Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

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

Junwen Ji

jijunwen343j@gmail.com

National Technical University of Ukraine “Igor Sikorsky Kyiv Polytechnic Institute”

Jie Liu

Nanjing University of Industry Technology

Anatoliy Zavdoveev

Paton Electric Welding Institute of the NAS of Ukraine

Viacheslav Kopylov

National Technical University of Ukraine “Igor Sikorsky Kyiv Polytechnic Institute”

Article

Keywords:

Posted Date: August 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10564290/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

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

Abstract

Most additive-manufacturing models predict final microstructure from process parameters or thermal-history descriptors. We ask a different question: how much of a material's thermal history must be carried forward to predict what the next heat cycle will do to it? We call this compact, continuously updated information a predictive material state. In Ti–6Al–4V simulations we hid the phase fractions and inferred state coordinates from noisy total-alpha measurements. Two coordinates identified the state 27.1% more accurately than one, while on thermal cycles held out from training a single coordinate already predicted the response to within the imposed measurement noise: how much memory is needed depends on what is being predicted. Transferred to a separately formulated kinetic model with only the readout recalibrated, a three-coordinate state cut median prediction error by 87.1% relative to ten conventional thermal-history descriptors with access to the entire past. Carried recursively through six unseen thermal segments with no intervening measurement, a two-coordinate state stayed 86.8% below the same descriptor baseline. A recurrent network given the full history forecast well but produced no compact updateable state. Two or three numbers, refreshed as heat arrives, carry more of what matters than ten summaries of everything that has happened.

Introduction

Directed energy deposition and other metal additive-manufacturing processes build material layer by layer, repeatedly exposing it to melting, solidification, solid-state transformation, coarsening, and recovery. Process parameters such as current, voltage, travel speed, interlayer dwell time, and interpass temperature determine the external thermal input during each cycle. Once a cycle has passed, however, these parameters alone no longer describe the current condition of a material point. The complete local thermal history records what the material has experienced, but it does not directly reveal which effects of that history still govern its subsequent evolution. Predicting the response to the next thermal cycle therefore requires the relevant information to be extracted from the complex thermal history and represented as a compact, updateable material state.

Existing studies commonly use process parameters or a small set of thermal-history descriptors to predict final microstructure. However, these quantities do not fully represent the current state of a material point. Similar process parameters can produce different local thermal histories because of differences in geometry, boundary conditions, and heat accumulation. Conversely, distinct thermal histories can have similar peak temperatures or cooling rates while producing different phase compositions and microstructural morphologies [1]. The mechanism is order. Build two histories from the same pair of thermal segments applied in opposite sequence: every order-invariant summary of them — peak, mean, duration, time above any threshold, integrated thermal exposure — is identical, yet because dissolution and formation do not commute the two end at transformed fractions of 0.874 and 0.126 (Supplementary Fig. 1). The central limitation is therefore not a lack of additional process parameters or thermal descriptors, but the absence of a material-state representation that retains the relevant effects of prior thermal exposure and can be updated as new cycles occur.[2–4] This problem is particularly pronounced in wire-fed arc deposition, where successive layers impose strong, non-monotonic, and position-dependent thermal cycles.[5–8]

This motivates a more specific question: how much information from previous thermal history must be retained to predict a material point's response to subsequent thermal input? We call this compact, updateable information a predictive material state. It can be viewed as a low-dimensional memory of thermal history: after each new thermal cycle, the state is updated and then used to predict the material response under future thermal input. Two histories are considered predictively equivalent, within a specified tolerance and input domain, when they produce approximately the same responses under the same future thermal continuations.

The idea that a few evolving variables can stand in for a long history is not new. Internal-state-variable thermodynamics has represented history-dependent constitutive behaviour this way for decades [9–11], and microstructural process models compress it into a few evolving phase variables [12–15]. Predictive-state and information-bottleneck approaches instead group histories by the futures they induce, or compress them towards a prediction target [16–19], while learned state-space and Koopman models seek coordinates with tractable evolution laws [20–23]. Neural differential equations and neural operators add flexible machinery for learning such dynamics [24, 25], and data-driven models have been applied directly to additive-manufacturing thermal histories and properties [26–30]. What none of them supplies is a way to test a proposed state against a specific prediction task. They select representations from physical assumptions, model architecture, reconstruction error or overall predictive performance, none of which reveals whether a state has quietly merged two histories that a later thermal input would separate.

Here, predictive material state is an operational concept tied to a prediction task rather than a new thermodynamic state variable. We identify candidate states from observable future responses, choose their dimension from prediction error on held-out histories, and then test sufficiency, transfer to a separately formulated kinetic model, and recursive update closure as three separate properties. Oracle propagation, thermal-cycle order permutations and state-reset controls in Supplementary Note 3 and Supplementary Fig. 3 calibrate these tests. Figure 1 summarizes the future-response equivalence test and the recursive state-update procedure.

Results

What a state must do: sufficiency, closure and minimality

Consider a material point or representative volume with fixed attributes$c$. Let$\mathcal{H}_t$ denote the admissible forcing and observations available before time t,$\mathcal{U}_{t:t+h}$ a future thermal or thermomechanical continuation, and $Y_{t:t+h}$ the response to be predicted. Two histories belong to the same predictive state only if they produce equivalent conditional responses under every admissible common continuation:

$$\mathcal{H}_t^{(1)} \sim_Y \mathcal{H}_t^{(2)} \Leftrightarrow P(Y_{t:t+h} \mid \mathcal{H}_t^{(1)}, \mathcal{U}, c) = P(Y_{t:t+h} \mid \mathcal{H}_t^{(2)}, \mathcal{U}, c), \qquad (1)$$

Equation (1) gives the ideal definition. In practice, exact equality is replaced by a predeclared discrepancy measure and a tolerance that reflects measurement and modelling uncertainty. The map $\boldsymbol{z}_t = \Phi(\mathcal{H}_t, c)$ assigns a common state coordinate to histories that satisfy this future-response equivalence test. The resulting state is therefore task-dependent: it varies with the alloy, spatial scale, admissible future-input domain, response target, and tolerance. For example, a state sufficient for predicting room-temperature phase fraction may not be sufficient for predicting morphology or fatigue response after another heat cycle.

Process variables such as current, voltage, travel speed, and interpass time determine the thermal forcing applied to the material. A phase fraction, grain descriptor, or learned latent coordinate should enter the state only when it retains information needed to predict what the material will do next. A predictive state must also be updateable when a new input arrives, without requiring the complete history to be read again:

$$\boldsymbol{z}_{n+1} = F(\boldsymbol{z}_n, u_n, c), \qquad Y_n = G(\boldsymbol{z}_n, c) + \eta_n, \qquad (2)$$

Here, zn is the current state, un is the next applied input, and ηn represents observation or modelling error. This recursive form distinguishes a material state from a terminal embedding: an embedding may predict well at one boundary but still fail to provide a closed update through the next thermal segment.

A candidate state has to pass three tests, and it can pass one while failing another. It is sufficient if two histories it treats as equivalent really do behave the same way under the same next heat cycle; if any single continuation separates them, the state has discarded memory it needed. It is closed if it can be carried through that cycle using only the applied temperature, without stopping to measure the material again — an embedding can predict well at a boundary and still have no update rule. It is minimal if no coordinate can be dropped without pushing held-out error past a tolerance fixed in advance. A task-level state architecture for Ti–6Al–4V is shown in Supplementary Fig. 2. Supplementary Note 1 and Supplementary Table 1 relate these requirements to

internal-variable thermodynamics, predictive-state theory and information-bottleneck methods; the order, reset and oracle controls are described in Supplementary Notes 2–3.

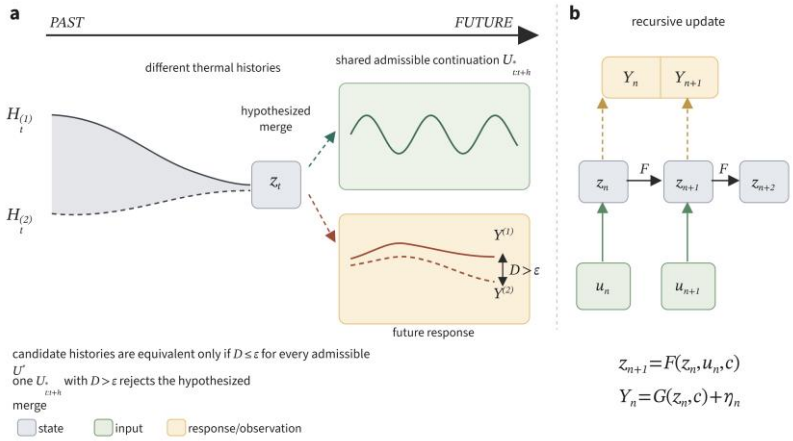


Figure 1. Future-response equivalence and recursive material-state update. a, Two thermal histories are provisionally assigned to the same candidate state. They are considered equivalent for target response Y only when every tested admissible common continuation U* produces future-response distributions whose discrepancy D does not exceed the predeclared tolerance ε. The condition D≤ε is the operational form of Eq. (1) when exact equality cannot be established because of measurement or modelling uncertainty. A single continuation with D>ε rejects the proposed merge. b, The candidate state is updated recursively from its current value, the next input, and the fixed context, and the updated state is then used to predict the specified response.

Hidden responses identify two coordinates for the diagnostic task

We first removed any route by which the identification could become circular. In the Nitzler kinetic model the stable-α and martensitic-α phase fractions were treated as hidden variables, and the state identifier saw only sparse, noisy observations of their sum (Fig. 2a). Each history was a synthetic multi-cycle temperature programme, with peaks, hold times and cooling rates spanning the ranges reported for wire-arc deposition of Ti–6Al–4V, and each was probed with a fixed panel of six standard heat–hold–cool cycles, seven total-alpha measurements being retained during every probe. Each history thus produced a 42-number response signature: the fingerprint by which two histories were judged equivalent. The model was fitted using 1,200 ascending-peak and mixed histories, while all 600 descending-peak histories were reserved for independent testing. All errors below are absolute phase fractions on a 0–1 scale, and each observation carried independent Gaussian noise of standard deviation 0.005.

Increasing the state dimension from one coordinate to two reduced the median identification RMSE from 0.00565 to 0.00412, a decrease of 27.1% that moved the estimator from above the observation noise down to the noise floor itself. The median transfer RMSE fell by the same proportion, from 0.00312 to 0.00228 (Fig. 2b). In the representative realization shown in Fig. 2c, the two-coordinate representation improved prediction for 76% of the 600 held-out transfer histories.

The two tasks disagreed about how much memory the material has. Identifying the state from noisy observations needed two coordinates. Predicting the response to thermal cycles held out from training needed only one, which already fell within the preset 0.001 RMSE margin — one fifth of the observation noise — of the best tested dimension. Neither number is the state dimension of Ti–6Al–4V; the question is incomplete until the prediction target and the tolerance are named. Paired intervals and statistical tests are reported in Supplementary Table 6.

Adding further coordinates provided no additional benefit. Median identification RMSE was 0.00412 at two coordinates and 0.00413 at three — no measurable change — while transfer RMSE increased from 0.00228 to 0.00237. Transfer error increased further for dimensions four through six. The complete comparison across one to six coordinates is given in Supplementary Table 2.

Only after predictive dimension had been evaluated were the learned coordinates aligned with the concealed phase state. The two-coordinate alignment achieved a median R^2 of 0.9912 and a phase-state RMSE of 0.01098, compared with 0.9711 and 0.01990, respectively, for the one-coordinate representation (Fig. 2d and Supplementary Table 2). The estimator never saw the phase labels, so the second coordinate is in effect the

stable-α/martensitic-α split that the total-alpha observation conceals. Any invertible transformation of the pair would align equally well, so this is not a claim to a privileged basis; it is evidence that the predictive test recovered a physically real distinction without being told it existed.

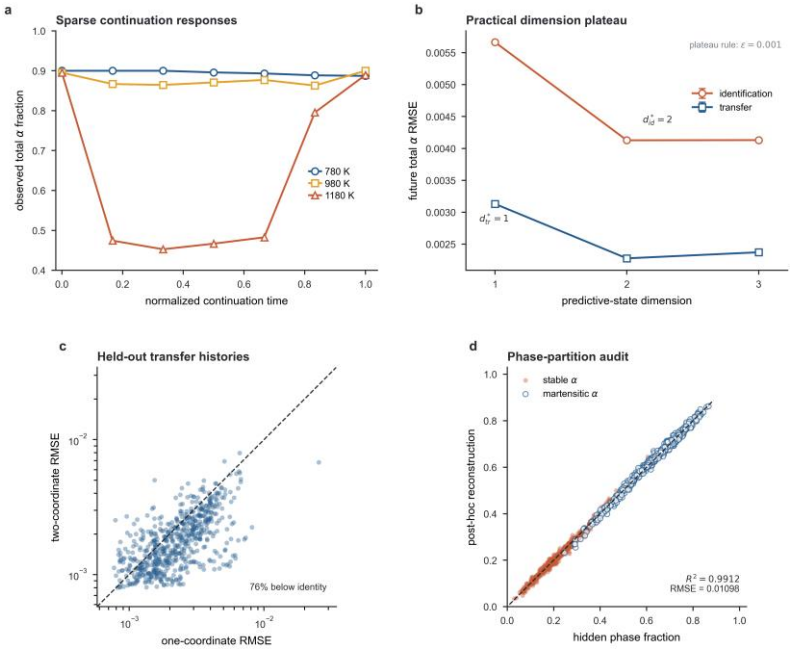


Figure 2. Observation-driven identification with the phase state hidden. a, Representative noisy total-alpha observations from the diagnostic continuations. b, Prediction error versus predictive-state dimension for hidden-state identification and transfer. Points show medians and error bars denote 95% bootstrap confidence intervals. c, Per-history transfer RMSE for one- and two-coordinate representations. d, Post-hoc alignment of the learned two-coordinate state with the hidden phase fractions.

A frozen state transfers to a different kinetic model once the readout is recalibrated

We next tested whether the learned representation retained predictive information after the underlying response law changed. The Nitzler response encoder was frozen and applied to total-alpha continuations generated by a separately formulated, literature-based Murgau-class model for Ti–6Al–4V. This target model resolves grain-boundary alpha, Widmanstätten alpha, and martensitic alpha, but none of these phase components was provided as a training label. A target-specific response decoder was calibrated from the transferred coordinates using 600 ascending-peak and mixed histories, while 300 descending-peak histories were reserved as external tests. This experiment evaluates representation transfer after limited target calibration; strict zero-shot transfer of the complete kinetic mapping was assessed separately.

With the source encoder fixed, the target continuation RMSE decreased from 0.00627 with one coordinate to 0.00299 with two coordinates and 0.00255 with three coordinates (Fig. 3a and Supplementary Table 3). Adding a second coordinate produced a median paired RMSE reduction of 0.00272 (95% bootstrap CI, 0.00133–0.00387). A third coordinate provided a smaller but consistent additional reduction of 0.000207 (0.000063–0.000470). However, at every tested dimension an encoder trained specifically on the target model performed better than the transferred encoder.

We then compared the transferred state with the way thermal history is normally summarized, using the same target-model split. Ten scalar descriptors — peak, mean and spread of temperature, total duration, integrated thermal exposure, residence times above three thresholds, maximum heating rate and cycle count — were mapped to the Murgau responses by an extremely randomized tree, again without access to target phase labels. By construction these descriptors are order-invariant: they treat the temperature trace as an unordered set, so two histories built from the same segments in different sequence receive identical inputs. Given the entire prior history in that form, the tree reached a median RMSE of 0.01982, against 0.00255 for the transferred three-coordinate state — an 87.1% reduction. Complete paired statistics are reported in Supplementary Table 6.

Strict zero-shot transfer of the complete kinetic mapping was unsuccessful. At three coordinates, the median zero-shot RMSE was 0.10065, compared with 0.00255 after calibrating the target response decoder (Fig. 3b). Post hoc alignment between the transferred coordinates and the concealed target-phase components was also weak, with a median R^2 of 0.0245 at three coordinates. The representation transferred once the readout was recalibrated; the source kinetic law itself did not.

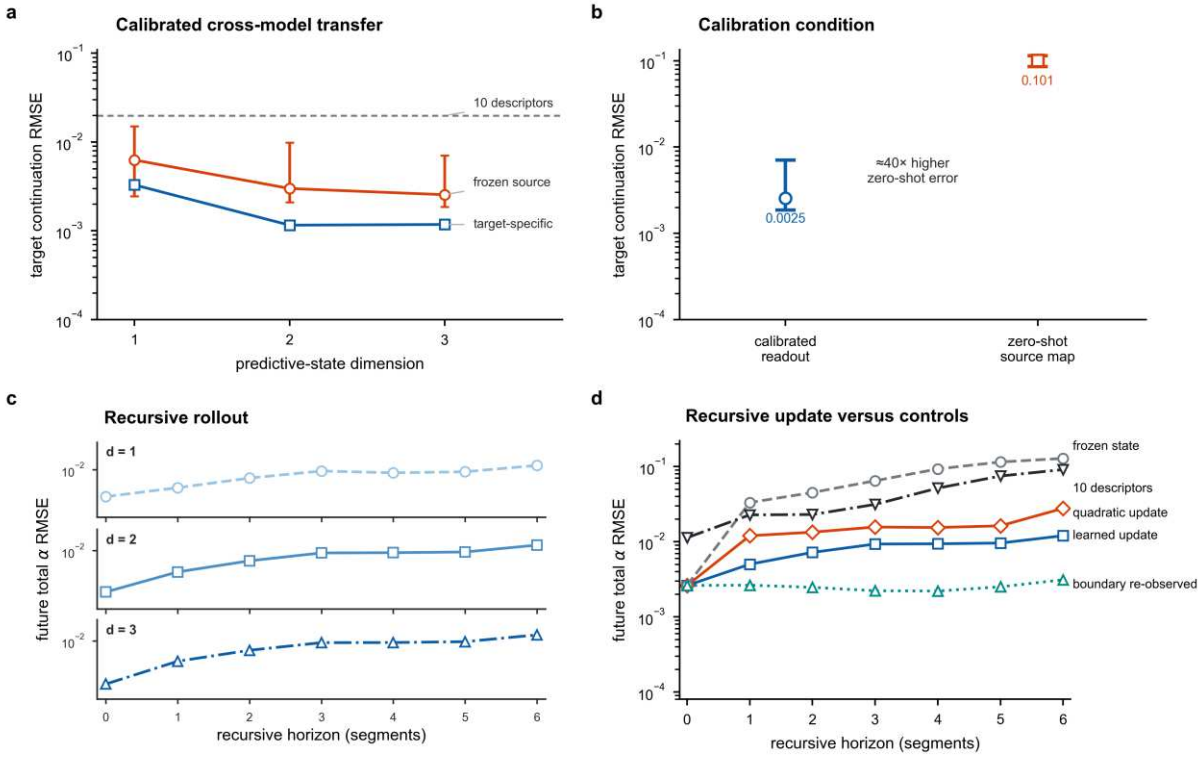


Figure 3. Source transfer and recursive update tests. a, Transfer RMSE for the frozen Nitzler encoder with a target-calibrated decoder and the target-specific encoder–decoder; the dotted line indicates the thermal-descriptor baseline. Points show medians and error bars denote 95% bootstrap confidence intervals. b, Comparison of the target-calibrated readout and the strict zero-shot source map at three coordinates. c, Response RMSE over six recursive segments for one-, two- and three-coordinate representations; shaded bands indicate 95% bootstrap confidence intervals. d, Comparison of the two-coordinate nonlinear rollout with the quadratic update, frozen state, thermal-descriptor baseline and direct boundary re-observation.

A learned update propagates state without boundary re-observation

A representation Z_d identified at a thermal-segment boundary is not yet a recursive state. To qualify as one, it must also support the update law in Eq. (2). We tested this requirement using sequences of six thermal segments and learned an update function $F(Z_n, u_n)$ that predicts the coordinate at the next boundary from the current coordinate and the next segment's peak temperature, hold time, and cooling rate.

The update model was trained on 360 ascending-order or mixed-order sequences. Tree regularization was selected by six-step rollout on an internal validation subset drawn from the training family. The 180 descending-order sequences were excluded from both model fitting and hyperparameter selection and were used only for external testing. A thermal-descriptor baseline followed the same split. At each boundary, this baseline received ten descriptors recomputed from the full accumulated thermal history, including all segments observed up to that point, but no material response measured at the boundary.

At test time the coordinate was estimated only once, before the first segment. All subsequent updates received control inputs alone; no diagnostic response or hidden phase value was provided at later boundaries. The learned update models carried the state across six unseen thermal segments. The horizon-six RMSE values were 0.01113, 0.01198 and 0.01217 for one-, two- and three-coordinate states. Adding coordinates did not help (Fig. 3c and Supplementary Table 4).

For the two-coordinate representation selected in the identification task, the validation-selected nonlinear update predicted the sixth-segment response to 0.01198 — about 2.4 times the observation noise, after six segments with no material measurement at all. A quadratic update law lost a factor of two (0.02771); freezing the state entirely lost a factor of ten (0.12796), which shows that the update rule, not the initial estimate, is doing the work. The ten-descriptor baseline, recomputed at every boundary from the whole accumulated thermal trace, reached only 0.09091 (Fig. 3d); the two-number recursive state reduced median error by 86.8% relative to it. Complete paired statistics are reported in Supplementary Table 6.

How far is this from perfect closure? Re-measuring the state at every boundary — the best any observer with a sensor at each step could do — brings horizon-six error down to 0.00311, against 0.01198 for the update running blind through all six segments. (The oracle descriptor control in Supplementary Note 3 uses a different target, history set and single continuation, and is not a comparator for this rollout.)

Discussion

Process settings remain indispensable because they determine the external thermal input passed to the state-update function F in Eq. (2). The question addressed here begins after part of that input has become history: what information must the current material state z retain to predict the next response? Once a thermal cycle has passed, prediction can no longer rely on nominal process settings alone. It must use the condition left in the material by that cycle together with the next applied input. Process settings describe what will be applied to the material; the predictive material state describes what previous inputs have left behind.

Across all three experiments the same pattern held: two or three numbers, updated as heat arrives, carried more predictive information than ten scalar summaries of the entire past — even though those descriptors had higher input dimension and access to the complete prior history. Thermal history is long, but the part of it that still governs what happens next is small, and it stays small when the kinetic law underneath is replaced, provided the readout is recalibrated. What does not carry over is the meaning of the coordinates: the transferred state predicted the target model's response well while showing almost no correspondence to its individual phase variables. It carries the memory, not the metallurgy. An encoder trained directly on the target model was more accurate at every tested dimension, so equal nominal dimension does not imply equal predictive content. Predictive-state dimension is not a fixed property of the alloy; it is set jointly by the response target, the admissible thermal-input domain, the error tolerance and the identification procedure.

The six-segment test exposed a separate requirement. A representation may predict well just beyond a segment boundary and still fail as a recursive state if it cannot be carried through the next segment. The learned update laws held their predictive information across all six segments and beat both the frozen-state control and the order-blind descriptor model, but horizon-six error did not fall as coordinates were added, and it stayed above the reference obtained by re-observing the state at every boundary. Within the simulated response, the sampled segment family and the stated tolerance, this is finite-precision recursive prediction rather than exact closure; a practical observer may still want intermittent measurements or stronger physical constraints on the update law.

A full-history sequence model answers a different question. Given 128 resampled prehistory values and 64 continuation values, a GRU reached a terminal-response RMSE of 0.07732 on the same held-out family, 71.8% below the source-training mean baseline. It forecasts well, but it rereads the entire history at every prediction rather than propagating a compact state, and it predicts a terminal response rather than the continuation-response signatures and recursive boundary forecasts used for the coordinate models, so the two sets of RMSE values are not numerically comparable. That distinction is the point of the exercise: forecasting accuracy and finite-dimensional closure are separate properties, and only the second says how much of the past a material has to keep.

In a control architecture, sensors provide partial observations of the hidden material condition, an observer estimates z_t, and a controller uses that estimate to select the next heat input, interpass time, or auxiliary treatment. The predictive material state then serves as a feedback variable, while the desired phase composition, microstructural morphology, or material property defines the control objective and the process window constrains the admissible inputs. Reset fits the same coordinate system but is not all-or-nothing: a strong

thermal excursion may erase phase-partition information while leaving parent-grain scale, chemistry or defect history intact, so it has to be tested coordinate by coordinate.

Two limits bound what these results can mean. Both kinetic systems describe Ti–6Al–4V and rest on related empirical and physical foundations, so their agreement and disagreement is cross-formulation computational evidence, not independent validation across specimens, laboratories or measurement modalities; the Murgau-class model is our own implementation of the published formulation, with limited numerical regularization. And real diffraction measurements will be messier than the observation model imposed here, which assumes independent, bounded Gaussian noise, whereas experimental phase estimates carry systematic bias, temporal correlation and temperature-dependent uncertainty. Branching a single prehistory into several diagnostic continuations would also demand highly reproducible specimens.

The framework could be tested experimentally using branched thermal paths. Replicated specimens would first receive controlled prehistories and would then be subjected to different diagnostic continuations. Histories judged to be state-equivalent would receive the same future thermal cycle. Co-located thermometry and time-resolved diffraction could test predictive and recursive closure of the phase state, while microscopy and diffraction on interrupted builds could probe coordinates associated with parent-grain scale and microstructural morphology. Holding out geometry, machine, and alloy batch would further test whether the learned state transfers beyond the response families used during identification to genuinely different material and processing conditions.

The practical consequence is a change of variable. Instead of asking only which process settings produced a microstructure, ask what condition the material is in now and what the next thermal input will do to it — then check that answer by predicting under shared future inputs, by updating across concatenated segments, and by keeping the smallest dimension that stays within a tolerance fixed in advance. No state vector is universally minimal. How much memory a material needs is a question about the prediction being made, and it can be measured.

Methods

Reference dynamical systems

The source system was the Ti–6Al–4V phase-fraction model developed by Nitzler et al. [15]. It describes the evolution of stable alpha, martensitic alpha, and beta under arbitrary temperature histories. The published equilibrium functions, characteristic temperatures, and calibrated kinetic parameters were implemented directly, and no parameter was fitted to the benchmarks used in this study. The source implementation was checked in three ways. First, cooling an initially pure-beta state at 410 K s^{-1} reproduced the high-cooling-rate regime reported for the original model [15, 31]. Second, the scalar and vectorized implementations agreed to machine precision on the validation histories. Third, all benchmark histories used a source-consistent continuous-cooling-transformation preconditioning step: the material was initialized as pure beta at 1400 K, cooled to 300 K at 500 K s^{-1}, held for 1 s, and then subjected to the prescribed thermal history. Production integrations used a time step of 0.0015625 s.

The transfer target was an independent implementation of the arbitrary-temperature Ti–6Al–4V model proposed by Murgau et al. [13]. Its state contains grain-boundary alpha, Widmanstätten alpha, martensitic alpha, and beta:

$$\boldsymbol{x}_{\mathrm{M}} = (X_{\alpha_{gb}}, X_{\alpha_W}, X_{\alpha_m}), \qquad X_\beta = 1 - X_{\alpha_{gb}} - X_{\alpha_W} - X_{\alpha_m}. \qquad (3)$$

For diffusional transformations, each integration step first inferred a fictive transformation time from the current transformed fraction and then applied the Johnson–Mehl–Avrami–Kolmogorov update. This procedure preserves the accumulated transformation history rather than restarting the kinetic clock at every time step. The equilibrium alpha fraction followed Eq. 23 of Murgau et al. [13]. Grain-boundary-alpha and Widmanstätten-alpha formation used the published fictive-time Johnson–Mehl–Avrami–Kolmogorov formulation. The Widmanstätten transformation rates and exponents were interpolated from Table 1 of Babu et al. [14]. Martensite formation used the Koistinen–Marburger relation with a start temperature of 575 °C. Stable-alpha

dissolution followed the published temperature-dependent rate law, with temperature expressed in kelvin, and martensite-recovery parameters were interpolated from Table 2 of Babu et al. [14].

Some published incremental expressions contain removable singularities when a phase fraction or its equilibrium capacity approaches zero. The implementation therefore used the corresponding analytical limits together with non-negative phase projection and phase-sum projection. The target implementation was evaluated using pure-beta cooling tests. It reproduced the expected transition from predominantly diffusional alpha at 0.167 and 20 K s^{-1} to predominantly martensitic alpha at 410 and 500 K s^{-1}. The total phase fraction was conserved to within 1.1×10^{-16}, and scalar and batch response propagation agreed exactly on the validation set. These checks establish internal consistency of the model transcription and numerical regime; they do not constitute experimental validation or demonstrate identity with the original authors' software.

History families and continuation banks

The hidden-state benchmark contained 1,800 three-cycle Nitzler histories. Each cycle heated from 300 K to 920–1360 K, held for 0.2–2.5 s, cooled at 35–520 K s^{-1} and included an inter-cycle dwell. Ascending-peak, mixed-order and descending-peak families contained 600 histories each. The first two families were used for fitting; the descending family was excluded until external evaluation.

Three disjoint banks of six peak–hold–cool continuations were used (main Table 1). Heating rate was 650 K s^{-1}, every continuation returned to 300 K, and seven equally spaced total-alpha observations were retained. Bank A supplied diagnostic observations, bank B supplied response signatures for coordinate identification, and bank C tested whether those coordinates remained informative under continuations absent from identification. The banks shared no complete peak–hold–cool parameter triple. Thermal parameters and generated response matrices are supplied as locked data.

Table 1. Diagnostic, identification and transfer continuation banks.

A: diagnostic				B: identification				C: transfer			
Index	Peak (K)	Hold (s)	Cooling (K s^{-1})	Index	Peak (K)	Hold (s)	Cooling (K s^{-1})	Index	Peak (K)	Hold (s)	Cooling (K s^{-1})
1	780	3	80	1	820	7	150	1	850	5	95
2	880	6	120	2	930	3	70	2	950	8	160
3	980	4	180	3	1030	9	220	3	1050	3	250
4	1080	8	90	4	1130	4	110	4	1150	10	130
5	1180	5	260	5	1220	7	340	5	1240	4	380
6	1260	3	400	6	1300	5	480	6	1340	6	520

Observation-driven coordinate identification

Stable and martensitic alpha were hidden from the estimator. Its only measured quantity was

$$y = X_{\alpha_s} + X_{\alpha_m}, \qquad (4)$$

which does not reveal the partition between stable and martensitic alpha.

Independent Gaussian noise with standard deviation 0.005 in absolute phase fraction was added to each observation and clipped to $[0,0.9]$. Fifty fixed noise realizations were analysed. Let A and B denote the standardized $N \times 42$ diagnostic and identification response matrices. A ridge-regularized map W was fitted from A to B, and its fitted response was decomposed as

$$AW = U\Sigma V^{\mathsf{T}}. \qquad (5)$$

The d-coordinate predictive representation was

$$Z_d = AWV_d, \qquad (6)$$

where V_d contains the first d right singular vectors. The procedure is reduced-rank predictive regression;[32] it does not reconstruct the hidden phases. Dimensions one to six were compared. Transfer to bank C used a cubic polynomial decoder whose ridge penalty was selected on an internal 80:20 training split and then refitted without descending histories. The plateau began at the smallest dimension whose median RMSE lay within 0.001 of the best tested dimension. Phase fractions appeared only inside the hidden simulator and, after model selection, in an affine alignment analysis. The 0.001 threshold was one fifth of the imposed observation-noise standard deviation and was fixed before comparing dimensions.

All reported calculations used a production integration step of 0.0015625 s, selected conservatively from a preliminary time-step refinement.

Thermal-history descriptor comparators and oracle controls

The ten order-invariant descriptors used throughout were maximum, mean and standard deviation of temperature; total duration; integrated temperature above 300 K; residence times above 700, 900 and 1100 K; maximum absolute temperature rate; and cycle count. Extremely randomized tree ensembles [33] mapped these descriptors to the relevant response target. Each ensemble used 160 trees; minimum leaf size was selected from 2, 5 and 10 on an internal subset of the ascending and mixed training families, then refitted on all training-family records. No descending-family record entered fitting or model selection.

For the oracle calibration, the complete Nitzler state was available only to construct targets and direct-propagation references. A descriptor ensemble reconstructed $(X_{\alpha_s}, X_{\alpha_m})$ from the ten thermal summaries, while a separate one-coordinate comparator reconstructed the same state from total alpha. Both were fitted on 1,200 ascending and mixed histories and evaluated on 600 descending histories. Each held-out state then received a newly sampled continuation. Exact order permutations and partial-reset controls used the same source implementation. Their purpose was calibration of the proposed tests, not discovery of state dimension.

Order and reset controls used the same CCT-preconditioned initial state and a production time step of 0.0015625 s. We generated 260 candidate three-segment histories with peaks from 920 to 1360 K, holds from 0.25 to 2.3 s and cooling rates from 40 to 500 K s^{-1}. Each candidate was run in ascending and descending peak order with the same segment multiset. The pair with the largest terminal phase-state separation was retained as a diagnostic illustration; it was not used for error, dimension or hypothesis-test calculations. Both terminal states then received the same sub-transus continuation, generated with seed 71 from a peak range of 760–1080 K, a hold range of 2–14 s and a cooling-rate range of 35–260 K s^{-1}. The reset control heated from 300 to 1400 K at 650 K s^{-1}, held for 25 s, cooled to 300 K at 500 K s^{-1} and then dwelled for 1 s.

Cross-model source transfer

The source-transfer benchmark used 900 newly generated histories: 600 ascending or mixed histories for fitting and 300 descending histories for testing. Nitzler and Murgau-class states were created by integrating identical prehistories. Both models were observed only through total-alpha responses to banks A–C.

For each noise realization, the source encoder was learned from Nitzler banks A and B and then frozen. Murgau bank-A responses were projected through this encoder. A cubic target decoder from the transferred coordinates to Murgau bank C was selected and fitted using only target total-alpha responses from the 600 training histories; target phase components were never used. A target-specific encoder–decoder was fitted as a comparison. Strict zero-shot prediction applied the Nitzler bank-B response reconstruction directly to Murgau bank B without target calibration and was reported separately. Dimensions one to three were tested over 50 fixed noise realizations.

For the matched descriptor comparison, the tree readout received the ten complete-prehistory summaries and predicted the same noisy Murgau bank-C responses. It used the shared 600-history training pool and the same internal validation histories as the transferred-state readout. The 300 descending histories and all target phase labels remained unavailable during selection and fitting.

Recursive update identification

The recursive benchmark contained 360 training-family and 180 descending-order test sequences, each with six segments. Segment controls were peak temperature, hold time and cooling rate; heating rate and terminal dwell were fixed by the continuation protocol. At every training boundary, Z_d was encoded from noisy bank-A responses. Candidate transitions paired (Z_n, u_n) with Z_{n+1}.

The primary update was an extremely randomized tree ensemble [33] with 160 trees. Minimum leaf size was selected from 2, 5 and 10 by complete six-step state rollout on training-family validation sequences, then refitted on all training-family transitions. A ridge-regularized quadratic update was selected on that training/validation partition as a functional-form sensitivity. A cubic response decoder mapped coordinates to clean bank-C total-alpha signatures.

For held-out sequences, only Z_0 was estimated from noisy diagnostics. Subsequent states were generated recursively from the predicted state and segment controls. No response was observed at intermediate boundaries. Direct re-encoding from a boundary's diagnostic responses was evaluated only as a lower reference. Errors were computed against clean bank-C response signatures at each horizon.

For the matched descriptor comparator, the ten descriptors were recomputed at each boundary from the complete accumulated temperature trace. The thermal input included the initial three-cycle prehistory and every subsequent segment; the summary discarded segment order. The readout was selected on complete training-family validation sequences and fitted to noisy bank-C responses. It used neither hidden phases nor intermediate material observations. Evaluation used the identical 180 descending-order sequences and clean response signatures.

Full-history GRU baseline. A gated recurrent unit received 128 uniformly resampled values from the complete initial prehistory and 64 values from the supplied continuation. It was trained on 1,200 ascending and mixed history families integrated at 0.00625 s, with an internal 960/240 training-validation split and evaluated on the held-out 600 descending histories. This baseline is included as a full-history forecast model; it does not define a finite-dimensional state or a recursively closed update.

Statistics and reproducibility

The statistical unit was an independently generated observation-noise realization applied to locked synthetic histories ($n = 50$ for each primary comparison and $n = 30$ for each sensitivity configuration). Reported point estimates are seed medians. Confidence intervals are deterministic 95% percentile-bootstrap intervals for the median, using 20,000 resamples for hidden-state identification and 10,000 for cross-model and recursive comparisons. They quantify sensitivity to imposed observation noise and do not represent material-specimen uncertainty.

Dimension reductions and descriptor comparisons used paired one-sided Wilcoxon signed-rank tests. The alternative direction was specified before evaluation, and the non-parametric test avoided assuming normally distributed RMSE differences across noise realizations. The familywise alpha level was 0.05. For each cross-model or recursive experiment, two adjacent-dimension tests and the descriptor comparison formed a three-test family adjusted by Holm's method. Hidden-state identification and transfer each formed a two-test adjacent-dimension family. When all 50 differences pointed in the prespecified direction, the exact one-sided probability reached 2^{-50}; Holm adjustment then gave the repeated values reported in Supplementary Table 6. Dimension was selected by the prespecified 0.001 RMSE tolerance rather than significance alone. No prospective sample-size calculation, random allocation or blinding was applicable to these deterministic simulations; history families and data splits were generated from fixed seeds and locked before comparison.

All history families, random seeds, temperature traces, accumulated thermal descriptors, response signatures and sequence archives were frozen before the reported comparisons and stored with SHA-256 digests. The executable software regenerates numerical summaries and vector, raster and print-resolution figure files from those locked inputs.

AI-assisted manuscript preparation

OpenAI Codex assisted with language editing. The authors checked the equations, model transcriptions, numerical results, references, wording and final artwork and remain accountable for the work. All final figures were generated deterministically from locked numerical inputs by the supplied Python and Matplotlib scripts; no generative-image system was used.

Data availability

No new experimental data were generated. The frozen random seeds, thermal trajectories, response signatures, descriptor arrays, sequence archives, benchmark tables and figure-source data required to reproduce the reported calculations are supplied with this submission as Supplementary Data and are openly available in Zenodo at https://doi.org/10.5281/zenodo.21827434.

Code availability

The custom Python code underlying the central results is supplied with this submission as Supplementary Software and is openly available in Zenodo at https://doi.org/10.5281/zenodo.21827434. It contains the Ti–6Al–4V demonstrator, observation-driven identification, recursive update and deterministic figure-generation scripts. Together with Supplementary Data, it regenerates the numerical summaries and final artwork. The implementation transcribes the published equations and parameters of Nitzler et al.; users who reuse it should cite the original kinetic-model publication.[15]

Funding

This study received no funding.

Author contributions

J.J. conceived the study, developed the methodology and software, performed the validation, investigation and formal analysis, curated the data, prepared the figures, and wrote the original draft. J.L. supervised the study and reviewed and edited the manuscript. A.Z. supervised the study and coordinated the project. V.K. supervised the study. All authors read and approved the final manuscript.

Competing interests

All authors declare no financial or non-financial competing interests.

References

1. Kenel, C. et al. In situ investigation of phase transformations in Ti–6Al–4V under additive manufacturing conditions combining laser melting and high-speed micro-X-ray diffraction. *Sci. Rep.* 7, 16358 (2017). https://doi.org/10.1038/s41598-017-16760-0.
2. DebRoy, T. et al. Additive manufacturing of metallic components—Process, structure and properties. *Prog. Mater. Sci.* 92, 112–224 (2018). https://doi.org/10.1016/j.pmatsci.2017.10.001.
3. Wei, H. L. et al. Mechanistic models for additive manufacturing of metallic components. *Prog. Mater. Sci.* 116, 100703 (2021). https://doi.org/10.1016/j.pmatsci.2020.100703.
4. Kouraytem, N., Li, X., Tan, W., Kappes, B. & Spear, A. D. Modeling process–structure–property relationships in metal additive manufacturing: a review on physics-driven versus data-driven approaches. *J. Phys. Mater.* 4, 032002 (2021). https://doi.org/10.1088/2515-7639/abca7b.
5. Williams, S. W. et al. Wire + Arc Additive Manufacturing. *Mater. Sci. Technol.* 32, 641–647 (2016). https://doi.org/10.1179/1743284715Y.0000000073.
6. Rodrigues, T. A., Duarte, V., Miranda, R. M., Santos, T. G. & Oliveira, J. P. Current status and perspectives on wire and arc additive manufacturing (WAAM). *Materials* 12, 1121 (2019). https://doi.org/10.3390/ma12071121.
7. Wu, B., Pan, Z., Ding, D., Cuiuri, D. & Li, H. Effects of heat accumulation on microstructure and mechanical properties of Ti6Al4V alloy deposited by wire arc additive manufacturing. *Addit. Manuf.* 23, 151–160 (2018). https://doi.org/10.1016/j.addma.2018.08.004.
8. Wu, B. et al. The effects of forced interpass cooling on the material properties of wire arc additively manufactured Ti6Al4V alloy. *J. Mater. Process. Technol.* 258, 97–105 (2018). https://doi.org/10.1016/j.jmatprotec.2018.03.024.
9. Coleman, B. D. & Gurtin, M. E. Thermodynamics with internal state variables. *J. Chem. Phys.* 47, 597–613 (1967). https://doi.org/10.1063/1.1711937.
10. Rice, J. R. Inelastic constitutive relations for solids: an internal-variable theory and its application to metal plasticity. *J. Mech. Phys. Solids* 19, 433–455 (1971). https://doi.org/10.1016/0022-5096(71)90010-X.
11. Horstemeyer, M. F. & Bammann, D. J. Historical review of internal state variable theory for inelasticity. *Int. J. Plast.* 26, 1310–1334 (2010). https://doi.org/10.1016/j.ijplas.2010.06.005.
12. Grong, Ø. & Shercliff, H. R. Microstructural modelling in metals processing. *Prog. Mater. Sci.* 47, 163–282 (2002). https://doi.org/10.1016/S0079-6425(00)00004-9.
13. Murgau, C. C., Pederson, R. & Lindgren, L.-E. A model for Ti–6Al–4V microstructure evolution for arbitrary temperature changes. *Model. Simul. Mater. Sci. Eng.* 20, 055006 (2012). https://doi.org/10.1088/0965-0393/20/5/055006.
14. Babu, B., Lundbäck, A. & Lindgren, L.-E. Simulation of Ti–6Al–4V additive manufacturing using coupled physically based flow stress and metallurgical model. *Materials* 12, 3844 (2019). https://doi.org/10.3390/ma12233844.
15. Nitzler, J., Meier, C., Müller, K. W., Wall, W. A. & Hodge, N. E. A novel physics-based and data-supported microstructure model for part-scale simulation of laser powder bed fusion of Ti–6Al–4V. *Adv. Model. Simul. Eng. Sci.* 8, 16 (2021). https://doi.org/10.1186/s40323-021-00201-9.
16. Shalizi, C. R. & Crutchfield, J. P. Computational mechanics: pattern and prediction, structure and simplicity. *J. Stat. Phys.* 104, 817–879 (2001). https://doi.org/10.1023/A:1010388907793.
17. Bialek, W., Nemenman, I. & Tishby, N. Predictability, complexity, and learning. *Neural Comput.* 13, 2409–2463 (2001). https://doi.org/10.1162/089976601753195969.

18. Tishby, N., Pereira, F. C. & Bialek, W. The information bottleneck method. In *Proc. 37th Allerton Conf. Communication, Control, and Computing* (eds Hajek, B. & Sreenivas, R. S.) 368–377 (University of Illinois, 1999).
19. Still, S., Sivak, D. A., Bell, A. J. & Crooks, G. E. Thermodynamics of prediction. *Phys. Rev. Lett.* 109, 120604 (2012). https://doi.org/10.1103/PhysRevLett.109.120604.
20. Bonatti, C. & Mohr, D. One for all: universal material model based on minimal state-space neural networks. *Sci. Adv.* 7, eabf3658 (2021). https://doi.org/10.1126/sciadv.abf3658.
21. Heidenreich, J. N. & Mohr, D. Extended minimal state cells (EMSC): self-consistent recurrent neural networks for rate- and temperature dependent plasticity. *Int. J. Plast.* 188, 104305 (2025). https://doi.org/10.1016/j.ijplas.2025.104305.
22. Ljung, L. System Identification: Theory for the User. 2nd edn (Prentice Hall PTR, 1999).
23. Brunton, S. L., Brunton, B. W., Proctor, J. L. & Kutz, J. N. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. PLoS ONE 11, e0150171 (2016). https://doi.org/10.1371/journal.pone.0150171.
24. Chen, R. T. Q., Rubanova, Y., Bettencourt, J. & Duvenaud, D. Neural ordinary differential equations. In Advances in Neural Information Processing Systems 31, 6571–6583 (2018).
25. Li, Z. et al. Fourier neural operator for parametric partial differential equations. In International Conference on Learning Representations (2021).
26. Mozaffar, M. et al. Data-driven prediction of the high-dimensional thermal history in directed energy deposition processes via recurrent neural networks. *Manuf. Lett.* 18, 35–39 (2018). https://doi.org/10.1016/j.mfglet.2018.10.002.
27. Roy, M. & Wodo, O. Data-driven modeling of thermal history in additive manufacturing. *Addit. Manuf.* 32, 101017 (2020). https://doi.org/10.1016/j.addma.2019.101017.
28. Xie, X. et al. Mechanistic data-driven prediction of as-built mechanical properties in metal additive manufacturing. *npj Comput. Mater.* 7, 86 (2021). https://doi.org/10.1038/s41524-021-00555-z.
29. Fang, L. et al. Data-driven analysis of process, structure, and properties of additively manufactured Inconel 718 thin walls. *npj Comput. Mater.* 8, 126 (2022). https://doi.org/10.1038/s41524-022-00808-5.
30. Fernandez-Zelaia, P., Dryepondt, S. N., Ziabari, A. K. & Kirka, M. M. Self-supervised learning of spatiotemporal thermal signatures in additive manufacturing using reduced order physics models and transformers. *Comput. Mater. Sci.* 232, 112603 (2024). https://doi.org/10.1016/j.commatsci.2023.112603.
31. Ahmed, T. & Rack, H. J. Phase transformations during cooling in α+β titanium alloys. *Mater. Sci. Eng. A* 243, 206–211 (1998). https://doi.org/10.1016/S0921-5093(97)00802-2.
32. Izenman, A. J. Reduced-rank regression for the multivariate linear model. *J. Multivar. Anal.* 5, 248–264 (1975). https://doi.org/10.1016/0047-259X(75)90042-1.
33. Geurts, P., Ernst, D. & Wehenkel, L. Extremely randomized trees. *Mach. Learn.* 63, 3–42 (2006). https://doi.org/10.1007/s10994-006-6226-1.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformation.pdf