Supplementary Information

Testing predictive-state sufficiency and recursive closure in additive-manufacturing microstructure models

Junwen Ji[1,*], Jie Liu[2], Anatoliy Zavdoveev[3], Viacheslav Kopylov[1]
[1] National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute", 35 Polytechnic St., 03056 Kyiv, Ukraine
[2]Nanjing University of Industry Technology, Nanjing 210023, China
[3]Paton Electric Welding Institute of the NAS of Ukraine, 11 K. Malevycha Street, 03150 Kyiv, Ukraine
[*]Correspondence: jijunwen343j@gmail.com

Supplementary Note 1: Theoretical inheritance and operational claim

The theory does not claim a new meaning for the word "state". It draws on four established lines of thought. Internal-variable thermodynamics supplies the constitutive precedent for finite evolving memory. Predictive-state theory defines histories by the futures they induce; information bottlenecks formalize task-relevant compression; and recurrent minimal-state models show how latent variables can carry path-dependent constitutive response. The contribution is the operational testing framework that connects these foundations to additive-manufacturing microstructure. Shared continuations test sufficiency, concatenated segments test recursive update, dimension plateaus test redundancy, and source transfer reveals when a representation ceases to be complete.

Supplementary Table 1. Relationship between the theory and its foundations.

Foundation	Established contribution	Operational use here
Internal variables	Finite evolving variables represent history-dependent constitutive response.	Candidate metallurgical variables must also pass held-out continuation and update tests.
Predictive or causal states	Histories are equivalent when they induce the same conditional future.	Equivalence is conditioned on alloy, target, scale, forcing domain and tolerance.
Information bottleneck	Compression retains information relevant to a prediction target.	A coordinate is retained only when it lowers held-out continuation risk beyond the declared tolerance.
Recurrent minimal-state models	Compact latent variables can reproduce path-dependent response.	The latent representation must accept the next thermal segment and survive source and order holdouts.

The definition is task-conditioned. Invertible coordinate transformations preserve predictive equivalence, so no single semantic basis is privileged. A state also need not be a short local vector. Solute transport, grain competition, stress redistribution and defect inheritance may require nonlocal or field-valued information. If no stable dimension plateau appears, the experiment has bounded the domain in which finite closure is credible.

Supplementary Note 2: Order and reset as intervention tests

Let U_A and U_B be state-update operators associated with two admissible thermal segments. If

$$U_B(U_A(\boldsymbol{z})) \neq U_A(U_B(\boldsymbol{z})),$$

then a representation invariant to exchanging A and B cannot be sufficient for a target that distinguishes the resulting states. This is a no-go diagnostic, not a new theorem about path dependence: non-commutation is already its mathematical statement.

The dimensionless example in Supplementary Fig. 1 isolates the logic. Starting from $x_0 = 0.5$, define dissolution $H(x) = x\exp(-2)$ and formation $M(x) = 1 - (1 - x)\exp(-2)$. The two orders give $M(H(x_0)) = 0.874$ and $H(M(x_0)) = 0.126$. Segment content, duration, extrema, histogram and time integral are unchanged. The example proves only that an order-erasing descriptor fails for this update system.

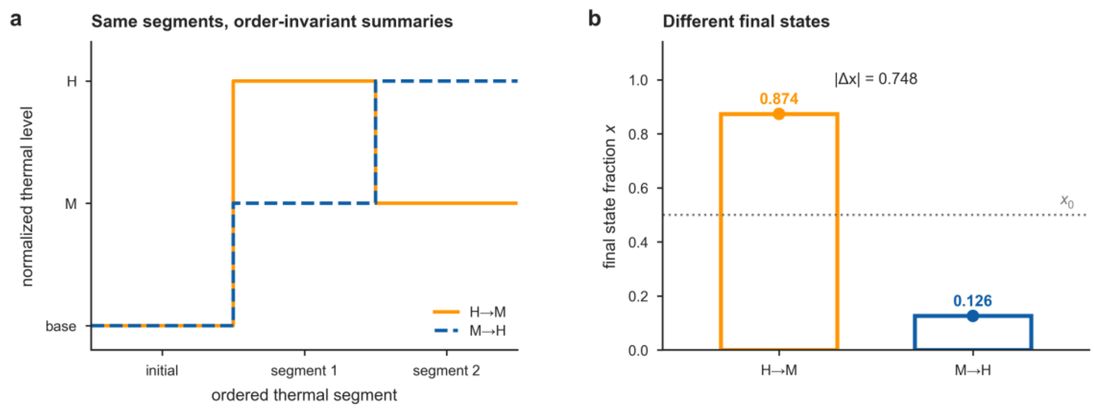

Supplementary Fig. 1. Order sensitivity as a diagnostic. (a) The histories contain identical high- and intermediate-temperature segments in opposite order. (b) Non-commuting dissolution and formation maps produce different terminal fractions. The model is dimensionless and is not fitted to an alloy.

Reset is similarly conditional. A high-temperature excursion can erase phase partition without erasing parent-grain dimensions, chemistry or defect lineage. The schematic in Supplementary Fig. 2 therefore treats parent phase, transformation progress, morphology and martensite as hypotheses to be admitted or rejected by response tests, not as a universal checklist.

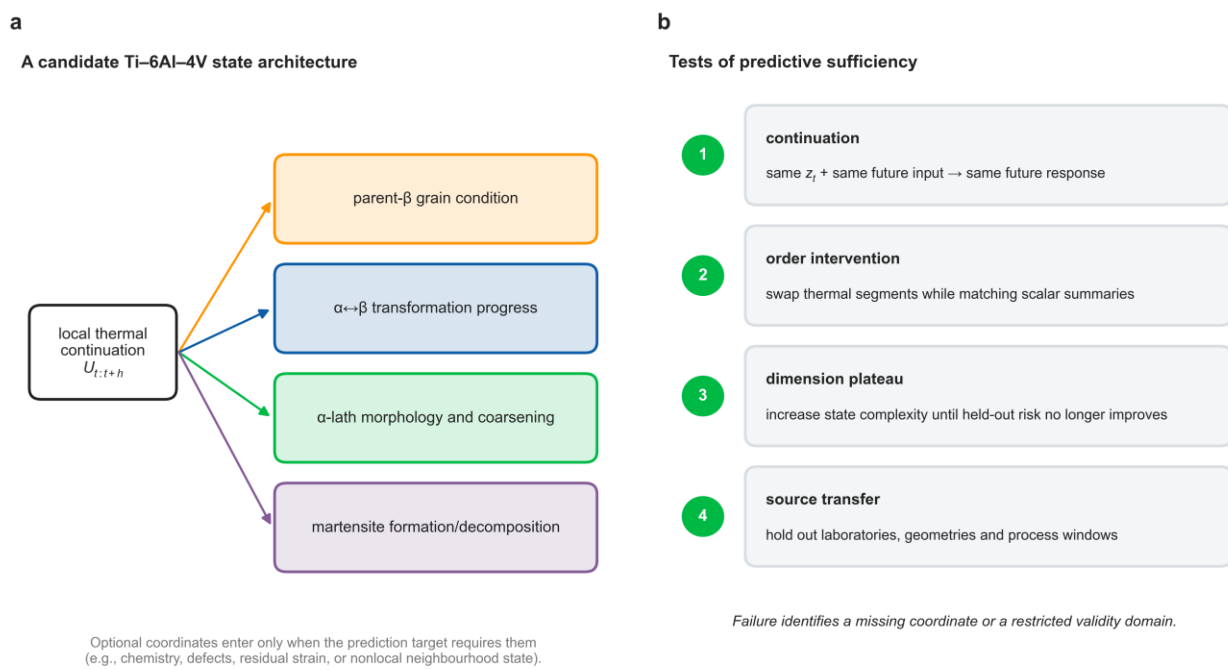

Supplementary Fig. 2. Hypothetical material-state structure and its tests. (a) A Ti–6Al–4V prediction task may require parent-beta condition, transformation progress, alpha-lath morphology and martensite; other targets may require chemistry, defects or nonlocal fields. (b) Shared continuations, order interventions, dimension plateaus and source transfer diagnose missing or redundant information. Only phase-state coordinates are exercised in the present calculations.

Supplementary Note 3: Oracle calibration, order pair and selective reset

The Nitzler phase state was made available only in a calibration benchmark. Direct propagation then reproduces the stated simulator response, whereas descriptors that erase order need not do so. This controlled comparison is a closure reference for the source model; it is not a claim that two variables define a universal material ontology.

The order and reset controls retain the same segment multiset while reversing its temperature order, then apply a common continuation. They show why an order-erasing summary can lose predictive information. The selective-reset result concerns only the phase variables represented by the source model and is not a general statement about all microstructural memory.

The oracle descriptor control uses a two-component phase target and a single continuation, whereas the main recursive experiment forecasts total-alpha response signatures across six segments. The two settings answer different questions and are therefore reported separately.

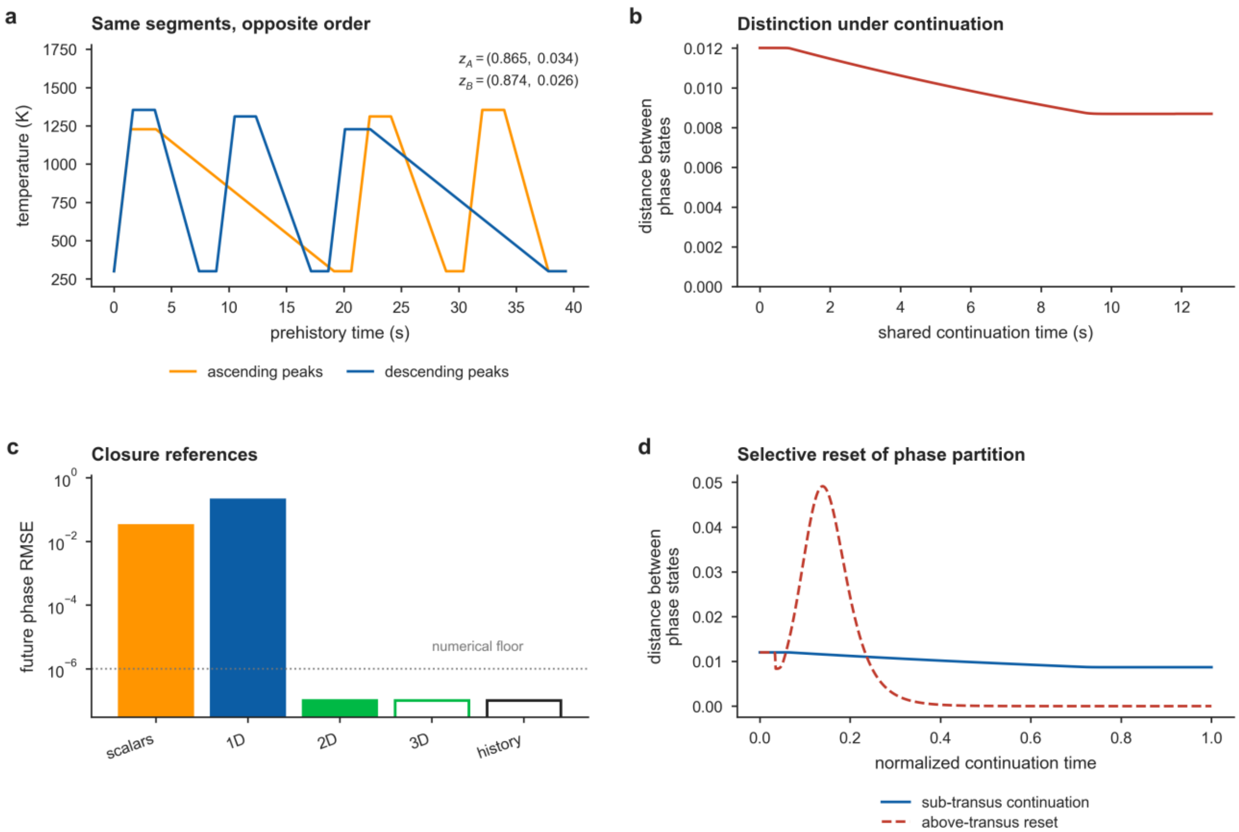

Supplementary Fig. 3. Oracle and intervention controls for the source model. (a) The same thermal segments in opposite order. (b) Their distinction under a shared continuation. (c) Direct state propagation and full-history prediction are shown as oracle references for the stated simulator. (d) Selective reset of the represented phase partition. These are in silico controls, not experimental validation.

Supplementary Result 1: Dimension robustness and full-history sequence baseline

A post-hoc continuation stress test sampled 248 additional continuations across the declared temperature, hold and cooling-rate domain. On this broader bank, two coordinates gave a median RMSE of 0.00127 and a 95th-percentile RMSE of 0.00329; three coordinates gave 0.00135 and 0.00337. This stress test is descriptive rather than a universal sufficiency result. Separately, a full-history GRU received the entire resampled prehistory and supplied continuation. Its held-out terminal-response RMSE was 0.07732, a 71.8% reduction relative to the source-training mean baseline. Because it rereads the history, the GRU is a forecasting baseline rather than a compact recursive state.

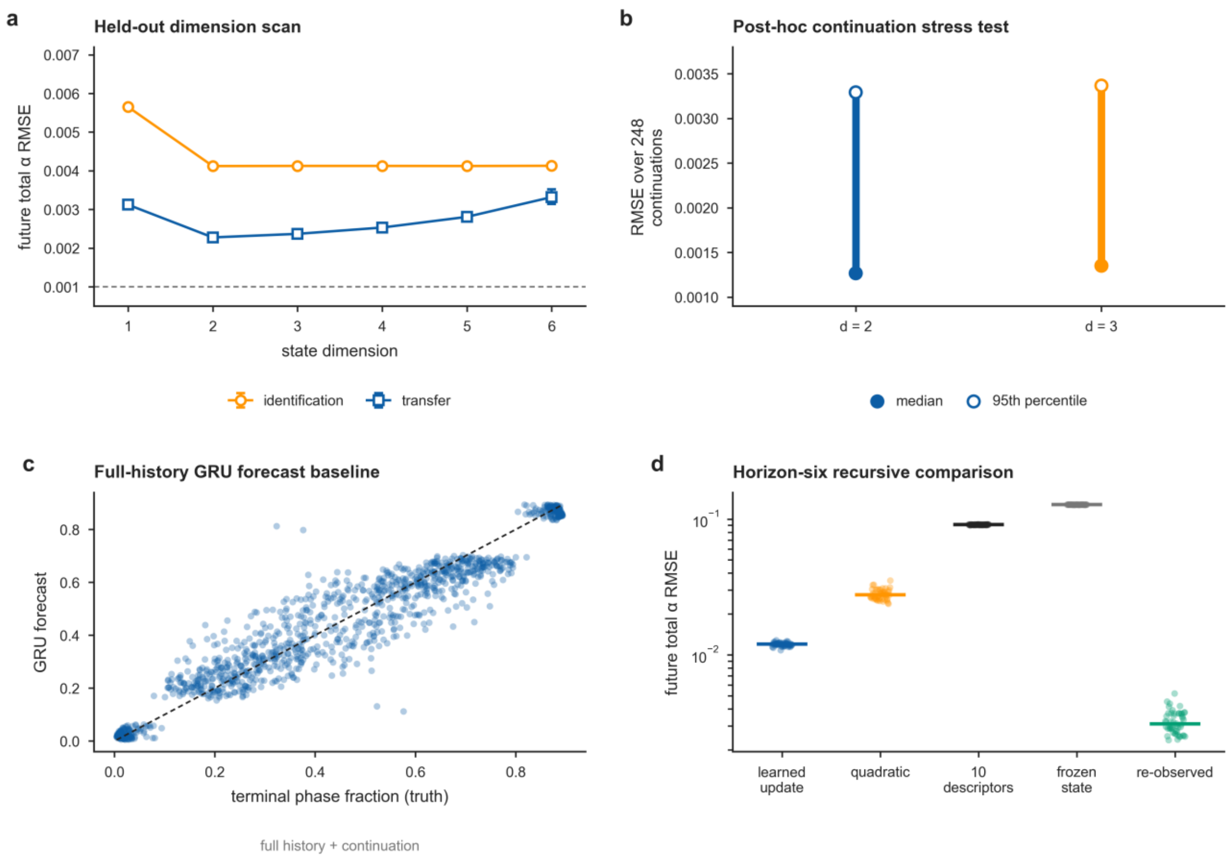

Supplementary Fig. 4. Extended robustness and sequence-baseline results. (a) Held-out identification and transfer errors over dimensions one to six. (b) Median and 95th-percentile error over 248 post-hoc continuations for dimensions two and three. (c) Full-history GRU terminal forecasts on held-out descending histories. (d) Horizon-six distributions for the selected two-coordinate recursive update and its controls. The post-hoc continuation bank is descriptive and does not alter the predeclared primary dimension-selection rule.

Supplementary Result 2: Cross-model transfer and recursive closure

With the Nitzler encoder frozen and only a Murgau response decoder calibrated, median held-out RMSE decreased from 0.006271 at one coordinate to 0.002994 at two and 0.002547 at three. The matched ten-descriptor baseline gave 0.019818, so the transferred three-coordinate representation reduced the median error by 87.1%. Strict zero-shot reuse remained near 0.10. Target-specific encoders are reported separately and do not test source-to-target transfer.

For recursive evaluation, the ten descriptors were recomputed at each boundary from the full accumulated thermal trace. At horizon six, descriptor RMSE was 0.090908 and the two-coordinate learned update gave 0.011980. The paired median reduction was 0.078896 (95% bootstrap CI, 0.078775–0.079197), equivalent to an 86.8% reduction by the ratio of medians. Direct boundary re-observation is shown separately because it observes the response at every boundary.

Supplementary Results tables

Supplementary Table 2. Observation-driven Nitzler benchmark. Values are medians over 50 observation-noise realizations; bracketed values are 95% percentile-bootstrap intervals for the seed median.

Dimension	Identification RMSE	Transfer RMSE	Post-hoc phase R^2
1	0.005651 [0.005646, 0.005659]	0.003122 [0.003108, 0.003151]	0.9711
2	0.004122 [0.004101, 0.004140]	0.002276 [0.002259, 0.002305]	0.9912
3	0.004126 [0.004106, 0.004145]	0.002369 [0.002343, 0.002411]	0.9912
4	0.004126 [0.004110, 0.004148]	0.002532 [0.002500, 0.002582]	0.9914
5	0.004123 [0.004112, 0.004150]	0.002808 [0.002729, 0.002876]	0.9914
6	0.004129 [0.004111, 0.004151]	0.003319 [0.003143, 0.003517]	0.9914

Supplementary Table 3. Cross-model source-transfer comparisons. Values are median held-out RMSE over 50 observation-noise realizations. ‘Calibrated’ freezes the Nitzler encoder and fits only a Murgau response decoder without phase labels.

Method	d = 1	d = 2	d = 3
Frozen encoder + calibrated target decoder	0.006271	0.002994	0.002547
Target-specific encoder	0.003288	0.001156	0.001177
Strict zero-shot	0.104126	0.100213	0.100652
10 thermal descriptors	0.019818	0.019818	0.019818

Supplementary Table 4. Six-segment Nitzler rollout at horizon six. Values are medians over 50 observation-noise realizations. Direct re-encoding observes the boundary response and is not recursive.

Method at horizon 6	d = 1	d = 2	d = 3
Learned recursive update	0.011127	0.011980	0.012169
Quadratic update	0.026921	0.027709	0.027332
Frozen state	0.126568	0.127957	0.128037
Direct boundary re-encoding	0.004037	0.003109	0.002926
10 thermal descriptors	0.090908	0.090908	0.090908

Supplementary Table 5. Robustness and sequence-baseline checks. The post-hoc continuation bank is descriptive; it does not replace the primary, predeclared dimension-selection analysis.

Check	Result
Murgau production time step	0.0015625 s
Post-hoc continuation bank	248 continuations across the declared domain
d = 2 continuation stress	median 0.001268; 95th percentile 0.003293
d = 3 continuation stress	median 0.001352; 95th percentile 0.003368
Full-history GRU	held-out RMSE 0.077321; 71.8% below source-training mean
GRU interpretation	forecast baseline; not a compact recursive state

Supplementary Table 6. Paired statistical details for the principal comparisons. ΔRMSE is defined as the RMSE of the first method minus that of the second method; intervals are deterministic 95% percentile-bootstrap intervals for the median paired difference.

Comparison	Median ΔRMSE	95% bootstrap CI	Holm-adjusted p
Identification d1 − d2	0.001532	[0.001512, 0.001542]	< 0.001
Identification d2 − d3	-0.000003	[-0.000005, -0.000002]	1
Transfer d1 − d2	0.000823	[0.000802, 0.000872]	< 0.001
Transfer d2 − d3	-0.000086	[-0.000096, -0.000074]	1
Cross-model d1 − d2	0.002717	[0.001333, 0.003867]	< 0.001
Cross-model d2 − d3	0.000207	[0.000063, 0.000470]	< 0.001
Descriptors − transferred d3	0.017273	[0.016876, 0.017533]	< 0.001
Recursive h6 d1 − d2	-0.000826	[-0.000919, -0.000708]	1
Recursive h6 d2 − d3	-0.000280	[-0.000457, -0.000184]	1
Descriptors − recursive d2	0.078896	[0.078775, 0.079197]	< 0.001