AtomWorld-Mem: Memory-Restored World States for Long-Horizon Atomistic Evolution

Tian Luo[1,5*], **Ruge Zhang**[2,3,5*], **Haozhi Han**[4,5], **Yifeng Chen**[4], **Yunquan Zhang**[2,3], **Ting Cao**[5], **Yunxin Liu**[5], **Kun Li**[5†]

[1]Sichuan University, Chengdu, China
[2]Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China
[3]University of Chinese Academy of Sciences, Beijing, China
[4]School of Computer Science, Peking University, Beijing, China
[5]Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China
[*]Equal contribution.
[†]Corresponding author: likun@air.tsinghua.edu.cn

Abstract

High-fidelity atomistic evolution over long timescales requires more than observing the current crystal configuration. Instantaneous atomistic snapshots are often incomplete: locally similar configurations can correspond to different hidden dynamical contexts, future event preferences, and waiting-time scales. We argue that this *snapshot ambiguity* makes long-horizon atomistic evolution fundamentally a memory-based world-state restoration problem. To address this, we introduce **AtomWorld-Mem**, a memory-restored atomistic world model that recovers the latent world state missing from instantaneous crystal snapshots. AtomWorld-Mem treats the evolving alloy as an *AtomWorld*: spatial encoders write multi-scale atomistic keyframes from dense local topology and sparse long-range defect context, while short-term event memory and long-term structural memory integrate these keyframes across time to restore a future-predictive evolutionary state. The restored state is used to prioritize legal vacancy-mediated events under single-event Kinetic Monte Carlo (KMC) constraints, while event legality, physical execution, and residence-time updates remain governed by the underlying simulator. Empirically, AtomWorld-Mem improves long-horizon atomistic progress under fixed microscopic event budgets while maintaining high-fidelity evolution across energetic, structural, and vacancy-transport observables. It further transfers zero-shot across diverse unseen alloy–temperature AtomWorlds, suggesting that the learned memory-restoration mechanism captures reusable principles of hidden-state inference rather than a system-specific local energy heuristic. These results position memory-restored world-state modeling as a promising route toward efficient, physically grounded, and transferable atomistic evolution.

1 Introduction

Simulating long-horizon atomistic evolution at microscopic fidelity is crucial across scientific domains, from radiation damage to catalysis and phase transformations [1–3]. Existing long-timescale atomistic methods advance a system through admissible local transitions and update physical time according to transition rates [4, 5]. This rate-driven principle preserves physical consistency, but also exposes a fundamental blind spot: instantaneous rates do not reveal whether an event is a reversible local fluctuation or a transition that drives structural evolution [6–8]. Consequently, sim-

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

ulation effort can be repeatedly consumed within locally metastable regions, delaying access to rare but consequential transitions that shape macroscopic structure under practical budgets [7, 8].

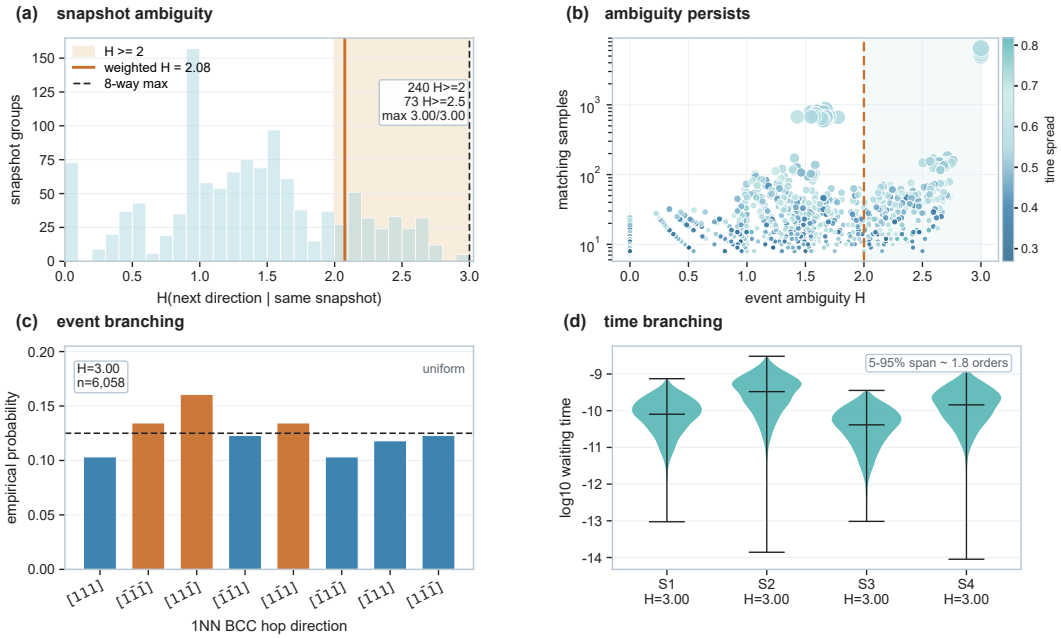


Figure 1: **Snapshot ambiguity in atomistic evolution.** Empirical analysis shows that identical local spatial observations can branch into different future trajectories and exhibit drastically different waiting times. This indicates that a single static snapshot is not a sufficient predictive state for long-horizon atomistic evolution.

A tempting direction is to make event selection aware of long-horizon consequences while remaining within the admissible event space. Yet this is ill-posed from the current spatial snapshot alone due to *snapshot ambiguity*: similar local configurations can branch into distinct future event distributions and markedly different waiting-time scales, as shown in Figure 1. In learning terms, an instantaneous atomic snapshot is not a sufficient statistic for the future [9, 10]; it aliases multiple hidden dynamical contexts into the same observation. What is missing is not another local descriptor, but the world state of the evolving atomic system: a latent dynamical representation that summarizes the hidden context needed for future event reasoning [11–13]. The central challenge is therefore to restore this world state from partial observations and temporal memory, enabling long-horizon event reasoning beyond instantaneous local rates.

To address this challenge, we model atomistic evolution as an evolving atomic world, where memory bridges instantaneous snapshots and long-horizon dynamics to recover the latent world state that snapshots alone cannot reveal. We introduce **AtomWorld-Mem**, a memory-restored atomistic world model that reconstructs this state across time, transforming snapshot-limited transition scoring into memory-grounded reasoning over atomistic evolution.

Concretely, AtomWorld-Mem implements this memory-restored world-state formulation through four coupled components: a *Keyframe Writer*, a *Latent-State Memory*, a *Legal Event Scorer*, and a *Physics-Guided Policy Learner*. The Keyframe Writer records compact, high-information keyframes from multi-scale atomistic observations, combining dense short-range local topology for immediate admissible transitions with sparse long-range defect context for slower structural evolution. The Latent-State Memory performs dual-scale temporal reasoning: **short-term memory** captures recent action history and local transition competition, while **long-term memory** accumulates structural biases and hidden defect-evolution trends across the trajectory. Together, these memories restore a latent world state from partial observations, forming a dynamically updated representation of the current crystal evolutionary context. The Legal Event Scorer evaluates admissible transitions under the restored state, while the Physics-Guided Policy Learner optimizes their prioritization under single-event KMC execution. By operating within the admissible event space and preserving physi-

cal time semantics, AtomWorld-Mem enables efficient, time-consistent evolution without sacrificing physical fidelity.

Empirically, the results support this state-restoration view. AtomWorld-Mem reaches up to 420× fixed-budget acceleration under matched legal-event budgets, showing that the restored world state enables long-horizon event reasoning beyond instantaneous local rates. The accelerated rollouts remain aligned with reference energetic, structural, and vacancy-transport trends in physical-time space. Memory ablations reveal complementary roles: short-term memory supports recent-event reasoning and vacancy arrival, whereas long-term memory supports structural-context accumulation and rollout fidelity. Zero-shot transfer across unseen alloy–temperature AtomWorlds further suggests that AtomWorld-Mem learns reusable hidden-state inference rather than a system-specific local energy heuristic.

2 Problem Formulation

We model long-horizon atomistic evolution as an *AtomWorld*: a partially observed atomistic world evolving under single-event KMC constraints. At step t, the complete microscopic configuration $x_t \in \mathcal{X}$ determines the legal event set, physical transition rates, and subsequent atomistic evolution. The learned model does not observe this complete state; it receives only structured partial observations and must infer the hidden evolutionary context needed for long-horizon decision making.

Under the single-event KMC constraint, each decision executes exactly one physically admissible vacancy-mediated exchange. For compactness, we write the legal event set as

$$\mathcal{A}_t = \mathcal{A}^{\mathrm{legal}}(x_t). \tag{1}$$

A legal event $a_t \in \mathcal{A}_t$ updates the AtomWorld by

$$x_{t+1} = f(x_t, a_t). \tag{2}$$

Each legal event has a physically defined base rate $r_t^{\mathrm{base}}(a)$. Conventional rejection-free KMC samples events proportionally to these rates,

$$p_{\mathrm{KMC}}(a \mid x_t) = \frac{r_t^{\mathrm{base}}(a)}{\sum_{a' \in \mathcal{A}_t} r_t^{\mathrm{base}}(a')}. \tag{3}$$

Eq. (2) defines the physically admissible state update, while Eq. (3) defines the passive rate-proportional event-selection rule. This preserves event legality and micro-kinetic consistency, but does not explicitly infer which admissible event is most consequential for long-horizon structural evolution.

2.1 Partial Observation and Legal Action Space

The complete configuration x_t determines legal events, base rates, and physical transitions, but the learned model only receives

$$o_t = \Omega(x_t) = \big(G_t^{\mathrm{local}},\, S_t^{\mathrm{def}}\big), \tag{4}$$

where G_t^{local} is a vacancy-centered local graph over structured 1NN–4NN shells, and S_t^{def} is a set of sparse defect tokens encoding longer-range defect context. AtomWorld-Mem is vacancy-centered in representation, but the control problem remains event-level: the policy does not generate unconstrained atomic moves, and all decisions are restricted to the legal set $\mathcal{A}_t$ defined in Eq. (1). Together, Eqs. (1) and (4) define the key partial-observation setting: the model acts only through legal KMC events, but must do so from incomplete observations.

2.2 Memory-Based State Restoration

A single local snapshot is generally not a sufficient predictive state for long-horizon atomistic evolution. Similar local observations may correspond to different hidden dynamical contexts, leading to different future event preferences, waiting-time scales, and structural trajectories. We therefore treat atomistic decision making as latent-state restoration under partial observability.

AtomWorld-Mem writes a compact spatial keyframe

$$k_t = \psi(o_t), \tag{5}$$

and infers a memory-restored latent state from keyframes and past actions,

$$z_t = \Phi(k_{\leq t}, a_{<t}). \tag{6}$$

Here, z_t represents hidden dynamical context that is missing from the current snapshot but useful for future atomistic evolution. It does not denote an unconstrained generative simulator. The policy acts on this restored state while remaining constrained to the legal KMC event set:

$$a_t \sim \pi_\theta(\cdot \mid z_t, \mathcal{A}_t), \qquad \operatorname{supp} \pi_\theta(\cdot \mid z_t, \mathcal{A}_t) \subseteq \mathcal{A}_t. \tag{7}$$

Thus, Eqs. (5)–(7) formalize the central state-restoration view of AtomWorld-Mem: spatial observations are written into keyframes, memory restores the latent evolutionary state, and the policy prioritizes only legal events.

2.3 Training Objective under KMC Constraints

The goal is to improve long-horizon atomistic progress under a fixed legal-event budget T. We define the immediate energy drop and shaped training reward as

$$\Delta E_t = E(x_t) - E(x_{t+1}), \qquad r_t^{\text{train}} = \lambda_E \Delta E_t + r_t^{\text{shape}}, \tag{8}$$

and optimize

$$J(\pi_\theta) = \mathbb{E}_{\pi_\theta}\left[\sum_{t=0}^{T-1} \gamma^t r_t^{\text{train}}\right]. \tag{9}$$

Learning changes the decision rule for prioritizing legal events under partial observability; event legality, physical execution, and time advancement remain governed by the underlying KMC process.

3 AtomWorld-Mem

AtomWorld-Mem implements the formulation in Section 2 with four modules: a Keyframe Writer, Latent-State Memory, Legal Event Scorer, and Physics-Guided Policy Learner.

3.1 Architecture and Formulation

AtomWorld-Mem follows a memory-centered principle: spatial encoders do not directly define the predictive state. Instead, they write high-information atomistic keyframes, and memory restores the missing evolutionary state by integrating these keyframes with past legal events. Space provides the instantaneous atomistic context, while memory converts partial observations into a future-predictive latent world state.

Keyframe Writer. The Keyframe Writer converts the partial observation o_t in Eq. (4) into a compact spatial keyframe. It uses two complementary spatial pathways:

$$h_t^{\text{dense}} = f_{\text{gnn}}(G_t^{\text{local}}), \qquad h_t^{\text{sparse}} = f_{\text{attn}}(S_t^{\text{def}}), \qquad k_t = \phi_{\text{fuse}}([h_t^{\text{dense}} \| h_t^{\text{sparse}}]). \tag{10}$$

The dense pathway captures short-range topology that determines immediate legal transitions, while the sparse pathway captures long-range defect context that shapes slower structural evolution. The keyframe k_t is not treated as a complete state; it is a spatial write into memory that must be interpreted together with past evolution.

Latent-State Memory. The Latent-State Memory restores hidden evolutionary context by integrating spatial keyframes with previous legal events, following the state-restoration view in Eq. (6). We decompose this restoration into two temporal scales. The short-term event memory summarizes recent action history and local transition competition:

$$c_t^{\text{short}} = \phi_{\text{hist}}(\phi_{\text{act}}(a_{t-H:t-1})). \tag{11}$$

It captures recent vacancy motion, local reversibility, repeated back-and-forth exchanges, and short-horizon competition among nearby legal events. The long-term structural memory accumulates slower evolutionary biases:

$$m_t^{\text{long}} = f_{\text{rnn}}\left(m_{t-1}^{\text{long}}, [k_t \| c_t^{\text{short}}]\right). \tag{12}$$

This recurrent state retains hidden defect-evolution trends, aggregation tendency, vacancy–solute interaction patterns, and structural biases over many KMC steps. The restored world state is obtained by combining instantaneous space, short-term event context, and long-term structural memory:

$$z_t = \phi_{\mathrm{mem}}\Big([k_t \| c_t^{\mathrm{short}} \| m_t^{\mathrm{long}}]\Big). \tag{13}$$

Thus, z_t is not merely an encoding of the current observation; it is a memory-restored latent state used for both future event preference and time-scale reasoning.

Legal Event Scorer. Given the restored state z_t, AtomWorld-Mem scores each legal event $a \in \mathcal{A}^{\mathrm{legal}}(x_t)$. For each event, we construct

$$\xi_t(a) = \big(\tilde{G}_t(a), d(a), \mu_t(a)\big), \tag{14}$$

where $\tilde{G}_t(a)$ is event-local structure, $d(a)$ is the event direction, and $\mu_t(a)$ denotes auxiliary energetic or kinetic features. The event descriptor is encoded into $e_t(a)$ and scored against the restored state:

$$s_t(a) = g(z_t, e_t(a)). \tag{15}$$

This memory-conditioned scoring is essential: two events with similar local descriptors may have different long-horizon consequences depending on the hidden evolutionary context stored in z_t. To preserve physical grounding, the learned score is fused with KMC and energetic priors:

$$\ell_t(a) = \alpha s_t(a) + \beta \log r_t^{\mathrm{base}}(a) + \eta q_t(a), \tag{16}$$

where $r_t^{\mathrm{base}}(a)$ is the base KMC rate and $q_t(a)$ is an energy-aware prior. These scores only prioritize among already legal KMC events; they do not generate unconstrained atomic moves.

Physics-Guided Policy Learner. The policy induced by Eq. (16) is trained with PPO using the shaped reward in Eq. (8). The clipped actor objective is

$$L_{\mathrm{ppo}} = -\mathbb{E}_t\left[\min\Big(\rho_t(\theta)\hat{A}_t,\ \mathrm{clip}(\rho_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\Big)\right], \tag{17}$$

where

$$\rho_t(\theta) = \frac{\pi_\theta(a_t \mid z_t, \mathcal{A}^{\mathrm{legal}}(x_t))}{\pi_{\theta_{\mathrm{old}}}(a_t \mid z_t, \mathcal{A}^{\mathrm{legal}}(x_t))}. \tag{18}$$

We further use a soft physics teacher $p_t^{\mathrm{phys}}(a)$ derived from explicit kinetic and energetic cues:

$$L_{\mathrm{phys}} = -\mathbb{E}_t \sum_{a \in \mathcal{A}^{\mathrm{legal}}(x_t)} p_t^{\mathrm{phys}}(a) \log \pi_\theta(a \mid z_t, \mathcal{A}^{\mathrm{legal}}(x_t)). \tag{19}$$

The full actor objective is

$$L_{\mathrm{actor}} = L_{\mathrm{ppo}} + \lambda_{\mathrm{phys}} L_{\mathrm{phys}} + \lambda_{\mathrm{time}} L_{\mathrm{time}} - \lambda_{\mathrm{ent}} H[\pi_\theta(\cdot \mid z_t)]. \tag{20}$$

Here L_{time} is the restored-state time-consistency objective in Eq. (25). The critic is also conditioned on z_t, so value estimation is based on the memory-restored evolutionary context rather than on the instantaneous snapshot alone.

3.2 Restored-State Time Consistency

AtomWorld-Mem does not replace the KMC physical clock. During execution, time is advanced by the standard residence-time rule from the base KMC rates:

$$\Gamma_t(x_t) = \sum_{a \in \mathcal{A}^{\mathrm{legal}}(x_t)} r_t^{\mathrm{base}}(a), \qquad \Delta t_t^{\mathrm{KMC}} = -\frac{\log u_t}{\Gamma_t(x_t)}. \tag{21}$$

Thus, the time unit, temperature dependence, attempt frequency, and barrier-to-rate conversion remain those of the underlying simulator. The learned policy changes which legal event is selected, not how the residence-time increment is computed at a visited state. We do not claim that AtomWorld-Mem is an exact kinetics-preserving replacement for passive KMC: the biased trajectory distribution is not guaranteed to satisfy detailed balance or the stationary path measure of unbiased KMC. Our

goal is rare-event discovery and low-energy structural evolution under legal-event constraints, while monitoring whether the induced trajectories remain physically meaningful.

Time prediction is induced indirectly through restored-state event hazards rather than by an independent clock head. We interpret the fused score in Eq. (16) as an event-wise log-hazard:

$$\lambda_\theta(a \mid z_t) = \exp(\ell_t(a)), \qquad \Lambda_\theta(z_t) = \sum_{a \in \mathcal{A}^{\mathrm{legal}}(x_t)} \lambda_\theta(a \mid z_t). \tag{22}$$

When clear, we suppress the dependence of Λ_θ on the current legal event set. The same hazard field defines both the legal-event policy and the expected waiting-time readout:

$$\pi_\theta(a \mid z_t) = \frac{\lambda_\theta(a \mid z_t)}{\Lambda_\theta(z_t)}, \qquad \widehat{\Delta t}_\theta(z_t) = \Lambda_\theta(z_t)^{-1}. \tag{23}$$

Thus, event preference and time-scale prediction are tied to the same memory-restored state.

We constrain this temporal readout using the base-rate prior, event-distribution alignment, and the executed KMC waiting time. Let

$$p_t^{\mathrm{base}}(a) = \frac{r_t^{\mathrm{base}}(a)}{\Gamma_t(x_t)}. \tag{24}$$

The auxiliary event-time consistency loss is

$$L_{\mathrm{time}} = D_{\mathrm{KL}}\big(p_t^{\mathrm{base}}(\cdot) \,\|\, \pi_\theta(\cdot \mid z_t)\big) + \log \frac{\Gamma_t(x_t)}{\Lambda_\theta(z_t)} + \big(\Lambda_\theta(z_t) - \Gamma_t(x_t)\big)\Delta t_t^{\mathrm{KMC}}. \tag{25}$$

The KL term softly regularizes event probabilities toward the physical rate structure, while the second term compares waiting-time statistics induced by the learned total hazard against the KMC residence-time target. This loss is not a detailed-balance constraint and does not make the biased policy path identical to passive KMC; it calibrates and diagnoses whether the restored state captures the local hazard landscape of the states it visits.

Throughout the paper, high fidelity does not mean exact reproduction of the unbiased KMC path measure. It refers to fidelity under three operational constraints: legal-event execution, KMC-grounded residence-time updates at visited states, and agreement with reference atomistic dynamics across energetic, structural, and vacancy-transport observables. Temporal fidelity is therefore a stringent probe of state-restoration quality: a model may improve energy descent through local heuristics, but matching the waiting-time scale requires recovering the broader legal-event hazard landscape.

4 Experiments

4.1 Experimental Setup

We train and evaluate AtomWorld-Mem on diffusion-driven Fe–Cu alloy evolution, a canonical vacancy-mediated setting for precipitation and defect-driven structural change. Each *AtomWorld* has microscopic states that determine legal events, base KMC rates, and transitions, while the model uses structured snapshots with internal memory. The actions are legal vacancy-mediated exchanges, and comparisons are step-controlled unless otherwise stated: methods start from matched initial-state families and use equal legal-event budgets, isolating fixed-budget event-prioritization quality. Learning experiments use NVIDIA A100 hardware.

4.2 Evaluating Memory-Restored AtomWorld Modeling

Our experiments test whether memory restores latent AtomWorld state by evaluating fixed-budget progress, trajectory-level fidelity, memory ablations, and zero-shot transfer.

4.2.1 Fixed-Budget Atomistic Progress

We first test whether restoring latent evolutionary context improves use of a fixed event budget. All methods start from matched initial-state families and execute the same number of legal vacancy-mediated events, so differences reflect event prioritization rather than rollout length.

Table 1: **Main results under matched legal-event budgets.** AF is the fixed-budget acceleration factor over Zacros (traditional KMC) [14], not a claim of path-measure-preserving kinetic speedup. Agg. Degree measures Cu aggregation, Arrival Hit measures vacancy arrival to aggregation-relevant Cu clusters, and DynErr measures deviation from reference KMC dynamics.

Method	Memory	Dilute setting: 0.18% Cu				Relatively concentrated setting: 1.34% Cu			
		AF ↑	Agg. Degree ↑	Arrival Hit ↑	DynErr ↓	AF ↑	Agg. Degree ↑	Arrival Hit ↑	DynErr ↓
Zacros (Traditional KMC)	✗	1.00×	0.016	5.0%	0.000	1.00×	0.072	43.8%	0.000
Energy-Greedy KMC	✗	11.03×	0.035	45.0%	0.466	2.16×	0.074	52.5%	0.371
Snapshot-only AtomWorld	✗	66.19×	0.116	67.8%	0.318	11.13×	0.166	67.5%	0.299
AtomWorld-Mem	✓	**420.05×**	**0.432**	**85.5%**	**0.221**	**51.25×**	**0.367**	**80.9%**	**0.262**

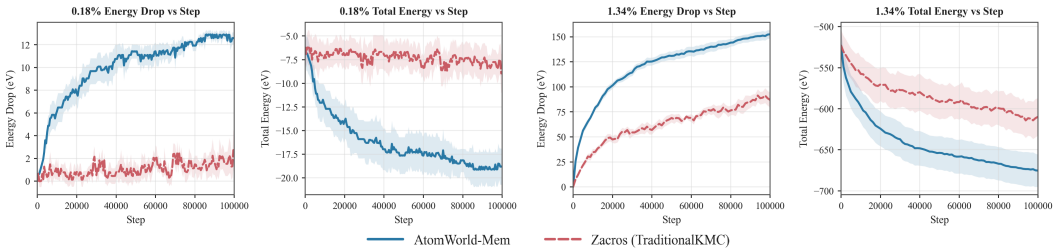


Figure 2: **Fixed-budget energy evolution over 100k legal KMC steps.** From left to right: energy drop and total energy at 0.18% Cu, followed by energy drop and total energy at 1.34% Cu. AtomWorld-Mem achieves larger energy drops and lower total energies than Zacros under the same legal-event budget. Shaded regions indicate trajectory variability across runs.

Table 1 separates passive sampling, local energy descent, snapshot-based spatial reasoning, and memory-restored state inference. Snapshot-only AtomWorld already improves over Zacros, showing that spatial world representations contain useful event context. However, adding memory yields the decisive gain: AF rises from 66.19× to 420.05× at 0.18% Cu and from 11.13× to 51.25× at 1.34% Cu. AtomWorld-Mem also improves Arrival Hit and DynErr, indicating that memory improves not only scalar progress but also long-horizon event prioritization quality.

Figure 2 shows that AtomWorld-Mem separates early from Zacros and maintains its advantage throughout the 100k-step rollout, ruling out an endpoint-only gain. The restored state repeatedly reallocates legal events toward structurally productive transitions. Figure 3 further shows that the advantage persists across Cu concentrations, averaging 186×. These results support the first prediction of our hypothesis: hidden-state restoration improves fixed-budget atomistic progress beyond snapshot-level event scoring or local energy heuristics.

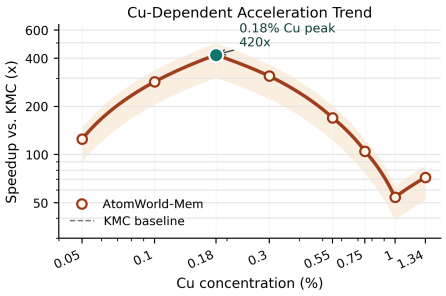


Figure 3: **Cu-dependent acceleration trend.** Gains remain positive across the tested Cu range, peak at 420× near 0.18% Cu, and average 186×.

4.2.2 Restored-State Fidelity and Correctness

Fast progress alone does not imply a high-fidelity atomistic world model: a policy may accelerate energy descent by exploiting local shortcuts. We therefore use trajectory-level fidelity as a stricter test of state restoration. Since AtomWorld-Mem intentionally biases legal-event selection, we do not claim exact reproduction of the unbiased KMC path measure or detailed balance. Instead, high fidelity here means that the induced rollout remains aligned with reference atomistic dynamics across energetic, structural, and vacancy-transport observables.

Figure 4 shows that AtomWorld-Mem remains broadly aligned with Zacros in physical-time space. Normalized energy follows the same relaxation trend; cluster count and largest-cluster size track

Table 2: **Memory ablation at 0.18% Cu.** Short/Long denote short-term event memory and long-term structural memory. TimeErr and DynErr measure temporal and rollout deviation from Zacros.

Variant	Short	Long	AF↑	Agg.↑	TimeErr↓	DynErr↓	Arr.Hit↑
Zacros (Traditional KMC)	✗	✗	1.00×	0.016	0.000	0.000	5.00%
Snapshot-only AtomWorld	✗	✗	62.16×	0.089	0.777	0.591	65.30%
No short-term memory	✗	✓	127.43×	0.183	0.475	0.293	71.64%
No long-term memory	✓	✗	148.25×	0.235	0.599	0.348	73.91%
Full AtomWorld-Mem	✓	✓	**418.44**×	**0.426**	**0.308**	**0.192**	**83.22%**

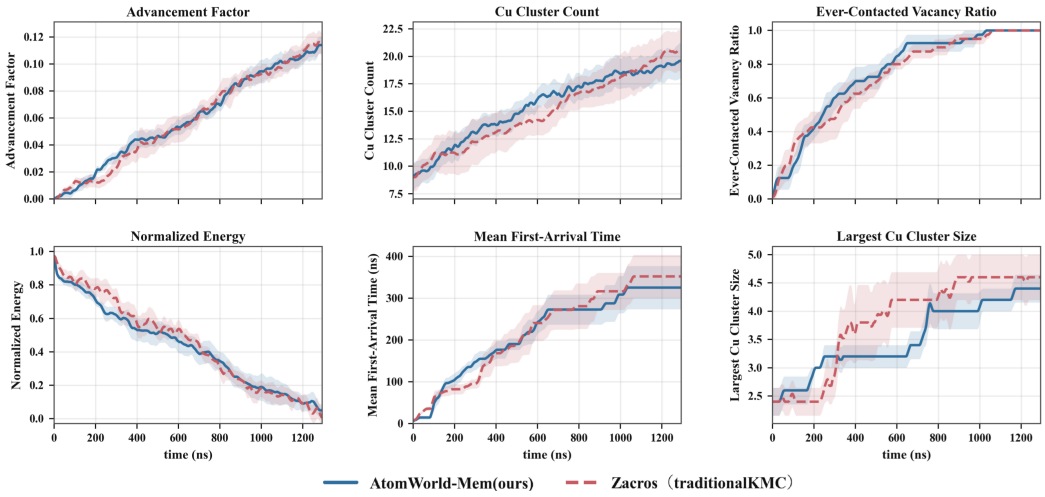


Figure 4: **High-fidelity atomistic evolution in physical-time space.** We compare AtomWorld-Mem with Zacros traditional KMC across six trajectory-level observables: advancement factor, Cu cluster count, ever-contacted vacancy ratio, normalized energy, mean first-arrival time, and largest Cu cluster size. AtomWorld-Mem follows the reference trends while maintaining faster fixed-budget progress, indicating that memory-restored event prioritization does not collapse into an unphysical local shortcut.

the aggregation regime; and vacancy-related metrics preserve the ordering and timing of vacancy–cluster interactions. These observables test fidelity beyond scalar energy descent. The result supports the restored-state interpretation: AtomWorld-Mem improves long-horizon event prioritization because memory recovers evolutionary context missing from instantaneous snapshots, rather than because the policy exploits a system-specific local energy shortcut.

4.2.3 Memory Ablation: Does Memory Restore the Missing State?

We next test the mechanism directly: if memory restores hidden evolutionary state, removing either memory scale should degrade both progress and fidelity. Table 2 shows that snapshot-only spatial reasoning is useful but incomplete, with improved AF but large TimeErr and DynErr. Adding either memory scale improves progress and fidelity, while the full model performs best across learned variants, increasing AF to 418.44× and reducing DynErr to 0.192. This indicates that AtomWorld-Mem benefits from temporal state restoration rather than only a stronger instantaneous event scorer.

4.2.4 Zero-Shot Transfer Across AtomWorlds

We finally evaluate zero-shot transfer to unseen alloy–temperature AtomWorlds without finetuning. We evaluate 16 unseen BCC alloy–temperature domains spanning Fe–Cr, Fe–Al, Fe–Si, Fe–V, Fe–Mo, Mo–Re, W–Re, and Ta–W. As shown in Figure 5, AtomWorld-Mem improves fixed-budget energy-drop trajectories in every domain, with a median AUC gain of **22.9**× and a maximum gain of **52.1**×. Dynamics-fidelity error remains small or moderate for most domains, with larger devia-

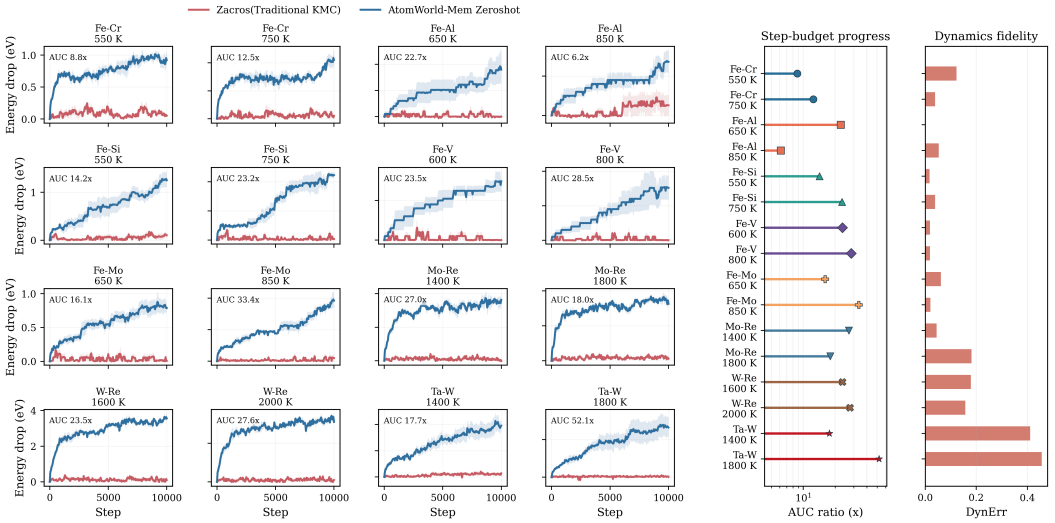


Figure 5: **Zero-shot transfer across 16 unseen BCC alloy–temperature AtomWorlds.** Left: energy-drop trajectories under a shared 10k-step budget. Right: per-domain AUC acceleration ratio and dynamics-fidelity error.

tions mainly in refractory systems. These consistent gains suggest transferable memory-based state inference rather than an Fe–Cu-specific local energy heuristic.

5 Related Work

Kinetic Monte Carlo and policy-guided atomistic kinetics. Kinetic Monte Carlo (KMC) simulates long-timescale atomistic evolution through physically admissible events and residence-time updates [15, 16, 5, 17, 2, 18]. Adaptive, accelerated, self-evolving, ML-assisted, and policy-guided variants improve event discovery, local updates, rate estimation, or transition selection under physical constraints [19–26]. However, these methods typically condition decisions on the current local configuration or emphasize aggregate sampling efficiency, leaving *state incompleteness* largely unaddressed: similar snapshots may hide different dynamical contexts and long-horizon event preferences [22, 20, 27]. AtomWorld-Mem retains KMC-grounded execution and timing, but uses memory to restore a future-predictive state for legal-event prioritization [15, 17].

World models and latent dynamics under partial observability. World models and latent-dynamics methods learn compact internal states from partial observations for prediction, planning, and control [28–32]. They often learn or imagine environment transitions in latent space, whereas atomistic KMC already provides a physically constrained transition system [33, 34, 17, 18]. AtomWorld-Mem adopts the latent-state view for a different purpose: restoring hidden dynamical context to rank admissible KMC events, while leaving event legality, execution, and residence-time updates to the underlying simulator [35–37, 5].

6 Conclusion and Limitations

We presented AtomWorld-Mem, a memory-restored atomistic world model for long-horizon evolution under physical constraints. The key idea is to treat atomistic evolution as hidden-state restoration rather than snapshot-level event sampling. Across progress, fidelity, ablation, and transfer tests, memory-restored states outperform snapshot-only reasoning while remaining physically grounded.

This work has limitations. Our experiments focus on vacancy-mediated on-lattice BCC alloy evolution, where the legal event space is explicitly defined by the underlying KMC simulator. Although the underlying memory-restored world-state formulation is potentially applicable to broader atomistic settings, including off-lattice or on-the-fly KMC scenarios, such extensions require integrating memory-based state restoration with dynamic event discovery, transition-state searches, or adaptive

event catalog construction, which are not demonstrated in the present work. Future research will investigate these directions, together with richer materials, more complex event catalogs, stronger clock-consistency objectives, and more expressive memory mechanisms.

7 Acknowledgements

This work is supported by Tecorigin and Wuxi Research Institute of Applied Technologies, Tsinghua University under Grant 20242001120.

References

[1] Cindy English and Jonathan Hyde. Radiation damage of reactor pressure vessel steels. *Comprehensive Nuclear Materials*, pages 151–180, 2012.

[2] Miguel Pineda and Michail Stamatakis. Kinetic monte carlo simulations for heterogeneous catalysis: Fundamentals, current status, and challenges. *The Journal of Chemical Physics*, 156(12), 2022.

[3] Georges Martin and Frédéric Soisson. Kinetic monte carlo method to model diffusion controlled phase transformations in the solid state. In *Handbook of Materials Modeling: Methods*, pages 2223–2248. Springer, 2005.

[4] Blas Pedro Uberuaga and Danny Perez. Computational methods for long-timescale atomistic simulations. In *Handbook of Materials Modeling: Methods: Theory and Modeling*, pages 683–688. Springer, 2020.

[5] Arthur F Voter. Introduction to the kinetic monte carlo method. In *Radiation effects in solids*, pages 1–23. Springer, 2007.

[6] Peter Kratzer. Monte carlo and kinetic monte carlo methods. *arXiv preprint arXiv:0904.2556*, 2009.

[7] Michail Stamatakis and Dionisios G Vlachos. Unraveling the complexity of catalytic reactions via kinetic monte carlo simulation: current status and frontiers. *Acs Catalysis*, 2(12):2648–2663, 2012.

[8] Karsten Reuter. First-principles kinetic monte carlo simulations for heterogeneous catalysis: Concepts, status, and frontiers. *ChemInform*, 43(36):no, 2012.

[9] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.

[10] Onno Eberhard, Michael Muehlebach, and Claire Vernade. Partially observable reinforcement learning with memory traces. *arXiv preprint arXiv:2503.15200*, 2025.

[11] Han Wang, Linfeng Zhang, Jiequn Han, et al. Deepmd-kit: A deep learning package for many-body potential energy representation and molecular dynamics. *Computer Physics Communications*, 228:178–184, 2018.

[12] Dedi Wang, Yihang Wang, Luke Evans, and Pratyush Tiwary. From latent dynamics to meaningful representations. *Journal of Chemical Theory and Computation*, 20(9):3503–3513, 2024.

[13] Imran Nasim and Michael E Henderson. Dynamically meaningful latent representations of dynamical systems. *Mathematics*, 12(3):476, 2024.

[14] Michail Stamatakis and Dionisios G Vlachos. A graph-theoretical kinetic monte carlo framework for on-lattice chemical kinetics. *The Journal of chemical physics*, 134(21), 2011.

[15] Daniel T Gillespie. Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry*, 81(25):2340–2361, 1977.

[16] Daniel T Gillespie. Stochastic simulation of chemical kinetics. *Annu. Rev. Phys. Chem.*, 58(1):35–55, 2007.

[17] Kristen A Fichthorn and W Hh Weinberg. Theoretical foundations of dynamical monte carlo simulations. *The Journal of chemical physics*, 95(2):1090–1096, 1991.

[18] Mie Andersen, Chiara Panosetti, and Karsten Reuter. A practical guide to surface kinetic monte carlo simulations. *Frontiers in chemistry*, 7:202, 2019.

[19] Graeme Henkelman and Hannes Jónsson. Long time scale kinetic monte carlo simulations without lattice approximation and predefined event table. *The Journal of Chemical Physics*, 115(21):9657–9666, 2001.

[20] Lijun Xu and Graeme Henkelman. Adaptive kinetic monte carlo for first-principles accelerated dynamics. *The Journal of chemical physics*, 129(11), 2008.

[21] Andreas Pedersen and Hannes Jónsson. Distributed implementation of the adaptive kinetic monte carlo method. *Mathematics and Computers in Simulation*, 80(7):1487–1498, 2010.

[22] Oleg Trushin, Altaf Karim, Abdelkader Kara, and Talat S Rahman. Self-learning kinetic monte carlo method: application to cu (111). *Physical Review BCondensed Matter and Materials Physics*, 72(11): 115401, 2005.

[23] Normand Mousseau, Laurent J Lewis, et al. The kinetic activation-relaxation technique: A powerful off-lattice on-the-fly kinetic monte carlo algorithm. *arXiv preprint arXiv:0805.2158*, 2008.

[24] Luca Messina, Nicolas Castin, Christophe Domain, and Pär Olsson. Introducing ab initio based neural networks for transition-rate prediction in kinetic monte carlo simulations. *Physical Review B*, 95(6): 064112, 2017.

[25] Jyri Kimari, Ville Jansson, Simon Vigonski, Ekaterina Baibuz, Roberto Domingos, Vahur Zadin, and Flyura Djurabekova. Application of artificial neural networks for rigid lattice kinetic monte carlo studies of cu surface diffusion. *Computational Materials Science*, 183:109789, 2020.

[26] Hao Tang, Boning Li, Yixuan Song, Mengren Liu, Haowei Xu, Guoqing Wang, Heejung Chung, and Ju Li. Reinforcement learning-guided long-timescale simulation of hydrogen transport in metals. *Advanced Science*, 11(5):2304122, 2024.

[27] Abhijit Chatterjee and Arthur F Voter. Accurate acceleration of kinetic monte carlo simulations through the modification of rate constants. *The Journal of chemical physics*, 132(19), 2010.

[28] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International conference on machine learning*, pages 2555–2565. PMLR, 2019.

[29] Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick Van der Smagt. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.

[30] Lars Buesing, Theophane Weber, Sébastien Racaniere, SM Eslami, Danilo Rezende, David P Reichert, Fabio Viola, Frederic Besse, Karol Gregor, Demis Hassabis, et al. Learning and querying fast generative models for reinforcement learning. *arXiv preprint arXiv:1802.03006*, 2018.

[31] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. *Advances in neural information processing systems*, 31, 2018.

[32] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.

[33] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.

[34] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.

[35] Michael Littman and Richard S Sutton. Predictive representations of state. *Advances in neural information processing systems*, 14, 2001.

[36] Ahmed Hefny, Zita Marinho, Wen Sun, Siddhartha Srinivasa, and Geoffrey Gordon. Recurrent predictive state policy networks. In *International Conference on Machine Learning*, pages 1949–1958. PMLR, 2018.

[37] Lingheng Meng, Rob Gorbet, and Dana Kulić. Memory-based deep reinforcement learning for pomdps. In *2021 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 5619–5626. IEEE, 2021.

A Experimental Details

A.1 Simulation Environments

The main experiments use vacancy-mediated Fe–Cu alloy evolution, a canonical setting for precipitation and defect-driven structural change. Each environment is treated as an AtomWorld: the full microscopic configuration determines legal events, base KMC rates, and physical transitions, while the learned model receives only structured partial observations and memory. The action space consists of vacancy-mediated atomic exchanges with local-environment-dependent Arrhenius rates.

Unless otherwise stated, comparisons are step-controlled. All methods start from matched initial-state families and execute the same number of legal KMC events. This isolates event-prioritization quality: differences in progress reflect how effectively a method uses the same microscopic event budget, rather than differences in rollout length.

A.2 Baselines and Ablations

Zacros traditional KMC. Zacros traditional KMC is the passive reference simulator. It samples legal events proportionally to their base transition rates and advances physical time using the standard residence-time rule. It provides the reference for both fixed-budget progress and trajectory-level fidelity.

Energy-Greedy KMC. Energy-Greedy KMC selects legal events using local energetic improvement. This baseline tests whether the gains of AtomWorld-Mem can be explained by a simple local energy-descent heuristic.

Snapshot-only AtomWorld. Snapshot-only AtomWorld keeps the instantaneous spatial representation but removes temporal memory. This baseline tests whether multi-scale spatial observations alone are sufficient, or whether hidden evolutionary context must be restored from history.

No short-term memory. This variant removes recent-event memory while retaining long-term structural memory. It tests the role of recent action history, local transition competition, vacancy backtracking, and short-horizon event reasoning.

No long-term memory. This variant removes recurrent structural memory while retaining short-term event history. It tests the role of slow structural-context accumulation, hidden defect-evolution trends, and long-horizon aggregation bias.

A.3 Evaluation Metrics

Aggregation degree (Agg. Degree). Let $\mathcal{C}(x)$ be the connected components of the Cu neighbor graph under periodic boundary conditions, N_{Cu} the total number of Cu atoms, and $n_{\min} \geq 2$ the minimum qualifying cluster size. The connectivity rule and $n_{\min}$ are fixed across methods. We report the mean endpoint value across evaluation runs:

$$\mathrm{Agg}(x) = \frac{1}{N_{\mathrm{Cu}}} \sum_{C \in \mathcal{C}(x)} |C| \, \mathbf{1}\{|C| \geq n_{\min}\}. \tag{26}$$

Vacancy arrival hit rate (Arrival Hit). Let R be the number of evaluation runs, $N_v^{(r)}$ the number of vacancies in run r, and B the legal-event budget. Let $I_{r,v}(n)$ indicate whether vacancy v contacts a qualifying Cu cluster after event n, using a fixed contact neighborhood. Each vacancy is counted once; the expression below includes initial contacts:

$$\mathrm{ArrivalHit}(B) = 100 \left[\frac{1}{R} \sum_{r=1}^{R} \frac{1}{N_v^{(r)}} \sum_{v=1}^{N_v^{(r)}} \mathbf{1}\Big\{ \max_{0 \leq n \leq B} I_{r,v}(n) = 1 \Big\} \right] \%. \tag{27}$$

Dynamics error (DynErr). Let $\mu_q^M(\tau_k)$ and $\mu_q^{\mathrm{ref}}(\tau_k)$ denote ensemble-mean values of observable q for method M and KMC, respectively. Here Q is the number of specified observables and $\{\tau_k\}_{k=1}^K$ is a common physical-time grid within the interval supported by all methods. The reference scale s_q uses a fixed positive floor ϵ_q with the same units as observable q:

$$s_q = \max\left\{ \max_k \mu_q^{\mathrm{ref}}(\tau_k) - \min_k \mu_q^{\mathrm{ref}}(\tau_k), \epsilon_q \right\}, \tag{28}$$

$$\mathrm{DynErr}(M) = \frac{1}{Q} \sum_{q=1}^{Q} \frac{\sqrt{\frac{1}{K}\sum_{k=1}^{K} \left[\mu_q^M(\tau_k) - \mu_q^{\mathrm{ref}}(\tau_k)\right]^2}}{s_q}. \tag{29}$$

Expected waiting-time error (TimeErr). Let N_{eval} be the number of evaluated states, $\Gamma(x_n)$ the base KMC total escape rate, and $\Lambda_M(z_n)$ the learned total hazard at the same state. Their conditional mean waiting times are $\overline{\Delta\tau}_n = 1/\Gamma(x_n)$ and $\widehat{\Delta\tau}_n = 1/\Lambda_M(z_n)$. Using natural logarithms, we define

$$\mathrm{TimeErr}(M) = \frac{1}{N_{\mathrm{eval}}} \sum_{n=1}^{N_{\mathrm{eval}}} \left| \ln \frac{\widehat{\Delta\tau}_n}{\overline{\Delta\tau}_n} \right| = \frac{1}{N_{\mathrm{eval}}} \sum_{n=1}^{N_{\mathrm{eval}}} \left| \ln \frac{\Gamma(x_n)}{\Lambda_M(z_n)} \right|. \tag{30}$$

B Additional Results

B.1 Additional Memory Ablations

The main text reports memory ablations in the dilute regime (0.18% Cu). Here we further provide the corresponding ablation in a relatively concentrated setting (1.34% Cu) to test whether the role of memory persists when the alloy exhibits richer local competition and stronger structural interactions.

Table 3: **Additional memory ablation at 1.34% Cu.** Short/Long denote short-term event memory and long-term structural memory. AF is the step-to-target acceleration factor over Zacros traditional KMC. TimeErr is log expected waiting-time MAE; DynErr is real-KMC-time aligned NRMSE to Zacros.

Variant	Short	Long	AF↑	Agg.↑	TimeErr↓	DynErr↓	Arr.Hit↑
Snapshot-only AtomWorld	✗	✗	11.35×	0.169	0.342	0.332	68.80%
No short-term memory	✗	✓	21.33×	0.202	0.219	0.206	70.85%
No long-term memory	✓	✗	28.36×	0.257	0.255	0.282	75.71%
Full AtomWorld-Mem	✓	✓	**51.38×**	**0.358**	**0.141**	**0.165**	**83.82%**

Table 3 shows the same qualitative pattern as the dilute regime, confirming that the benefit of memory is not concentration-specific. Snapshot-only spatial reasoning already improves over passive sampling, but remains substantially weaker than the full model. Adding either memory scale improves both progress and fidelity, while the full AtomWorld-Mem performs best across all metrics.

The ablation also clarifies the complementary roles of the two memory scales. The full model combines both: short-term memory supports local temporal reasoning, whereas long-term memory preserves slow structural context across the trajectory. Together, these results further support the state-restoration hypothesis: the advantage of AtomWorld-Mem comes from recovering hidden evolutionary state across time, rather than from a stronger snapshot-level event scorer.

B.2 Long-Horizon Visualization Details

We further visualize long-horizon structural evolution to verify that accelerated progress corresponds to physically meaningful mesoscale organization rather than only improved scalar energy descent.

Figure 6 provides a qualitative complement to the quantitative metrics in the main text. Trajectories are rendered at matched physical-time anchors, with fixed visualization settings across methods. Cu atoms are grouped into clusters using the same neighborhood-based rule used for aggregation metrics, so that visual differences reflect trajectory behavior rather than rendering choices.

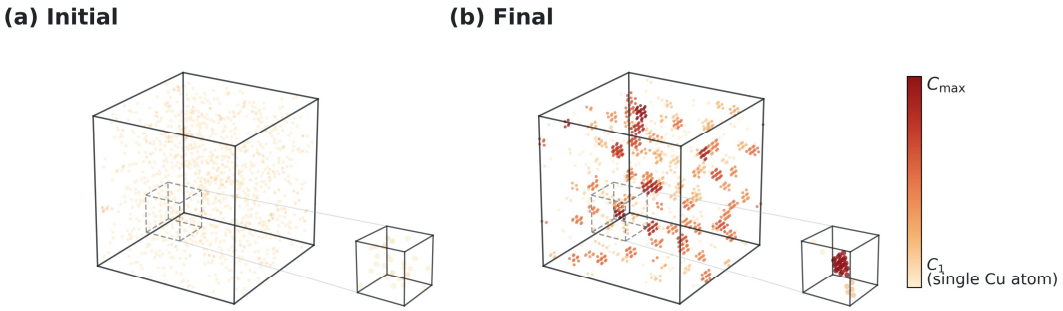


Figure 6: **Long-horizon Cu clustering visualization.** Cu clustering evolution in Fe–0.67 at.% Cu over a 50-year physical-time window. The visualization illustrates whether memory-restored event prioritization produces coherent long-horizon structural organization rather than only faster energy reduction.

The visualization supports the same conclusion as the trajectory-level metrics: AtomWorld-Mem does not merely exploit a local energy shortcut. Instead, the memory-restored state guides legal vacancy-mediated events toward coherent clustering and sustained long-horizon structural evolution. In this view, short-term memory helps capture recent vacancy–cluster interactions, while long-term memory preserves the slow structural context needed for persistent aggregation over physical time.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [N/A].
- [N/A] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for [N/A]).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will also be asked to include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While [Yes] is generally preferable to [No], it is perfectly acceptable to answer [No] provided a proper justification is given (e.g., error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering [No] or [N/A] is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist"**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers**.

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state the main claim that long-horizon atomistic evolution is a memory-based world-state restoration problem, and the experiments evaluate fixed-budget progress, trajectory-level fidelity, memory ablations, and zero-shot transfer. The scope is limited to KMC-constrained atomistic evolution and is further clarified in the conclusion and limitations.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification:The conclusion explicitly discusses that the current experiments focus on vacancy-mediated BCC alloy evolution and that future work should extend the framework to richer materials, more complex event catalogs, stronger clock-consistency objectives, and more expressive memory mechanisms.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification:The paper does not present formal theoretical theorems or proofs. The main mathematical definitions, objectives, event-selection rules, time-consistency equations, and evaluation metrics are numbered and cross-referenced in the main text and appendix.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The paper describes the AtomWorld/KMC setting, baselines, model components, reward shaping, metrics, and training/evaluation configurations in the main text and appendix. The appendix includes configuration tables and formal definitions of AF, TimeErr, and DynErr.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification:We are actively preparing the repository to ensure clarity and usability for the community

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.

- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The experimental setup describes the Fe–Cu KMC environment, step-controlled protocol, baselines, and evaluation metrics. The appendix reports training/evaluation settings, observation and memory configuration, reward shaping, optimizer, PPO/rMAPPO hyperparameters, and compute resources.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The main trajectory plots include shaded variability across runs, and the reported metrics are averaged over evaluation seeds. The appendix defines how trajectory-level metrics are aligned and averaged over seeds.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.

- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The paper reports that learning experiments use NVIDIA A100 GPUs, and the appendix summarizes the training and evaluation configuration. Wall-clock training time is recorded as part of the experimental configuration.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The work studies physics-constrained atomistic simulation and does not involve human subjects, private data, surveillance, sensitive personal attributes, or deceptive generation. We have reviewed the NeurIPS Code of Ethics and do not identify violations.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The work may positively impact efficient materials simulation and scientific discovery by reducing the cost of long-horizon atomistic evolution. Potential negative impacts are indirect, such as misuse of accelerated materials simulation for undesirable materials design; the method itself does not operate on human data or deploy a user-facing decision system.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper does not release a high-risk generative model, scraped dataset, human-subject dataset, or pretrained model with clear misuse risks. The method is a physics-constrained atomistic simulation framework.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The paper credits the KMC reference simulator and related methodological foundations in the references. No human-subject datasets or scraped assets are used. Any third-party software dependencies should be cited or documented in the released code package.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has

curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The submission introduces a method and experimental protocol, but does not release a new dataset or standalone public benchmark asset at submission time. If code or generated trajectories are released later, documentation and usage instructions will be provided with the release.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing, user studies, or human-subject research.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve human subjects, personal data, or user studies, so IRB approval is not applicable.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs are not used as part of the core method, model architecture, training pipeline, data generation, or evaluation. Any ordinary writing or editing assistance, if used, does not affect the scientific contribution or methodology.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.