AI as a Biological Primitive

Individual-First Biological Simulation and Support-Constrained Inverse Interrogation

M. Antony Ewing, PhD

Conquer Medical Health; University of Illinois Chicago

ORCID: 0009-0005-1550-1718 maewing@alumni.princeton.edu

September 2026

Abstract

Digital twins commonly construct a patient-specific computational object by conditioning or calibrating response structure learned substantially from external populations. AI as a Biological Primitive (ABP) tests a different direction of construction: a focal biological system's own longitudinal states, perturbations, contexts, and realized responses materially define the intervention-facing artificial primitive used for virtual interrogation. Using 194,010 five-minute observations from 12 people with type 1 diabetes, the original inversion experiment reduced mean target gap by 76.5% for activity-only, 60.6% for insulin-only, 87.8% for clinically constrained combined, and 96.4% for unconstrained combined routes; compact structured vectors retained 79.0% reduction. Separately instantiated individual and population primitives had mean response-vector cosine similarity 0.078 and identical inverse solutions in only 24.49% of evaluations. Across 3,739 observed therapeutic perturbations, median subject-level Curvature Shift was 1.85-fold and Deformation Risk 1.457-fold higher than matched within-person stable references, with discovery/replication support. Among 4,835 nonzero virtual intervention-coordinate changes, 78.0% had adequate same-person historical support; among grounded directions, 65.6% were historically favorable and 9.8% unfavorable. Selected intervention vectors changed across adjacent states in 65.5% of comparisons and across target horizons at the same state in 74.4%. Cross-scale validation includes COVID-19 population dynamics and approximately 1.2 million Allen Brain Observatory stimulus-response observations. The architecture is non-causal and support-constrained: candidate interventions are virtual experimental hypotheses requiring physical or clinical validation, and unsupported regions can remain unresolved.

Validation scope. ABP is validated as an individual-first experimental representation and inverse-interrogation architecture. Same-person support is historical evidence, not proof of treatment effect. All current intervention findings are retrospective and non-causal.

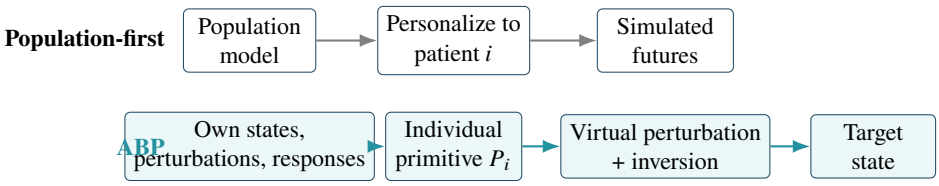

Figure 1: Architectural distinction. Population models may assist ABP, but the individualized embodiment makes the focal unit's own perturbation-response history materially determine the intervention-facing primitive.

Contents

1 From patient-as-score to patient-as-experimental-primitive

Conventional biomedical AI commonly maps observed variables to labels, risks, or future outcomes. A digital twin extends prediction into simulation but often retains a population-to-individual direction: learn a response structure from many units, then personalize it. ABP asks whether a sufficiently observed focal biological system can instead become its own artificial experimental primitive.

The distinction is architectural, not terminological. External cohorts, pretrained models, mechanistic models, and priors may assist initialization, regularization, representation learning, uncertainty estimation, or support assessment. They need not define the focal unit's local intervention-response identity.

2 Operational object

For focal biological unit i, ABP distinguishes state $x_i(t)$, intervention or perturbation $u_i(t)$, context $\phi_i(t)$, regime $R_i(t)$, and target x_i^*. A public response abstraction is

$$\widehat{x}_{i,t+\Delta} = h_i(H_{x,i}, H_{u,i}, H_{\phi,i}, R_i).$$

Target-directed interrogation reverses the scientific question:

$$u_i^* \in \arg\min_{u \in \mathcal{U}_i} d\big(h_i(H_{x,i}, H_{u,i} \cup u, H_{\phi,i}, R_i), x_i^*\big)\,.$$

The output is a model-indicated candidate, not a causal treatment effect. The production construction of $\mathcal{U}_i$, search scheduling, ranking, tie-breaking, and support calibration are not disclosed.

3 Support before recommendation

A mathematically attractive inverse solution can be biologically unsupported. ABP therefore separates optimization from empirical support and admissibility. Public conceptual states include supported-favorable, supported-mixed/qualified, supported-unfavorable, unsupported/unresolved, and abstain. Same-person support indicates that the focal unit's own history contains comparable perturbation-response evidence; it does not make that evidence causal.

4 Public Python interface

```
!pip install -q \
  git+https://github.com/ConquerMedical/ABP.git

from abp import ABP

system = ABP()
result = system.evaluate(
    substrate="type1_diabetes",
    analysis="individual_vs_population"
)
print(result.to_dict())
```

Candidate-support evidence is also callable:

```
result = system.evaluate(
    substrate="type1_diabetes",
    analysis="candidate_support"
)
print(result.same_person_support)
print(result.eligible_fraction)
```

The SDK returns frozen released evidence only. It does not instantiate a new production biological primitive from arbitrary user data and does not perform protected inverse optimization.

5 Data substrates and validation design

The primary individual-first substrate is OhioT1DM: 194,010 five-minute observations from 12 people with type 1 diabetes. Six 2018 participants (559, 563, 570, 575, 588, 591) form the discovery cohort; six 2020 participants (540, 544, 552, 567, 584, 596) form the independent replication cohort. COVID-19 population dynamics provide 56,487 population-day observations for cross-scale inversion, and the neuronal embodiment uses approximately 1.2 million Allen Brain Observatory stimulus-response observations in an analyzed 400-cell subset.

6 Validation results

6.1 Original scalar and structured inversion

The initial diabetes experiment interrogated 400 high-glucose states. The cross-scale COVID experiment used analogous route-specific and structured inversion.

Table 1: Original target-directed inversion results.

Domain / route	Mean target-gap reduction	Cases improved
Diabetes: activity only	76.5%	97.25%
Diabetes: insulin only	60.6%	97.50%
Diabetes: clinically constrained combined	87.8%	96.00%
Diabetes: unconstrained combined	96.4%	99.75%
Diabetes: compact structured vectors	79.0%	93.75%
Diabetes: full sensitivity vectors	81.2%	93.75%
COVID-19: combined scalar route	25.1%	99.75%
COVID-19: compact structured vectors	85.9%	95.00%

COVID route-specific scalar target-gap reductions were 17.9% for lockdown, 5.8% for testing, 3.8% for vaccination, and 25.1% for the combined route. Thirty-six of 39 compact COVID vectors were supported by the available substrate, and compact inversion was within approximately 0.5 percentage points of the larger sensitivity construction. In diabetes, 20 of 46 compact vectors were usable; 26 were skipped because they contained components absent from the Ohio substrate.

6.2 Individual and population primitives differ

Across the 12 separately instantiated people, the mean individual-versus-population response-vector distance was 0.1372 and mean cosine similarity was 0.078. Mean divergence between individualized and population inverse solutions was 0.4936, and only 24.49% of inverse solutions were identical. The result tests intervention-facing geometry, not merely predictive fit.

6.3 Observed therapeutic perturbations deform within-person trajectory relations

The deformation analysis identified 3,739 supported therapeutic perturbations, 16,533 stable within-person reference windows, and 3,739 matched event-control pairs. Curvature Shift (CS) measures the absolute change in pre- versus post-event trajectory slope; Deformation Risk (DR) measures mean absolute post-event displacement from the continuation of the pre-event relation.

Table 2: Replicated within-person deformation.

Analysis	Events	CS event	CS stable	CS ratio	DR event	DR stable	DR ratio
2018 discovery	1,775	0.614	0.326	1.847	21.32	14.46	1.411
2020 replication	1,964	0.585	0.269	1.857	21.32	12.67	1.578
All 12	3,739	0.614	0.308	1.850	21.32	13.00	1.457

In the 2018 cohort, CS was higher in 5/6 subjects (paired one-sided Wilcoxon $p = 0.03125$) and DR in 6/6 ($p = 0.015625$). In the independent 2020 cohort, both CS and DR were higher in 6/6 subjects ($p = 0.015625$ for each). Across all 12, CS was higher in 11/12 ($p = 0.000488$) and DR in 12/12 ($p = 0.000244$). Subject-level CS and DR ratios were strongly aligned in the

all-12 analysis (Spearman $\rho = 0.909$, $p = 4.19 \times 10^{-5}$). These are observed deformation diagnostics, not causal treatment effects.

6.4 Same-person grounding of virtual intervention directions

The inversion-to-history linkage contained 12,960 coordinate evaluations; 8,125 were no-change cases and 4,835 were nonzero proposed coordinate changes. Of those nonzero directions, 3,772 (78.0%) had adequate same-person historical support.

Table 3: Same-person grounding by discovery/replication cohort. Favorable and unfavorable percentages are among grounded directions.

Cohort	Nonzero	Grounded	Favorable	Unfavorable
2018 discovery	2,413	69.3%	50.5%	14.1%
2020 replication	2,422	86.7%	77.6%	6.3%
All 12	4,835	78.0%	65.6%	9.8%

Among all grounded directions, 2,473 (65.6%) were historically favorable, 772 (20.5%) mixed, 368 (9.8%) unfavorable, and 159 (4.2%) supported without adequate high-glucose context.

6.5 Candidate triage

Across 4,320 state-by-intervention-space candidate experiments, 1,163 produced no change and 3,157 were nontrivial. Of the nontrivial experiments, 1,270 were favorable on all changed coordinates, 402 were favorable with residual uncertainty, 368 were rejected as historically unfavorable, 749 were unresolved for lack of support, and 368 were supported but not directionally positive. Thus 1,672 of 3,157 nontrivial experiments (52.96%) met the defined development-eligibility rule.

At least one biologically grounded candidate was available for 41.0% of adverse-state queries in the discovery cohort, 73.5% in the replication cohort, and 57.1% overall. The unconstrained winning candidate passed the screen in 55.2% of all queries. Among queries with an eligible candidate, the screen changed the winner in only 3.34% of cases; median within-query target-redirection cost was zero under the reported metric.

6.6 State- and objective-specific re-interrogation

The individual primitive did not behave as a static patient-type classifier. Across patient/horizon/space sets, the median number of distinct selected intervention vectors was 7.0 overall (7.5 in 2018; 7.0 in 2020), and 77.8% of sets produced more than one solution. Selected intervention vectors changed across adjacent states in 65.5% of comparisons. Changing target horizon at the same state changed the selected solution in 74.4% of 1,440 comparisons. The supported interpretation is that the primitive must be re-interrogated as biological state or objective changes.

6.7 Cross-scale evidence

Reported forward fits were $R^2 = 0.934$ in diabetes and $R^2 = 0.912$ in COVID-19 population dynamics. Compact structured inversion retained 79.0% target-gap reduction in diabetes and 85.9% in COVID-19. The neuronal embodiment used approximately 1.2 million Allen Brain Observatory cell-stimulus observations. These experiments support scale generality of the computational object, not common biological mechanism across domains.

7 Drug-development interpretation

The intervention vector can conceptually be expanded to

$$u = (\text{compound}, \text{dose}, \text{schedule}, \text{combination}),$$

allowing support-constrained screening, dose/schedule interrogation, combination testing, trial simulation, and adaptive re-screening before physical validation. Current evidence demonstrates mechanics within observed therapeutic or perturbational spaces. It does not establish de novo molecule generation, causal efficacy, or clinical treatment recommendation.

8 Implementation boundary and limitations

The public release excludes production focal-unit primitive compilation, response-operator training pipelines, candidate-grid construction, admissible-set generation, same-unit support calibration, target-directed inverse optimization, search scheduling/ranking/tie-breaking, private drug-development screening logic, and deployment orchestration. All current evidence is retrospective. Intervention choice is endogenous in longitudinal physiology, and support-grounded inversion does not establish a treatment effect. The individual-first result has been demonstrated in a densely observed substrate and does not imply that every patient or disease contains enough within-person information for separate instantiation; when information is inadequate, the appropriate state is insufficient or unresolved.

9 Public artifacts and conclusion

The repository is https://github.com/ConquerMedical/ABP; the interactive demonstrator is https://conquermedical.github.io/ABP/. ABP's central validated distinction is the direction of biological construction: when a focal unit is observed deeply enough, its own realized perturbation history can materially define the experimental object used for virtual interrogation. The empirical program further shows that this object differs from a population-conditioned response surface, changes as state and objective change, and can be constrained by same-person historical support before candidate advancement.

References

[1] Marling C, Bunescu R. The OhioT1DM Dataset for Blood Glucose Level Prediction: Update 2020. *CEUR Workshop Proceedings*. 2020;2675:71–74.
[2] Hale T, Angrist N, Goldszmidt R, et al. A global panel database of pandemic policies. *Nature Human Behaviour*. 2021;5:529–538.
[3] de Vries SEJ, Lecoq JA, Buice MA, et al. A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature Neuroscience*. 2020;23:138–151.
[4] Katsoulakis E, Wang Q, Wu H, et al. Digital twins for health: a scoping review. *npj Digital Medicine*. 2024;7:77.
[5] Laubenbacher R, Mehrad B, Shmulevich I, et al. Digital twins in medicine. *Nature Computational Science*. 2024;4:184–191.
[6] Venkatapurapu SP, Clegg L, Nowojewski A, et al. Digital twins for accelerating drug discovery and development: opportunities and challenges. *Drug Discovery Today*. 2026;31(2):104617.

Intellectual property disclosure

Relevant methods and computational systems are the subject of previously filed United States provisional patent applications. This paper discloses the mathematical objects, empirical definitions, validation logic, public interfaces, and results necessary to evaluate the reported scientific analyses. Proprietary operational implementations and system components not necessary to reproduce or evaluate the reported scientific claims remain outside the scope of disclosure. Publication of this paper does not grant a patent license or a commercial software license.

Competing interests

The author is Founder of Conquer Medical Health and has financial and intellectual-property interests in the technologies described here. Relevant methods and systems are patent pending. The author declares no other competing interests.

Funding

No external funding supported preparation of this paper.

Author contributions

M.A.E. conceived the architecture, developed the methods and software, designed and conducted the reported validation program, interpreted the results, and wrote the paper.