Available online www.ijpras.com

International Journal of Pharmaceutical Research & Allied Sciences, 2026, 15(2):20-29
https://doi.org/10.51847/5GTArIIafD

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

From Platform Readiness to Patient Readiness through a Decision Architecture for Next-Generation Pharmaceutical Translation under Scientific Uncertainty

James Walker[1]*, Olivia Harris[1], George Wilson[2]

[1]*Department of Platform-to-Patient Readiness and Decision Architecture, Faculty of Pharmacy, University of Edinburgh, Edinburgh, United Kingdom.*
[2]*Department of Pharmaceutical Translation and Scientific Uncertainty, Faculty of Pharmacy, University of Leeds, Leeds, United Kingdom.*

***Email:** james.walker@ed.ac.uk*

ABSTRACT

Next-generation pharmaceutical platforms can generate compounds, models, assays, manufacturing processes, and human-relevant experimental evidence with increasing speed and technical sophistication. Yet platform performance is frequently interpreted as evidence of translational maturity before product-specific quality, biological consequence, pharmacological behaviour, clinical utility, or implementation feasibility has been established. This article develops an original Translational Decision Architecture for distinguishing platform readiness from patient readiness under scientific uncertainty. The proposed architecture treats translation as a sequence of connected but non-substitutable evidence domains: platform performance and reproducibility; product-specific manufacturability and quality; biological consequence and mechanism validation; pharmacological and safety translation; clinical utility and usability; and implementation readiness. Advancement between domains is governed by conditional evidence-transition gates rather than by a single maturity score or linear development ladder. Patient readiness is defined not as an approved status but as a context-specific decision condition requiring adequate evidence across all domains relevant to the proposed use. The architecture further introduces evidence warrants, explicit context-of-use statements, non-compensatory constraints, uncertainty-weighted decision gates, readiness regression, stopping rules, and an evidence-renewal function. Clinical utility, usability, and implementation are treated as constitutive translational requirements rather than downstream operational considerations. Validation would require retrospective and prospective testing across computational, experimental, manufacturing, pharmacological, clinical, and implementation settings. The architecture does not establish predictive accuracy, therapeutic effectiveness, regulatory acceptability, or universal applicability. Its principal contribution is a falsifiable conceptual structure for preventing technical achievement from being mistaken for patient-relevant pharmaceutical readiness.

Key words: *Translational pharmaceutical research, Translation readiness, Reproducibility, Manufacturing feasibility, Clinical relevance, Evidence maturity*

Received: 01 February 2026; **Revised:** 17 April 2026; **Accepted:** 19 April 2026

INTRODUCTION

Pharmaceutical translation is often represented as a progression from discovery through preclinical development, clinical investigation, and implementation. This representation can obscure the heterogeneity of the evidence required at each transition. Development success varies across phases and therapeutic contexts, while programme

failure may arise from deficiencies in target validity, exposure, safety, patient selection, product quality, clinical differentiation, or development strategy rather than from one universal bottleneck [1-3]. Technical progress in one part of this chain therefore cannot be assumed to resolve weaknesses elsewhere.

Multidimensional development frameworks have shown the value of integrating evidence concerning targets, tissues, safety, patient populations, and differentiated therapeutic value [2]. Nevertheless, such frameworks do not fully address a growing translational problem: the expanding use of general-purpose or modality-specific platforms whose technical performance may be demonstrated before the evidentiary consequences of a particular product, mechanism, population, or healthcare context are understood. A platform may generate reproducible outputs, support rapid experimentation, or operate within a controlled technical envelope while remaining poorly qualified for a patient-relevant decision.

This article is an Original Translational Decision Architecture Article. Its objective is to develop a non-empirical conceptual structure for determining how evidence should be organised when a pharmaceutical platform progresses from technical capability toward product-specific and patient-relevant use. The proposed contribution is not a new development phase model, technology-readiness scale, regulatory standard, clinical recommendation, or validated prediction system. It is a decision architecture intended to clarify which claims may legitimately be made at different evidence states and which transitions require additional validation.

The central argument is that platform readiness and patient readiness are related but non-equivalent constructs. The architecture separates six readiness domains, connects them through conditional evidence-transition gates, and introduces explicit mechanisms for uncertainty weighting, evidence renewal, stopping, and regression. It is designed to prevent evidence generated for one purpose from being transferred to another without examining context, assumptions, failure consequences, and the strength of the connecting inference. Every new component and decision rule proposed here requires empirical evaluation before operational use.

Why technical platform readiness differs from patient readiness

Technical platforms can improve the consistency, scale, speed, or informational density of pharmaceutical research, but these capabilities do not independently establish that a target is therapeutically valid, that a computational output is decision-relevant, or that an experimental model reproduces the conditions necessary for human translation. Structured target assessment, bounded use of machine-learning systems, and qualification of human-relevant models each reveal a different reason why technical capability cannot be treated as sufficient evidence of patient benefit [4-6].

In this architecture, platform readiness means that a declared technical function can be performed credibly within a specified context of use. Its components include an explicit operating envelope, controlled inputs, repeatability, reproducibility, traceable outputs, characterised failure modes, and evidence that the measured endpoint corresponds to the platform claim. Target-assessment frameworks illustrate why even technically robust assays must be connected to evidence concerning biological relevance, safety, feasibility, and the intended therapeutic hypothesis [4]. Platform readiness therefore establishes the credibility of a technical claim, not the credibility of every downstream pharmaceutical inference.

Patient readiness is defined more narrowly and more demanding: a context-specific decision condition in which the intervention, evidence package, target population, clinical objective, delivery pathway, and implementation environment are sufficiently aligned to justify a defined next action. It does not mean that uncertainty has been eliminated or that the intervention is approved, clinically effective, or suitable for unrestricted deployment. It means that unresolved uncertainty has been made visible and judged in relation to the consequences of the particular decision.

The difference between these constructs becomes especially important for computational and experimental platforms. Machine-learning systems may identify patterns, rank compounds, generate hypotheses, or estimate properties, yet their usefulness depends on data provenance, domain shift, external validation, uncertainty characterisation, and whether the predicted quantity is causally or operationally connected to the subsequent pharmaceutical decision [5]. Similarly, organs-on-chips and related human-relevant systems may improve selected aspects of physiological representation, but standardisation, reproducibility, model qualification, and context-of-use limitations remain central to their translational interpretation [6].

Platform readiness is therefore neither irrelevant nor merely preliminary. It is a necessary evidentiary condition whose meaning must remain bounded. The proposed distinction rejects two opposing errors: treating technical performance as patient readiness, and treating technical evidence as unimportant until clinical testing. Instead, technical evidence should be preserved as a domain-specific warrant whose downstream significance depends on

product-specific, biological, pharmacological, clinical, and implementation evidence. This relationship is proposed as conditional rather than automatic.

Proposed readiness-transition architecture

Translational science has increasingly been framed as a systematic effort to identify and overcome recurrent barriers, integrate heterogeneous evidence, and manage risk across the development pathway. These perspectives support an architecture in which translation is governed by linked scientific and decision problems rather than by isolated technical milestones [7-9]. The proposed Readiness-Transition Decision Architecture builds on that foundation while introducing explicit boundaries between evidence domains and explicit rules for movement among them.

The first domain is platform performance and reproducibility, which concerns whether the underlying technical system operates credibly within a stated context. The second is product-specific manufacturability and quality, which concerns whether a particular intervention can be produced, controlled, stored, delivered, and maintained with attributes appropriate to its intended function. The third is biological consequence and mechanism validation, which concerns whether the intervention produces the intended causal or mechanistically relevant effect in systems appropriate to the disease and population. The remaining domains are pharmacological and safety translation, clinical utility and usability, and implementation readiness. This decomposition extends the view that translational science must address recurrent scientific and operational barriers across the full path to health impact [7].

The architecture does not require that every programme follow an identical sequence. Instead, it proposes conditional evidence-transition gates. Each gate asks whether the evidence generated in one domain is sufficiently credible, relevant, and portable to justify action in another. A transition may permit advancement, restrict the scope of the next claim, require additional evidence, trigger redesign, pause progression, or support stopping. Because pharmaceutical risks are distributed across targets, products, processes, clinical strategies, and development operations, the architecture treats a serious unresolved deficit as a potential veto condition rather than assuming that strengths elsewhere can compensate for it [9].

Every gate is informed by an Evidence Maturity and Provenance Ledger, a proposed construct recording what is known, how it was generated, which assumptions connect it to the decision, where uncertainty remains, and what changes would invalidate its use. Evidence maturity is not reduced to publication count, data volume, or a single hierarchy. It includes internal validity, reproducibility, independence, directness, relevance to the intended population and product, temporal currency, and fitness for the decision consequence. Broader conceptions of evidence-based medicine support the need to integrate heterogeneous forms of evidence while preserving their distinct purposes and limitations [8].

The architecture also applies a non-substitutability principle. Strong platform reproducibility cannot neutralise unmanageable product instability; favourable manufacturability cannot establish the intended biological consequence; mechanistic plausibility cannot replace exposure and safety evidence; clinical efficacy cannot by itself establish usability or equitable implementation. These relationships are proposed as testable decision rules rather than empirical findings. Their value must be evaluated by examining whether domain-resolved decisions are better calibrated, more transparent, and more responsive to contradictory evidence than decisions based on aggregate maturity labels.

Platform performance and reproducibility

Platform performance concerns the capacity of a technical system to produce a specified output under declared conditions. Reproducibility is one part of that capacity, but it is not a single property. It may refer to consistency across runs, analysts, instruments, sites, datasets, laboratories, or time. Reproducible science also depends on study design, transparent methods, appropriate incentives, accessible materials and data, and cumulative correction mechanisms rather than replication alone [10]. Accordingly, a platform claim should specify what must remain invariant, what variation is permitted, and which output is expected to reproduce.

Reporting quality is an essential but limited element of this assessment. Detailed reporting allows others to judge design adequacy, reproduce procedures, interpret exclusions, and identify possible sources of bias. Updated animal-research guidance demonstrates the importance of reporting design, sample-size reasoning, inclusion and exclusion criteria, randomisation, blinding, outcome measures, and statistical methods [11]. However, complete reporting does not transform a weak experimental design into a strong one, establish the biological relevance of the endpoint, or prove that results will transfer to another product or population.

Independent replication provides a stronger challenge to claim portability, but its interpretation also requires care. Multisite efforts in preclinical cancer biology have demonstrated that influential findings may be difficult to reproduce completely and that incomplete methods, material differences, design constraints, and statistical uncertainty can complicate replication conclusions [12]. Such evidence supports caution toward unexamined platform claims, but it does not justify assuming that all preclinical findings are irreproducible or that one replication outcome determines the validity of an entire platform.

The architecture therefore proposes a Platform Evidence Warrant. Before platform outputs are used to justify a transition, the warrant would state the intended context of use, technical operating envelope, input specifications, output definition, reproducibility dimensions, data and material provenance, known failure modes, external replication expectations, and conditions under which the evidence becomes non-portable. The warrant does not certify product quality, mechanism validity, human pharmacology, safety, or clinical utility. It establishes only that the platform claim is technically credible enough to support the next explicitly bounded question. Its structure, minimum contents, and effect on subsequent decisions require prospective validation.

Product-specific manufacturability and quality

A platform can perform reliably while a particular product remains unsuitable for manufacture, storage, delivery, or clinical use. Product readiness therefore concerns the intervention's specific material properties, critical quality attributes, formulation, stability, process sensitivity, and compatibility with the intended route of administration.

Clinical-stage antibodies occupy constrained biophysical-property ranges, illustrating that target activity alone does not establish developability [13]. These observations are modality-specific, but they support a broader principle: product characteristics must be evaluated independently of the platform that discovered, designed, or tested the intervention.

Computational profiling may identify aggregation, viscosity, charge, solubility, or interaction risks before extensive experimental development [14]. Such predictions can prioritise testing, but they cannot substitute for product-specific measurements, scale-dependent studies, stability evaluation, or confirmation that manufacturing operations preserve biological function.

Manufacturing readiness additionally requires process understanding, control strategies, quality monitoring, and lifecycle management [15]. The architecture therefore treats manufacturability as a separate evidence warrant. A technically mature platform cannot confer readiness on a product whose quality attributes remain unstable, poorly controlled, or disconnected from biological and clinical consequences.

Biological consequence and mechanism validation

Biological readiness concerns whether an intervention produces the intended consequence through a mechanism relevant to the disease, target population, and proposed use. It requires more than target binding, pathway modulation, or correlation between a platform output and a biological label.

Human genetic evidence can strengthen confidence that modulation of a target is related to disease biology, but it does not guarantee that a particular intervention, direction of effect, exposure profile, or patient population will produce clinical benefit [16]. Genetic support must therefore remain one component of mechanism validation.

Patient-derived organoids can reproduce selected treatment-response patterns within defined malignancies and experimental settings [17]. Their translational contribution depends on tissue representation, sampling, culture conditions, endpoints, treatment context, and external comparison. Human origin alone does not establish universal patient prediction.

Liver-chip systems have reproduced selected human and cross-species toxicity patterns, demonstrating the potential value of contextually qualified human-relevant models [18]. The architecture consequently proposes a Biological Consequence Warrant documenting causal assumptions, target engagement, model relevance, orthogonal evidence, failure modes, and the conditions under which biological findings may inform downstream pharmacological decisions.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in from platform readiness to patient readiness through a decision architecture for.

Table 1. Definitions, components, relationships, and intended uses in From Platform Readiness to Patient Readiness through a Decision Architecture for.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Platform readiness	Unreliable technical claims	Operating envelope, inputs, outputs, provenance	Proposed Platform Evidence Warrant	Credible bounded technical use	Reproducibility, reporting, independent challenge [10]	Reproducible but irrelevant endpoint	Does not establish product or patient readiness
Product readiness	Platform–product mismatch	Material properties, formulation, process, stability	Product-specific quality warrant	Manufacturable and controlled intervention	Experimental developability and process evidence [13]	Scale sensitivity or unstable attributes	Analytical quality does not establish clinical equivalence
Biological readiness	Plausibility mistaken for consequence	Target support, perturbation, human-relevant models	Proposed Biological Consequence Warrant	Mechanistically interpretable evidence	Genetic, causal, orthogonal, and human-context evidence [16]	Association, model bias, context mismatch	Does not establish efficacy
Pharmacological and safety readiness	Biological activity without usable exposure or safety	Dose, exposure, response, distribution, toxicity	Integrated exposure–response and safety assessment	Bounded human-use hypothesis	Qualified PK and safety models	Incorrect dose translation or unrecognised toxicity	Does not establish favourable benefit–risk
Clinical-utility readiness	Efficacy detached from patient value	Outcomes, population, comparators, usability	Context-specific utility assessment	Evidence relevant to a defined care decision	Prospective clinical and usability evidence	Surrogate benefit or poor actionability	Does not establish implementation
Implementation readiness	Effective intervention not usable in practice	Workflow, infrastructure, personnel, access	Context-sensitive implementation assessment	Feasible and sustainable use	Multilevel implementation evaluation	Inequitable access or workflow failure	Adoption does not prove clinical effectiveness
Evidence maturity	Data volume mistaken for credibility	Validity, independence, directness, currency	Proposed Evidence Maturity and Provenance Ledger	Transparent evidentiary basis	Traceable and decision-relevant evidence	Hidden assumptions or obsolete evidence	No universal maturity score is proposed
Decision transition	Automatic progression between domains	Evidence warrants, uncertainty, consequences	Proposed conditional gate with veto constraints	Calibrated advance, pause, redesign, or stop	Prospective decision validation	Compensation for a critical deficit	Gate rules are conceptual and unvalidated

Pharmacological and safety translation

Biological activity becomes pharmacologically relevant only when exposure, concentration, timing, tissue distribution, response, and variability are connected to the intended intervention. A convincing mechanism may still fail because the necessary exposure cannot be achieved safely or maintained in the target tissue.

Fluidically connected organ-chip systems have supported quantitative investigation of human pharmacokinetic responses under defined conditions [19]. Such studies demonstrate a possible bridge between experimental platforms and pharmacological prediction, but their portability depends on compound properties, system configuration, model assumptions, and external validation.

Safety translation requires an equally explicit context of use. Predictive models for drug-induced liver injury illustrate the need to align mechanisms, endpoints, assay performance, qualification standards, and development decisions [20]. A model that detects one hazard class cannot be assumed to characterise all toxicities or determine overall benefit–risk.

The proposed Pharmacological and Safety Warrant would therefore document dose rationale, exposure assumptions, uncertainty, therapeutic-window evidence, off-target effects, model qualification, and unresolved safety signals. Advancement would remain conditional on the consequence of error rather than on technical performance alone.

Clinical utility, usability, and implementation

Clinical utility concerns whether an intervention can improve a meaningful decision or outcome for a defined population. Usability concerns whether patients and professionals can understand, deliver, monitor, and act on it. Implementation concerns whether these functions can be sustained within real healthcare settings.

Multicountry pharmacogenomic implementation required coordination among laboratories, prescribing workflows, informatics, education, governance, and clinical evaluation [21]. This experience shows that patient relevance depends on organisational and behavioural systems as well as scientific evidence.

Implementation frameworks further identify determinants associated with the intervention, inner and outer settings, individuals, and implementation processes [22]. These determinants may explain variation in uptake, but they do not guarantee effectiveness, sustainability, or equity.

Patient readiness is therefore proposed as a configuration rather than a final technology status. It requires sufficient alignment among product quality, mechanism, exposure, safety, meaningful benefit, usability, infrastructure, access, and monitoring for the specific decision under consideration.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for testable propositions, evidence requirements, and validation criteria for from platform readiness to patient readiness through a decision architecture for.

Table 2. Testable propositions, evidence requirements, and validation criteria for From Platform Readiness to Patient Readiness through a Decision Architecture for.

Architecture layer or process stage	Core function	Information or material flow	Interaction with other components	Validation criterion	Potential failure mode	Human or experimental responsibility	Readiness boundary
Platform Evidence Warrant	Establish bounded technical credibility	Inputs to reproducible output	Supplies evidence to product assessment	Cross-run, site, operator, and time performance	Technical drift	Define context and challenge assumptions	Technical credibility only
Product-quality transition	Test product-specific portability	Platform output to controlled product	Links process with biological function	Stable critical attributes under intended conditions [15]	Process-induced change	Characterise product and process	Manufacturability is not therapeutic benefit
Biological Consequence Warrant	Test causal relevance	Product exposure to biological effect	Informs pharmacological hypotheses	Orthogonal and human-context confirmation [17]	Model-specific artefact	Select relevant models and endpoints	Plausibility is not validation
Pharmacological–safety transition	Connect dose with consequence	Product to exposure, response, and toxicity	Constrains clinical evaluation	Qualified exposure–response and safety evidence [20]	Unsafe or unattainable exposure	Define stopping and uncertainty rules	Preclinical evidence is not benefit–risk
Clinical-utility transition	Determine meaningful patient value	Intervention to clinical decision and outcome	Depends on usability and implementation	Prospective intended-population evaluation	Surrogate success without utility	Select outcomes and comparators	Efficacy is not implementation
Implementation transition	Test practical and equitable use	Evidence into healthcare workflow	Feeds real-world evidence renewal	Feasibility, acceptability, access, sustainability [22]	Workflow or access failure	Evaluate context and affected groups	Uptake is not validation

Proposed non-substitutability rule	Prevent cross-domain compensation	Domain-resolved evidence to gate decision	Applies across all transitions	Compare domain-specific and aggregate decisions	Strong metrics conceal critical deficit	Record reasons for advance or stop	Requires empirical testing
Proposed regression rule	Reopen prior conclusions	New evidence back to earlier domains	Activates renewal and stopping	Detection of contradiction, drift, or context change	Irreversible commitment to stale evidence	Preserve authority to reassess	No readiness state is permanent

Uncertainty-weighted decision gates

A readiness gate should not ask only whether evidence exists. It should ask how directly the evidence supports the decision, how independent and reproducible it is, which assumptions connect it to the next action, and what harm could follow from error.

Probability-of-pharmacological-success approaches illustrate how heterogeneous pharmacological evidence may be organised into a decision construct [23]. Their usefulness depends on transparent definitions, calibration, and context; no probability estimate should be treated as a universal readiness threshold.

Bayesian methods provide a formal basis for updating beliefs as new evidence accumulates [24]. They can support cumulative learning, but posterior estimates remain conditional on model structure, prior assumptions, data quality, and the correspondence between measured endpoints and decision consequences.

The architecture therefore proposes uncertainty-weighted gates that permit advance, restricted advance, evidence generation, redesign, pause, or stop. Gate weights and thresholds are deliberately unspecified. They must be calibrated prospectively to the intervention, uncertainty source, reversibility of the action, and consequences of false progression or premature termination.

Readiness regression, stopping, and evidence renewal

Readiness should not be treated as permanently accumulated. A new manufacturing process, population, formulation, assay, dataset, safety signal, or implementation environment can invalidate assumptions that previously supported progression.

Predictive validity concerns whether a model or decision tool forecasts later usefulness rather than merely fitting an earlier technical endpoint [25]. When that relationship weakens, the appropriate response may be reassessment rather than continued progression.

Stopping is not synonymous with scientific failure. It may represent a rational response to a non-correctable product limitation, an unsafe exposure requirement, absent clinical differentiation, irreducible uncertainty, or an implementation condition that defeats the intended use.

Longitudinal organisational experience suggests that disciplined evidence use, learning, and progression decisions can influence pharmaceutical R&D productivity [26]. Such observations support evidence renewal but do not validate the proposed architecture or establish that one organisational model is universally superior.

The architecture therefore introduces readiness regression and a proposed translation observatory. The observatory would monitor contradictory evidence, drift, process changes, new populations, altered contexts of use, and implementation feedback, while preserving the provenance of earlier decisions.

Figure 1 presents the original conceptual synthesis for decision architecture connecting platform readiness to patient readiness, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

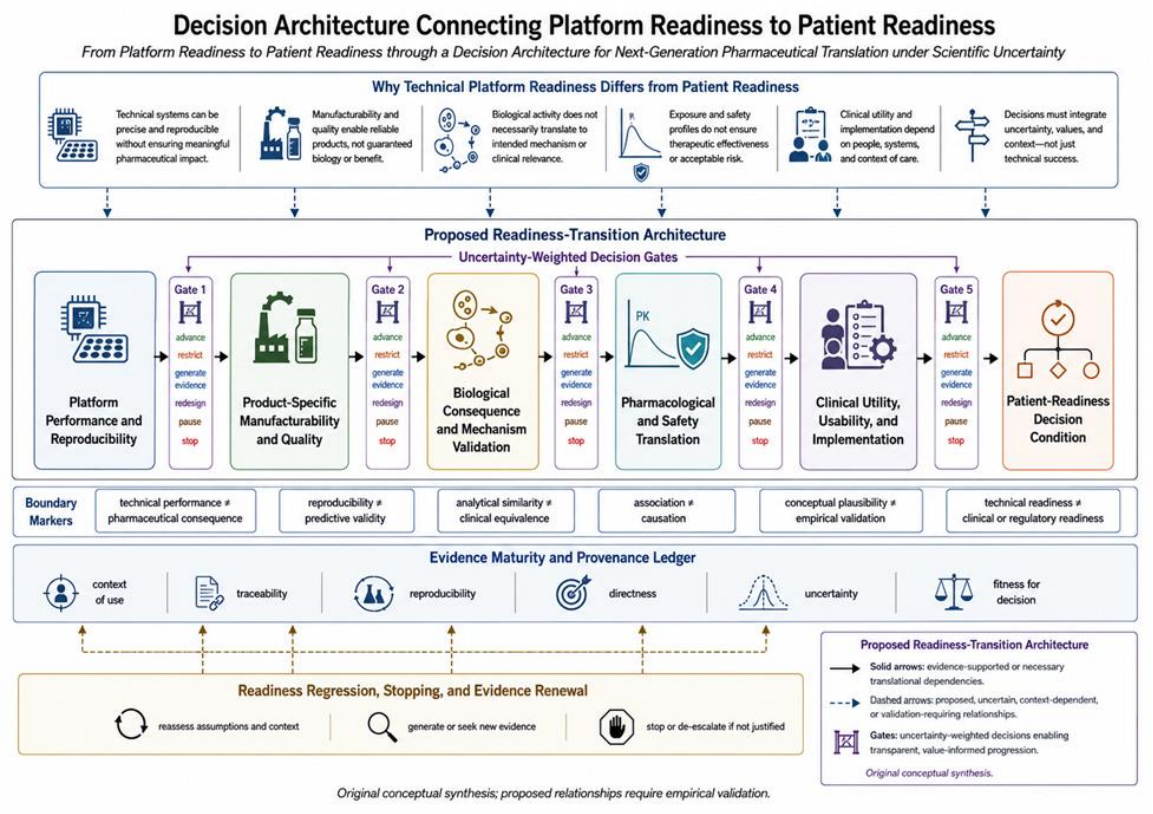

Figure 1. Decision Architecture Connecting Platform Readiness to Patient Readiness

Limitations and research priorities

The architecture is a conceptual synthesis rather than an empirically validated decision system. Its domains may overlap, their relative importance may vary by modality, and formal separation could create administrative burden or false precision if applied mechanically.

Self-driving laboratories can accelerate iterative experiment selection and execution [27]. However, faster experimentation cannot repair invalid endpoints, biased datasets, poorly characterised materials, or weak links between experimental outputs and patient consequences.

Advanced platforms also face dissemination barriers involving standardisation, access, interoperability, expertise, and workflow integration [28]. These limitations caution against interpreting technical availability as translational or implementation readiness.

Research should test whether domain-resolved gates improve decision calibration, transparency, reversibility, and detection of hidden dependencies. Comparative evaluations should examine computational, experimental, manufacturing, clinical, and implementation contexts while allowing the architecture to be falsified, simplified, or rejected.

CONCLUSION

The proposed Translational Decision Architecture distinguishes technically credible platform performance from product-specific, biological, pharmacological, clinical, and implementation evidence required for patient-relevant decisions. Its central contribution is a conditional and reversible structure built around bounded evidence warrants, non-substitutability, uncertainty-weighted gates, readiness regression, stopping, and evidence renewal. The architecture does not establish therapeutic effectiveness, regulatory acceptability, universal applicability, or deployment readiness. Its value depends on prospective validation demonstrating that it improves the quality, transparency, and revisability of pharmaceutical decisions without imposing unjustified complexity or delaying beneficial translation.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Wong CH, Siah KW, Lo AW. Estimation of clinical trial success rates and related parameters. Biostatistics. 2019;20(2):273-86. doi:10.1093/biostatistics/kxx069
2. Morgan P, Brown DG, Lennard S, Anderton MJ, Barrett JC, Eriksson U, et al. Impact of a five-dimensional framework on R&D productivity at AstraZeneca. Nat Rev Drug Discov. 2018;17(3):167-81. doi:10.1038/nrd.2017.244
3. Sun D, Gao W, Hu H, Zhou S. Why 90% of clinical drug development fails and how to improve it? Acta Pharm Sin B. 2022;12(7):3049-62. doi:10.1016/j.apsb.2022.02.002
4. Emmerich CH, Gamboa LM, Hofmann MCJ, Bonin-Andresen M, Arbach O, Schendel P, et al. Improving target assessment in biomedical research: The GOT-IT recommendations. Nat Rev Drug Discov. 2021;20(1):64-81. doi:10.1038/s41573-020-0087-3
5. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463-77. doi:10.1038/s41573-019-0024-5
6. Low LA, Mummery C, Berridge BR, Austin CP, Tagle DA. Organs-on-chips: Into the next decade. Nat Rev Drug Discov. 2021;20(5):345-61. doi:10.1038/s41573-020-0079-3
7. Austin CP. Opportunities and challenges in translational science. Clin Transl Sci. 2021;14(5):1629-47. doi:10.1111/cts.13055
8. Subbiah V. The next generation of evidence-based medicine. Nat Med. 2023;29(1):49-58. doi:10.1038/s41591-022-02160-z
9. Schuhmacher A, Brieke C, Gassmann O, Hinder M, Hartl D. Systematic risk identification and assessment using a new risk map in pharmaceutical R&D. Drug Discov Today. 2021;26(12):2786-93. doi:10.1016/j.drudis.2021.06.015
10. Munafò MR, Nosek BA, Bishop DVM, Button KS, Chambers CD, Percie du Sert N, et al. A manifesto for reproducible science. Nat Hum Behav. 2017;1(1):0021. doi:10.1038/s41562-016-0021
11. Percie du Sert N, Hurst V, Ahluwalia A, Alam S, Avey MT, Baker M, et al. The ARRIVE guidelines 2.0: Updated guidelines for reporting animal research. PLoS Biol. 2020;18(7). doi:10.1371/journal.pbio.3000410
12. Errington TM, Mathur M, Soderberg CK, Denis A, Perfito N, Iorns E, et al. Investigating the replicability of preclinical cancer biology. eLife. 2021;10. doi:10.7554/eLife.71601
13. Jain T, Sun T, Durand S, Hall A, Houston NR, Nett JH, et al. Biophysical properties of the clinical-stage antibody landscape. Proc Natl Acad Sci U S A. 2017;114(5):944-9. doi:10.1073/pnas.1616408114
14. Raybould MIJ, Marks C, Krawczyk K, Taddese B, Nowak J, Lewis AP, et al. Five computational developability guidelines for therapeutic antibody profiling. Proc Natl Acad Sci U S A. 2019;116(10):4025-30. doi:10.1073/pnas.1810576116
15. Nasr MM, Krumme M, Matsuda Y, Trout BL, Badman C, Mascia S, et al. Regulatory perspectives on continuous pharmaceutical manufacturing: Moving from theory to practice: September 26-27, 2016, International Symposium on the Continuous Manufacturing of Pharmaceuticals. J Pharm Sci. 2017;106(11):3199-206. doi:10.1016/j.xphs.2017.06.015
16. King EA, Davis JW, Degner JF. Are drug targets with genetic support twice as likely to be approved? Revised estimates of the impact of genetic support for drug mechanisms on the probability of drug approval. PLoS Genet. 2019;15(12). doi:10.1371/journal.pgen.1008489
17. Vlachogiannis G, Hedayat S, Vatsiou A, Jamin Y, Fernández-Mateos J, Khan K, et al. Patient-derived organoids model treatment response of metastatic gastrointestinal cancers. Science. 2018;359(6378):920-6. doi:10.1126/science.aao2774
18. Jang KJ, Otieno MA, Ronxhi J, Lim HK, Ewart L, Kodella KR, et al. Reproducing human and cross-species drug toxicities using a Liver-Chip. Sci Transl Med. 2019;11(517). doi:10.1126/scitranslmed.aax5516
19. Herland A, Maoz BM, Das D, Somayaji MR, Prantil-Baun R, Novak R, et al. Quantitative prediction of human pharmacokinetic responses to drugs via fluidically coupled vascularized organ chips. Nat Biomed Eng. 2020;4(4):421-36. doi:10.1038/s41551-019-0498-9

20. Weaver RJ, Blomme EA, Chadwick AE, Copple IM, Gerets HHJ, Goldring CE, et al. Managing the challenge of drug-induced liver injury: A roadmap for the development and deployment of preclinical predictive models. Nat Rev Drug Discov. 2020;19(2):131-48. doi:10.1038/s41573-019-0048-x
21. van der Wouden CH, Cambon-Thomsen A, Cecchin E, Cheung KC, Dávila-Fajardo CL, Deneer VHM, et al. Implementing pharmacogenomics in Europe: Design and implementation strategy of the Ubiquitous Pharmacogenomics Consortium. Clin Pharmacol Ther. 2017;101(3):341-58. doi:10.1002/cpt.602
22. Damschroder LJ, Reardon CM, Widerquist MAO, Lowery J. The updated Consolidated Framework for Implementation Research based on user feedback. Implement Sci. 2022;17(1):75. doi:10.1186/s13012-022-01245-0
23. Chen C, Zhou X, Lavezzi SM, Arshad U, Sharma R. Concept and application of the probability of pharmacological success as a decision tool in drug development: A position paper. J Transl Med. 2023;21(1):17. doi:10.1186/s12967-022-03849-y
24. Ruberg SJ, Beckers F, Hemmings R, Honig P, Irony T, LaVange L, et al. Application of Bayesian approaches in drug development: Starting a virtuous cycle. Nat Rev Drug Discov. 2023;22(3):235-50. doi:10.1038/s41573-023-00638-0
25. Scannell JW, Bosley J, Hickman JA, Dawson GR, Truebel H, Ferreira GS, et al. Predictive validity in drug discovery: What it is, why it matters and how to improve it. Nat Rev Drug Discov. 2022;21(12):915-31. doi:10.1038/s41573-022-00552-x
26. Fernando K, Menon S, Jansen K, Naik P, Nucci G, Roberts J, et al. Achieving end-to-end success in the clinic: Pfizer's learnings on R&D productivity. Drug Discov Today. 2022;27(3):697-704. doi:10.1016/j.drudis.2021.12.010
27. Häse F, Roch LM, Aspuru-Guzik A. Next-generation experimentation with self-driving laboratories. Trends Chem. 2019;1(3):282-91. doi:10.1016/j.trechm.2019.02.007
28. Alver CG, Drabbe E, Ishahak M, Agarwal A. Roadblocks confronting widespread dissemination and deployment of organs on chips. Nat Commun. 2024;15(1):5118. doi:10.1038/s41467-024-48864-3