biomarkers may not establish clinical benefit or causality	Mechanistic pharmacology and prospectively defined biological studies
Responsive-system performance	Stimulus range, kinetics, specificity, reversibility, repeated activation, material aging	Evaluates conditional and time-varying delivery	Defined operating region under realistic physiological change	Laboratory triggers may not reproduce disease-state variability	Dynamic, disease-relevant, repeated-cycle and off-target testing
Computational and AI modeling	Curated formulation, mechanism, transport, PK, PD, manufacturing, and usability data	Supports integration, simulation, uncertainty analysis, and candidate prioritization	Bounded predictions linked to explicit decisions	Dataset bias, missing mechanisms, non-transferability, and overconfidence	Calibration, sensitivity analysis, external validation, and human expert review
Manufacturing and product integration	Scale-up, process controls, stability, device compatibility, sterility, packaging	Converts a formulation concept into a reproducible product	Manufacturable and quality-controlled dosage form	Scale-up may alter microstructure and release mechanism	Process validation and clinically representative product characterization
Patient and use-context assessment	Administration burden, acceptability, adherence, retrieval, residual drug, care pathway	Determines whether technical performance can become practical therapeutic value	Context-appropriate use profile	Preferences and care infrastructure vary among populations	Human-factor, usability, acceptability, and clinical studies
Translation boundary governance	Intended claim, evidence level, uncertainty, decision authority	Prevents conclusions from exceeding the available evidence	Explicitly bounded pharmaceutical claim	Pressure to infer therapeutic value from surrogate performance	Independent review and stage-specific evidence criteria
Research implementation	Shared terminology, interoperable data,	Enables cumulative and reproducible theory testing	Comparable evidence across platforms and studies	Heterogeneous methods and incomplete reporting	Multicenter, cross-platform, and

prospective protocols, reporting standards	prospective validation programs

Figure 2 presents the evidence, validation, and translation boundary observatory for release profiles do not equal therapeutic performance, showing how the article's key components and boundaries are connected within design and translation implications for advanced delivery systems.

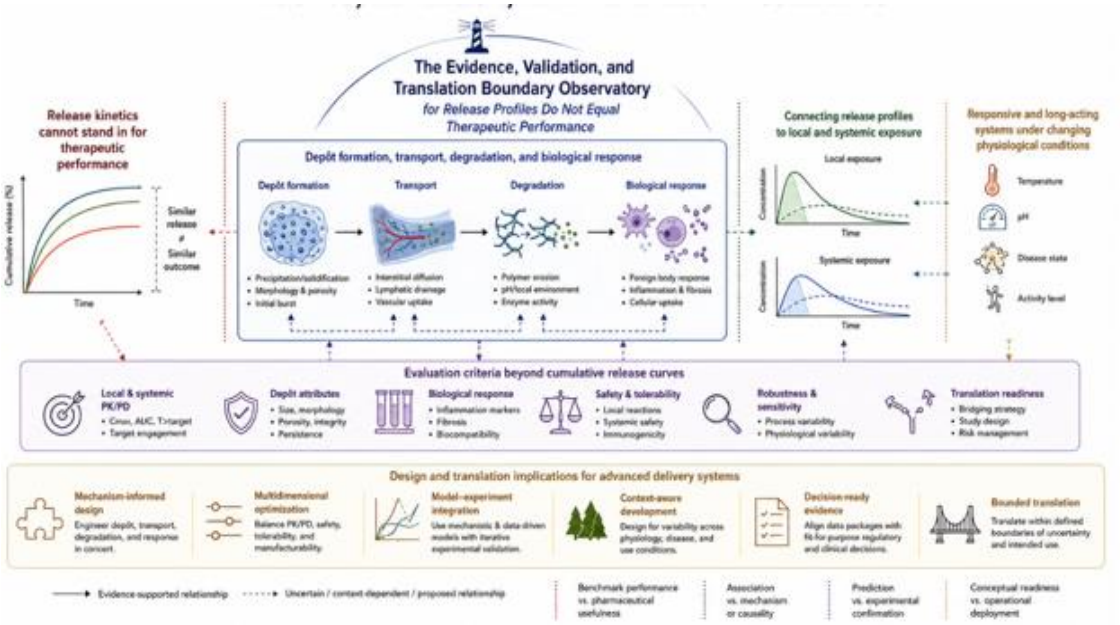

Figure 2. Evidence, Validation, and Translation Boundaries.

The figure is an original conceptual synthesis that shows how claims move from conceptual plausibility through evaluation, uncertainty assessment, and bounded pharmaceutical use. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Limitations and Boundary Conditions

RERTF is a conceptual contribution and has not been validated as a predictive model, scoring system, regulatory standard, or universal development framework. Its components are intentionally broad enough to apply across multiple controlled, responsive, and long-acting platforms, but this breadth limits platform-specific precision. The dominant mechanisms for a PLGA microsphere, lipid depot, implant, microneedle patch, prodrug suspension, or enzyme-responsive nanocarrier may differ substantially. The framework should therefore be instantiated for a defined dosage form, route, molecule, target tissue, and intended claim. It cannot establish which mechanism dominates without direct evidence, nor can it determine universal acceptance thresholds for release similarity, exposure comparability, biological response, or therapeutic utility.

The available evidence base is heterogeneous in platform, experimental design, analytical method, biological model, and translational maturity. Some relationships are supported by direct experimental observations, whereas others are represented mainly through mechanistic modeling or review-level synthesis. Local tissue exposure, lymphatic disposition, evolving depot structure, and long-duration biological feedback remain particularly difficult to observe directly. Computational agreement with plasma pharmacokinetics may be achieved despite uncertainty in local mechanisms, and a fitted model may not remain transferable across formulation changes or patient populations [12]. Similarly, a discriminatory in vitro method can improve product understanding without establishing clinical predictiveness [25]. The framework therefore requires explicit separation of measured evidence, model-derived estimates, mechanistic hypotheses, and downstream therapeutic interpretations.

Misuse may occur if the framework is converted prematurely into a checklist whose completion is treated as proof of therapeutic performance. Not every development question requires direct measurement of every component, and evidence should remain proportional to the intended decision. Conversely, omitting a component because it is difficult to measure does not make its influence negligible. The framework also cannot resolve regulatory equivalence, clinical benefit, treatment adherence, patient preference, or benefit–risk acceptability without the corresponding empirical evidence. It should be used to expose missing links and restrict overclaiming, not to authorize claims that have not been experimentally or clinically established.

Research Agenda and Implementation Priorities

The first research priority is to generate matched datasets spanning multiple RERTF layers. Studies should connect depot structure, mechanism-sensitive release, local disposition, systemic exposure, biological response, and relevant safety observations within the same formulation comparison. Longitudinal imaging and chemical characterization are needed to determine how depot state changes across the intended dosing interval, while paired local and systemic measurements are needed to distinguish true formulation-controlled release from tissue-controlled absorption. Dynamic testing systems should represent changing physiological variables rather than only fixed pH, enzyme, glucose, or redox conditions. Responsive systems should be assessed for activation threshold, response lag, reversibility, repeated cycling, false activation, depletion, and hysteresis. These studies should be designed prospectively around specific translation hypotheses instead of adding mechanistic measurements only after unexpected pharmacokinetics emerge.

The second priority is to improve computational credibility and data suitability. Models should specify the chemical species, anatomical compartments, mechanisms, parameter sources, and intended decisions they represent. Structural and practical identifiability should be evaluated before mechanistic interpretation is claimed. Sensitivity analysis should determine which depot, tissue, and physiological parameters control predicted exposure, and external evaluation should test transferability across batches, formulations, doses, routes, or populations. Artificial-intelligence models require datasets that link formulation attributes to mechanistic and translational outcomes rather than release curves alone. Reporting should distinguish interpolation from extrapolation, association from mechanism, calibrated uncertainty from raw model confidence, and prediction from experimental confirmation. Human pharmaceutical oversight remains necessary when model objectives omit safety, manufacturability, usability, or patient-context constraints.

The third priority is staged implementation through shared terminology, interoperable evidence structures, and decision-specific reporting. Early-stage use should focus on hypothesis organization, experiment selection, and identification of missing evidence. Intermediate use may include mechanistic model qualification, formulation ranking, and design-space exploration when external checks are available. Later translational use requires integration with manufacturing, safety, PK/PD, administration, usability, and clinical evidence. Cross-platform consortia could develop minimum reporting elements for depot state, release mechanism, assay conditions, local transport assumptions, active-species exposure, biological response, uncertainty, and context of use. Such infrastructure may support comparison and cumulative learning, but implementation should remain conditional on prospective evaluation and should not be interpreted as evidence that the proposed framework is operationally ready.

Conclusion

Release profiles are scientifically important but therapeutically incomplete representations of controlled, responsive, and long-acting drug-delivery systems. The proposed Release–Exposure–Response Translation Framework separates formulation and depot state, release-generating processes, local transport and transformation, local and systemic exposure, biological response, and therapeutic performance while making the evidentiary boundaries among them explicit. Its central contribution is the proposition that similarity, optimization, or predictability at one layer cannot be propagated automatically to the next. Depot evolution, physiological variability, biological feedback, active-moiety formation, route-specific transport, uncertainty, usability, and patient context can each alter the therapeutic meaning of an apparently favorable release profile. The framework is not a validated predictive model, equivalence standard, or deployment-ready decision tool. It is a theory-building structure intended to improve experimental design, computational modeling, data interpretation, and translational discipline by ensuring that claims remain proportional to the evidence connecting release to exposure, response, and therapeutic value.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Nkanga CI, Fisch A, Rad-Malekshahi M, Duvnjak Romic M, Kittel B, Ullrich T, et al. Clinically established biodegradable long acting injectables: An industry perspective. Adv Drug Deliv Rev. 2020;167:19-46. doi:10.1016/j.addr.2020.11.008
2. Abdelkader H, Fathalla Z, Seyfoddin A, Farahani M, Thrimawithana T, Allahham A, et al. Polymeric long-acting drug delivery systems (LADDS) for treatment of chronic diseases: Inserts, patches, wafers, and implants. Adv Drug Deliv Rev. 2021;177:113957. doi:10.1016/j.addr.2021.113957
3. Jindal AB, Bhide AR, Salave S, Rana D, Benival D. Long-acting parenteral drug delivery systems for the treatment of chronic diseases. Adv Drug Deliv Rev. 2023;198:114862. doi:10.1016/j.addr.2023.114862

4. Pandya AK, Vora LK, Umeyor C, Surve D, Patel A, Biswas S, et al. Polymeric in situ forming depots for long-acting drug delivery systems. Adv Drug Deliv Rev. 2023;200:115003. doi:10.1016/j.addr.2023.115003
5. Sharma R, Yadav S, Yadav V, Akhtar J, Katari O, Kuche K, et al. Recent advances in lipid-based long-acting injectable depot formulations. Adv Drug Deliv Rev. 2023;199:114901. doi:10.1016/j.addr.2023.114901
6. Andhariya JV, Shen J, Choi S, Wang Y, Zou Y, Burgess DJ. Development of in vitro-in vivo correlation of parenteral naltrexone loaded polymeric microspheres. J Control Release. 2017;255:27-35. doi:10.1016/j.jconrel.2017.03.396
7. Andhariya JV, Jog R, Shen J, Choi S, Wang Y, Zou Y, et al. In vitro-in vivo correlation of parenteral PLGA microspheres: Effect of variable burst release. J Control Release. 2019;314:25-37. doi:10.1016/j.jconrel.2019.10.014
8. Ho MJ, Jeong MY, Jeong HT, Kim MS, Park HJ, Kim DY, et al. Effect of particle size on in vivo performances of long-acting injectable drug suspension. J Control Release. 2022;341:533-47. doi:10.1016/j.jconrel.2021.12.011
9. Wang X, Burgess DJ. Drug release from in situ forming implants and advances in release testing. Adv Drug Deliv Rev. 2021;178:113912. doi:10.1016/j.addr.2021.113912
10. Gao GF, Thurn M, Wendt B, Parnham MJ, Wacker MG. A sensitive in vitro performance assay reveals the in vivo drug release mechanisms of long-acting medroxyprogesterone acetate microparticles. Int J Pharm. 2020;586:119540. doi:10.1016/j.ijpharm.2020.119540
11. Perazzolo S, Shireman LM, McConnachie LA, Koehn J, Kinman L, Lee W, et al. Integration of computational and experimental approaches to elucidate mechanisms of first-pass lymphatic drug sequestration and long-acting pharmacokinetics of the injectable triple-HIV drug combination TLC-ART 101. J Pharm Sci. 2020;109(5):1789-801. doi:10.1016/j.xphs.2020.01.016
12. Zheng F, Hou P, Corpstein CD, Park K, Li T. Multiscale pharmacokinetic modeling of systemic exposure of subcutaneously injected biotherapeutics. J Control Release. 2021;337:407-16. doi:10.1016/j.jconrel.2021.07.043
13. Corpstein CD, Hou P, Park K, Li T. Multiphysics simulation of local transport and absorption coupled with pharmacokinetic modeling of systemic exposure of subcutaneously injected drug solution. Pharm Res. 2023;40(12):2873-86. doi:10.1007/s11095-023-03546-5
14. Shah JC, Hong J. Model for long acting injectables (depot formulation) based on pharmacokinetics and physical chemical properties. AAPS J. 2022;24(3):44. doi:10.1208/s12248-022-00695-0
15. Bettonte S, Berton M, Battegay M, Stader F, Marzolini C. Development of a physiologically-based pharmacokinetic model to simulate the pharmacokinetics of intramuscular antiretroviral drugs. CPT Pharmacometrics Syst Pharmacol. 2024;13(5):781-94. doi:10.1002/psp4.13118
16. Majumder J, Minko T. Multifunctional and stimuli-responsive nanocarriers for targeted therapeutic delivery. Expert Opin Drug Deliv. 2021;18(2):205-27. doi:10.1080/17425247.2021.1828339
17. Wang J, Wang Z, Yu J, Kahkoska AR, Buse JB, Gu Z. Glucose-responsive insulin and delivery systems: Innovation and translation. Adv Mater. 2020;32(13). doi:10.1002/adma.201902004
18. Agrawal M, Saraf S, Saraf S, Dubey SK, Puri A, Gupta U, et al. Stimuli-responsive in situ gelling system for nose-to-brain drug delivery. J Control Release. 2020;327:235-65. doi:10.1016/j.jconrel.2020.07.044
19. Zhang C, Jiang G, Gao X. Matrix metalloproteinase-responsive drug delivery systems. Bioconjug Chem. 2023;34(8):1349-65. doi:10.1021/acs.bioconjchem.3c00266
20. Li R, Peng F, Cai J, Yang D, Zhang P. Redox dual-stimuli responsive drug delivery systems for improving tumor-targeting ability and reducing adverse side effects. Asian J Pharm Sci. 2020;15(3):311-25. doi:10.1016/j.ajps.2019.06.003
21. Nothnagel L, Wacker MG. How to measure release from nanosized carriers? Eur J Pharm Sci. 2018;120:199-211. doi:10.1016/j.ejps.2018.05.004
22. O'Brien MN, Jiang W, Wang Y, Loffredo DM. Challenges and opportunities in the development of complex generic long-acting injectable drug products. J Control Release. 2021;336:144-58. doi:10.1016/j.jconrel.2021.06.017
23. Wan B, Bao Q, Burgess DJ. Long-acting PLGA microspheres: Advances in excipient and product analysis toward improved product understanding. Adv Drug Deliv Rev. 2023;198:114857. doi:10.1016/j.addr.2023.114857
24. Zhang C, Yang L, Wan F, Bera H, Cun D, Rantanen J, et al. Quality by design thinking in the development of long-acting injectable PLGA/PLA-based microspheres for peptide and protein drug delivery. Int J Pharm. 2020;585:119441. doi:10.1016/j.ijpharm.2020.119441
25. Malavia N, Bao Q, Burgess DJ. Novel dissolution methods for drug release testing of long-acting injectables. Int J Pharm. 2024;664:124634. doi:10.1016/j.ijpharm.2024.124634
26. Magill E, Demartis S, Gavini E, Permana AD, Thakur RRS, Adrianto MFA, et al. Solid implantable devices for sustained drug delivery. Adv Drug Deliv Rev. 2023;199:114950. doi:10.1016/j.addr.2023.114950
27. Chien ST, Suydam IT, Woodrow KA. Prodrug approaches for the development of a long-acting drug delivery systems. Adv Drug Deliv Rev. 2023;198:114860. doi:10.1016/j.addr.2023.114860
28. Vora LK, Sabri AH, Naser Y, Himawan A, Hutton ARJ, Anjani QK, et al. Long-acting microneedle formulations. Adv Drug Deliv Rev. 2023;201:115055. doi:10.1016/j.addr.2023.115055
29. Li W, Tang J, Lee D, Tice TR, Schwendeman SP, Prausnitz MR. Clinical translation of long-acting drug delivery formulations. Nat Rev Mater. 2022;7(5):406-20. doi:10.1038/s41578-021-00405-w
30. Alidori S, Subramanian R, Holm R. Patient-centric long-acting injectable and implantable platforms—an industrial perspective. Mol Pharm. 2024;21(9):4238-58. doi:10.1021/acs.molpharmaceut.4c00665