

GLAM: Global-Local Variation Awareness in Mamba-based World Model

Qian He^{1,2,3*}, Wenqi Liang^{1,2,3*}, Chunhui Hao⁴, Gan Sun⁵, Jiandong Tian^{1†}

¹State Key Laboratory of Robotics, Shenyang Institute of Automation, Chinese Academy of Sciences.

²Institutes for Robotics and Intelligent Manufacturing, Chinese Academy of Sciences.

³University of Chinese Academy of Sciences.

⁴Shenyang University of Chemical Technology.

⁵School of Automation Science and Engineering, South China University of Technology.
{heqian0325, liangwenqi0123, sungan1412}@gmail.com, {haochunhui, tianjd}@sia.cn

Abstract

Mimicking the real interaction trajectory in the inference of the world model has been shown to improve the sample efficiency of model-based reinforcement learning (MBRL) algorithms. Many methods directly use known state sequences for reasoning. However, this approach fails to enhance the quality of reasoning by capturing the subtle variation between states. Much like how humans infer trends in event development from this variation, in this work, we introduce **Global-Local variation Awareness Mamba-based world model (GLAM)** that improves reasoning quality by perceiving and predicting variation between states. GLAM comprises two Mamba-based parallel reasoning modules, GMamba and LMamba, which focus on perceiving variation from global and local perspectives, respectively, during the reasoning process. GMamba focuses on identifying patterns of variation between states in the input sequence and leverages these patterns to enhance the prediction of future state variation. LMamba emphasizes reasoning about unknown information, such as rewards, termination signals, and visual representations, by perceiving variation in adjacent states. By integrating the strengths of the two modules, GLAM accounts for higher-value variation in environmental changes, providing the agent with more efficient imagination-based training. We demonstrate that our method outperforms existing methods in normalized human scores on the Atari 100k benchmark.

Code — <https://github.com/GLAM2025/glam>

Introduction

A key challenge in deep reinforcement learning (DRL) is the lack of sample efficiency (Schrittwieser et al. 2020; Vinyals et al. 2019). Many model-based reinforcement learning (MBRL) (Sutton 1991) algorithms tackle this challenge by learning models of the environment to improve the training efficiency of agents. Inspired by human cognitive processes, an MBRL approach trains a network known as a *world model* to simulate the environment and subsequently train agents using this model. The concept of world model

*These authors contributed equally.

†Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

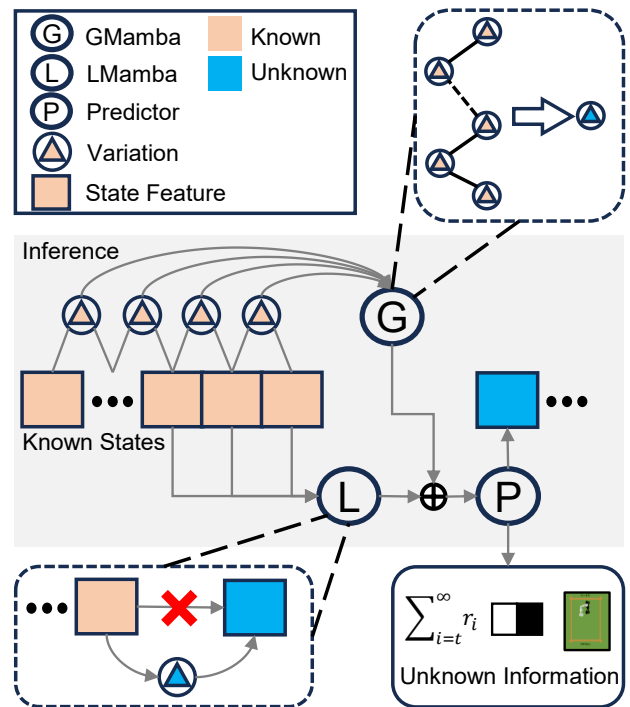


Figure 1: A single-step demonstration of GLAM inference. GLAM leverages GMamba and LMamba for global and local state variation, respectively. GMamba captures patterns of variation changes within the sequence, while LMamba incorporates variation awareness when inferring unknown information, rather than relying on direct inference alone.

has achieved breakthroughs in various decision-making domains.

These algorithms utilize interaction data between the agent and the real environment to train a parameterized world model. The agent can generate an unlimited amount of training data through interactions with this world model, thus enhancing its decision-making capabilities. This method of generating data is analogous to *imagination* of the human. However, this imagination-based self-supervised learning often fails if the data provided by the world model are not sufficiently accurate. Therefore, the lack of accuracy

in imagination is a bottleneck for the world model to improve sample efficiency.

Numerous endeavors have been made to improve the imaginative reasoning quality of the world model. For instance, STORM (Zhang et al. 2024) proposed a Transformer-based (Vaswani et al. 2017) world model that introduces random noise into the Transformer’s prediction of environmental states to enhance the robustness of agent training in imagination reasoning. However, it predicts unknown information directly from known states during imagination reasoning, leading to the inability to leverage higher-value differential information for high-quality reasoning. On the other hand, some studies have shown that using long sequence trajectories as input can enhance the imaginative capabilities of the world model, such as in the cases of IRIS (Micheli, Alonso, and Fleuret 2022) and TWM (Robine et al. 2023). However, predicting based on entire long sequences requires processing a large amount of low value information, making it difficult to efficiently infer in long sequence trajectories. Although there have been significant advances in these methods, they have not been able to efficiently leverage the state variation in known trajectories to further improve the imaginative quality of the world model.

To address the challenges outlined above, we explore a novel world model inference framework. Inspired by the human cognitive tendency to focus on variation between pieces of information during reasoning, we propose that enhancing the quality of reasoning-based imaginative interaction in the world model requires equipping them with a similar capacity to infer key information from variation. This potential can be characterized in the world model through the following two capabilities:

- **Local variation awareness:** This refers to the world model’s ability to improve its predictions of unknown information (*i.e.*, reward, termination signal) by perceiving variation between states in adjacent short sequences.
- **Global variation awareness:** The model needs to efficiently summarize patterns of variation in long sequences, without being distracted by low value information. The quality of imagination in world model can be significantly enhanced by leveraging these patterns to predict future variation.

In this work, we introduce Global-Local variation Awareness Mamba-based world model (GLAM), a novel model-based reinforcement learning (MBRL) framework that enhances sample efficiency in agent training through high-quality imagination reasoning. In the GLAM, we use Mamba (Gu and Dao 2023) as the inference network and design two modules: the Local variation awareness Mamba module (LMamba) and the Global variation awareness Mamba module (GMamba) to capture variation information. In the inference process GLAM, as illustrated in Fig. 1, LMamba infers unknown information by predicting variation between adjacent states. Meanwhile, GMamba captures patterns of variation across environmental states from long-sequence trajectories and utilizes these patterns to predict future variation distributions. Specifically, GLAM combines LMamba and GMamba to jointly infer future states and un-

known information, thereby achieving variation awareness imagination reasoning. To the best of our knowledge, our work is among the first to emphasize the importance of variation awareness in world model reasoning. Additionally, we design a phased approach to increase the number of interactions between the agent and the world model in each imagined training to enhance the training efficiency of the agent within the world model.

Specifically, our contributions are as follows:

- We introduce LMamba, a Mamba-based inference module that enhances the efficiency of world model reasoning for unknown information by perceiving variation.
- To improve the accuracy of imagined trajectories, we design GMamba, which enhances the world model’s ability to predict variation in future states by capturing the global patterns of variation within known trajectories.
- To achieve global-local variation awareness, we design GLAM, a world model framework first employs parallel modules in inference. This framework provides a novel insight into the development of effective world model.

Related Work

World Model

The *world model* is a type of reinforcement learning method within the realm of Model-Based Reinforcement Learning (MBRL), which is originally proposed in the context of world model (Ha and Schmidhuber 2018) for exploring generative neural network models in RL. Under the advantage of training agent in imagination, PlaNet (Hafner et al. 2019b) trains agent using fewer training data than model-free RL methods and even outperforms them. Besides, PlaNet introduces the recurrent state space model (RSSM) as a core component. Many subsequent works (Hafner et al. 2019a, 2020, 2023) adopt RSSM as inference module. Recent research use state space models (SSMs) as the central structure for the world model. S4WM (Deng, Park, and Ahn 2024) builds on S4 (Gu, Goel, and Ré 2021) and proposes the first universally applicable world model framework compatible with SSMs. Other world model approaches, such as IRIS (Micheli, Alonso, and Fleuret 2022), TWM (Robine et al. 2023), STORM, and TransDreamer (Chen et al. 2022), integrate Transformers into the world model, enabling efficient parallel training and leveraging self-attention (Vaswani et al. 2017; Dosovitskiy 2020; Dong et al. 2023) mechanisms for imaginative reasoning. Some works (Alonso et al. 2024; Ding et al. 2024) utilize the latent diffusion model (Rombach et al. 2022; Sun et al. 2024) to predict the next state as the world model and achieve advancements in gaming. World model has also made significant breakthroughs in other areas, such as robotic manipulation (Wu et al. 2024; Seo et al. 2022; Liang et al. 2024), autonomous driving (Gao et al. 2024; Min et al. 2024; Wang et al. 2024; Pan et al. 2023), classic control (Ha, Kim, and Kim 2023).

State Space Model

State space sequence models (SSMs) originated from classical state-space models (Kalman et al. 1960). In recent research, SSMs have proven to be effective in handling long

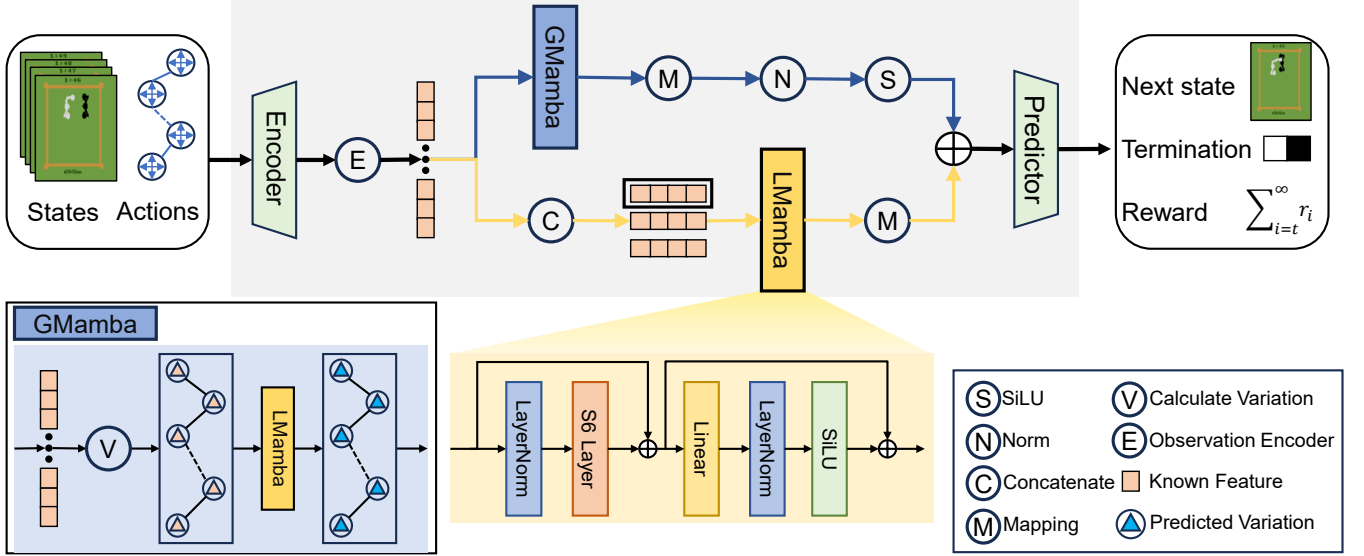


Figure 2: Overview of the Global-Local variation Awareness Mamba-based world model (GLAM). Based on the known sequence, GLAM infers future information to train the agent in imagination. In each inference, GLAM captures the awareness of global and local variation from long and short sequences and integrates both for prediction.

sequences. The latest SSM, Mamba (Gu and Dao 2023), introduces selective scanning and hardware-aware algorithms to enhance the awareness of sequential data and computational efficiency in the SSM framework. Some works (Yan, Gu, and Rush 2024; Zhu et al. 2024; Shi, Dong, and Xu 2024) use SSMs in computer vision tasks, particularly those involving image data. For example, (Islam and Bertasius 2022; Wang et al. 2023; Nguyen et al. 2022) achieve breakthrough results by incorporating S4 into various video tasks.

Method

Preliminary

State space models (SSMs) originate from a method in continuous systems for handling continuous sequence data. This method accomplishes the mapping of sequences from $x(t) \in \mathbb{R} \Rightarrow y(t) \in \mathbb{R}$ through four weight parameters $\mathbf{A} \in \mathbb{R}^{N \times N}$, $\mathbf{B} \in \mathbb{R}^{N \times 1}$, $\mathbf{C} \in \mathbb{R}^{1 \times N}$ and $\mathbf{D} \in \mathbb{R}$, including a hidden state $h(t) \in \mathbb{R}^{N \times 1}$. Here N represents the dimension of parameters. This method can be described by the ordinary differential equation (ODE):

$$\begin{aligned} h(t) &= \mathbf{A}h(t-1) + \mathbf{B}x(t), \\ y(t) &= \mathbf{C}h(t) + \mathbf{D}x(t), \end{aligned} \quad (1)$$

where $x(t)$ and $y(t)$ represent the continuous input signal and continuous output signal in the time domain. In order to adapt to computers, SSMs discretize the parameters in the above method. $\bar{\mathbf{A}}, \bar{\mathbf{B}}, \bar{\mathbf{C}}, \bar{\mathbf{D}}$ represent the corresponding parameters obtained through the zero-order holder, respectively.

Mamba, also known as S6 layer, removes the parameter matrix $\bar{\mathbf{D}}$, and its mapping formula is as follows:

$$\begin{aligned} h_t &= \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}x_t, \\ y_t &= \bar{\mathbf{C}}h_t. \end{aligned} \quad (2)$$

On the other hand, Mamba introduces a selective scanning algorithm in order to focus on more important segments in the sequence. It replaces $\bar{\mathbf{B}}, \bar{\mathbf{C}}$ with two higher-dimensional parameter matrices $\bar{\mathbf{B}}, \bar{\mathbf{C}} \in \mathbb{R}^{N \times N}$ that are mapped by the discretized input x_t . The specific formula is as follows:

$$\begin{aligned} \bar{\mathbf{B}} &= \text{Linear}(x_t), \\ \bar{\mathbf{C}} &= \text{Linear}(x_t). \end{aligned} \quad (3)$$

This design results in a historical feature vector $h_t \in \mathbb{R}^{N \times N}$ of higher dimension than previous SSMs. For efficient calculation, Mamba introduces the hardware-aware algorithm in inference on sequence. As a result, Mamba is unable to output the historical feature vector h_t in parallel like S5 (Smith, Warrington, and Linderman 2022).

Besides, the selective scanning algorithm enables Mamba to focus more effectively on important segments of the data, compared to the previous SSMs. Although Transformers use a self-attention mechanism for selective scanning, the causal inference form inside Mamba aligns better with the causal relationship where variation triggers environmental information changes in the interaction process. Therefore, we utilize Mamba as the inference network.

We introduce GLAM, a world model that improves the quality of inference through awareness and prediction of variation between states. An overview of GLAM is shown in Fig. 2. GLAM leverages two modules, LMamba and GMamba, to perform differential-based reasoning. LMamba infers unknown information and predicts future state by estimating potential variation in the future states. GMamba captures the patterns of variation in environmental state variation over long sequences and uses these patterns to predict the distribution of future variation. Finally, GLAM combines the inference results based on both local and global variation to generate the final prediction of the model.

Mamba-based World Model

The overall structure of GLAM is shown in Fig. 2. The input to the model includes the sequence of observations of the environment $o_{\leq T} = \{o_t\}_{t=0}^T$ and the corresponding action sequences $a_{\leq T} = \{a_t\}_{t=0}^T$ determined by the agent. Consistent with previous works, we use the image encoder q_ϕ and the image decoder p_ϕ , which are implemented with convolutional neural networks (CNNs) (LeCun et al. 1989), to process image observation o_t . The feature vector z_t represents the features sampled from \mathcal{Z}_t to represent o_t .

$$\begin{aligned} \text{Image Encoder: } z_t &\sim \mathcal{Z}_t = q_\phi(o_t), \\ \text{Image Decoder: } \hat{o}_t &= p_\phi(z_t). \end{aligned} \quad (4)$$

We design two parallel inference modules, LMamba and GMamba, denoted by M_ϕ^l and M_ϕ^s , to process the global and local variation in sequence, respectively. During the inference process, the known action sequence $a_{\leq T}$, and the image feature sequence $z_{\leq T} = \{z_t\}_{t=0}^T$ are mixed by the feature encoder f_ϕ to obtain the feature vector $e_{\leq T} = \{e_t\}_{t=0}^T$. Based on the feature vector, the two modules predict two distributions of the t -th moment, u_t^l and u_t^g . Here, u_t^l represents the prediction based on the local variation awareness, while u_t^g reflects the prediction based on the variation awareness of the entire known trajectory. Finally, we combine the results to predict future states distribution $\hat{\mathcal{Z}}_{t+1}$ and unknown information $\hat{r}_t, \hat{c}_t, \hat{o}_t$.

$$\begin{aligned} \text{Feature Encoder: } e_t &= f_\phi(z_{\leq t}, a_{\leq t}), \\ \text{GMamba: } u_t^g &= M_\phi^g(e_{\leq t}^g), \\ \text{LMamba: } u_t^l &= M_\phi^l(e_{\leq t}^l), \\ \text{Dynamics Predictor: } \hat{\mathcal{Z}}_{t+1} &= g_\phi^D(u_t^g, u_t^l), \\ \text{Reward Predictor: } \hat{r}_t &= g_\phi^R(u_t^g, u_t^l), \\ \text{Continuation Predictor: } \hat{c}_t &= g_\phi^C(u_t^g, u_t^l). \end{aligned} \quad (5)$$

Here, $f_\phi, g_\phi^D, g_\phi^R, g_\phi^C$ are implemented with multi-layer perceptrons (MLPs).

Inputting a longer known sequence can enhance the quality of imagination in world model. However, the presence of large amounts of redundant information in long sequences consumes significant computational resources and can mislead the model’s ability to perceive critical variation. To address this, we clearly define the roles of LMamba and GMamba in processing long sequences. LMamba focuses on inferring unknown information through local discrepancy awareness, and thus, its input is a shorter sequence $e_{\leq T}^l = \{e_t\}_{t=T+1-s}^T$ of length s . GMamba aims to capture the patterns of global discrepancy changes, and its input is a longer sequence $e_{\leq T}^g = \{e_t\}_{t=T+1-l}^T$ of length l . In this work, we use fixed length $l = 16$ and $s = 4$.

Local variation awareness Mamba Aiming to enforce the local variation awareness ability of world model, we propose local variation awareness Mamba (LMamba). Unlike other world models that directly inference future state and unknown information, LMamba focuses more on the prediction of variation in feature. It infers future states and un-

known information of the current state by predicting the distribution of variation between future states and the current state. This inference method allows LMamba to focus more on improving its awareness of local variation of the environment during training.

As the input $e_{\leq T}^l$ is a feature sequence that contains both states and actions, LMamba can infer the possible change distribution of the environmental state based on the actions from it, and finally output a sequence, $u_{\leq T}^l = \{u_t^l\}_{t=T-s}^T$, of the same length, each vector u_t^l in the sequence serves as the module’s prediction for the t -th moment. The inference formula of LMamba is as follows:

$$\begin{aligned} \text{Mamba: } s_{\leq t} &= M_\phi(Norm(e_{\leq t}^l)), \\ u_{\leq t}^l &= e_{\leq t} + Norm(s_{\leq t}), \end{aligned} \quad (6)$$

where, $Norm$ refers to a normalization method in (Ba, Kiros, and Hinton 2016) and M_ϕ represents the Mamba.

Parallel scanning is a key feature for efficiently training the world model (Deng, Park, and Ahn 2024), and modifications to the dimensions of internal parameter in Mamba pose challenges to this feature in existing SSM-based world model frameworks. Therefore, we design LMamba to directly use the final output of Mamba for prediction during inference, rather than relying on the historical features h_t in Eq. (2). Our ablation studies demonstrate that even using LMamba as the sole inference module, the world model outperforms the baseline.

Global variation awareness Mamba To efficiently summarize patterns of variation in long sequences without being distracted by low value information, we developed GMamba, a module designed to be globally aware of variation changes within historical state sequences.

The single-step inference process of GMamba is illustrated in Eq. (7). Based on the input feature sequence $e_{\leq T}^g = \{e_t\}_{t=T-l}^T$, GMamba computes the sequence of variation between states, denoted as $d_{\leq T-1}^g = \{d_t\}_{t=T-l}^{T-1}$. Then it analyzes this variation sequence to summarize the distribution patterns of state variation induced by actions in the environment. Finally, GMamba outputs predictions $u_{\leq T}^g = \{u_t^g\}_{t=T+1-l}^T$ based on the awareness of global variation.

$$\begin{aligned} \text{Global Variation: } d_{\leq t-1}^g &= e_{\leq t}^g - e_{\leq t-1}^g, \\ \text{Mamba: } s_t &= M_\phi(Norm(d_{\leq t-1}^g)), \\ u_t^g &= d_{t-1}^g + Norm(s_t). \end{aligned} \quad (7)$$

Actually, the changes in variation distribution are more like subtle adjustments on states feature. Therefore, we map u_t^g and u_t^l in different ways. After being mapped by MLPs, the prediction result u_t^g from M_ϕ^g is normalized by Layer-Norm (Ba, Kiros, and Hinton 2016) and activated by SiLU (Elfwing, Uchibe, and Doya 2018) to retain only the important variation contained in u_t^l . The prediction result u_t^l from M_ϕ^l just goes through a simple mapping of MLPs, as we need it to retain more state information. Specific algorithmic details are demonstrated in Algorithm 1.

During agent training, GLAM only needs to output the prediction of the next moment, as shown in Fig. 2. But in

the training process of the world model, to achieve parallel inference for M_ϕ^l , we concatenate the sequence $\{e_i\}_{i=0}^t$ into a short sequence block $E^s = \{\{e_{i+1-s}, \dots, e_i\}\}_{i=3}^t$ and input E^s to M_ϕ^l for parallel inference. We demonstrate the parallel inference steps in GLAM in Algorithm 1.

Loss Function The world model of GLAM incorporates multiple network structures, including Mamba, CNNs, and MLPs. We jointly optimize all networks within the world model in a self-supervised manner. The total loss function value is calculated by the formula in Eq. (8).

$$\mathcal{L}(\phi) = \mathcal{L}_{PRED}(\phi) + \mathcal{L}_{DYN}(\phi) + \mathcal{L}_{REP}(\phi) + \mathcal{L}_{VAR}(\phi). \quad (8)$$

The total loss consists of four components. Following DreamerV3, \mathcal{L}_{PRED} represents the world model’s predictive loss for reward \hat{r}_t , observation \hat{o}_t and termination signals \hat{c}_t . The dynamics loss \mathcal{L}_{DYN} emphasizes the model’s prediction of the feature distribution of future states, while the representation loss \mathcal{L}_{REP} slightly adjusts the encoder to approach the predicted feature distribution. \mathcal{L}_{DYN} and \mathcal{L}_{REP} are expressed as Kullback–Leibler (KL) divergences. The formula of these three components is as follows:

$$\begin{aligned} \mathcal{L}_{PRED}(\phi) &= \mathcal{L}_{SYM}(\hat{r}_t, r_t) \\ &\quad + \|\hat{o}_t - o_t\|_2 \\ &\quad + [c_t \log(\hat{c}_t) + (1 - c_t) \log(1 - \hat{c}_t)], \\ \mathcal{L}_{DYN}(\phi) &= \max(1, \text{KL}[sg(q_\phi(o_t)) \| g_\phi^D(u_t^g, u_t^l)]), \\ \mathcal{L}_{REP}(\phi) &= \max(1, \text{KL}[q_\phi(o_t) \| sg(g_\phi^D(u_t^g, u_t^l))]), \end{aligned} \quad (9)$$

where \mathcal{L}_{SYM} denotes the symlog two-hot loss, $sg(\cdot)$ denotes the operation of stop-gradients.

We specifically design variation loss function for M_ϕ^g to constrain its convergence direction relative to M_ϕ^l . The optimization objective for M_ϕ^g is to predict variation information in the next frame based on the pattern in the global sequence. We quantify this objective by measuring the variation between feature vectors across consecutive frames, noted as Δo_t . This loss function is noted as \mathcal{L}_{VAR} in Eq. (10).

$$\begin{aligned} \Delta o_t &= sg(z_{t+1} - z_t), \\ \mathcal{L}_{VAR}(\phi) &= \max(1, \text{KL}[\Delta o_t \| g_\phi^D(u_t^g)]). \end{aligned} \quad (10)$$

Variable Number of Interaction In GLAM, agent is trained entirely in the imagined interaction with world model. We found the number of interactions between the agent and the world model in each training session influences the agent’s final training performance. Consequently, we incorporate a variable number of interaction during imagination in GLAM. As the steps of world model training increases, the quality of the world model’s imagination gradually improves. Therefore, we adopt a phased approach to increase the number of imagined training within the world model based on the training step. The specific formula for calculating the variable number of interaction is as follows:

$$n_t = \min[n_0 + \text{Int}(t/\Delta n) * n_i, n_{max}], \quad (11)$$

Algorithm 1: Parallel inference in GLAM

Initialize: The input sequence length $l = t$, the length of the concatenate short sequences $s = 4$.

Input: Observation sequence $o_{0:t}$, Action sequence $a_{0:t}$.

▷ **Encode feature sequence** $e_{0:t}$:

$$\begin{aligned} z_{0:t} &\sim \mathcal{Z}_{0:t} = q_\phi(o_{0:t}); \\ e_{0:t} &= f_\phi(z_{0:t}, a_{0:t}); \end{aligned}$$

▷ **Global variation inference:**

$$\begin{aligned} d_{0:t-1} &= e_{1:l} - e_{0:t-1}; \\ u_{1:t}^g &= LMamba(d_{0:t-1}); \\ u_{1:t}^g &= LayerNorm(u_{1:t}^g); \\ u_{1:t}^g &= SiLU(u_{1:t}^g); \end{aligned}$$

▷ **Local variation inference:**

$$\begin{aligned} &\text{Concatenate } e_{0:t} \text{ into a short sequence block } E^s; \\ \{e_0 \dots e_t\} &\Rightarrow E^s = \{\{e_{i+1-s}, \dots, e_i\}\}_{i=3}^t; \\ u_{3:t}^g &= LMamba(E^s); \end{aligned}$$

▷ **Predict unknown information and next state:**

$$\begin{aligned} \hat{\mathcal{Z}}_{4:t+1} &= g_\phi^D(u_{3:t}^g, u_{3:t}^l); \\ \hat{r}_{3:t} &= g_\phi^R(u_{3:t}^g, u_{3:t}^l); \\ \hat{c}_{3:t} &= g_\phi^C(u_{3:t}^g, u_{3:t}^l); \end{aligned}$$

Return: The unknown information sequence $\hat{r}_{3:t}, \hat{c}_{3:t}$ and the distribution $\hat{\mathcal{Z}}_{4:t+1}$ of the state sequence.

where n_t represents the variable number of interaction, t denotes the current training step, n_f represents the change frequency of n_t , Δn corresponds to the increment of number at each update and n_{max} represents the maximum number of interaction during imagined training. For other parts of the agent’s training, we refer to the method in STORM.

Experiment

We evaluate GLAM on a subset of 26 games from Atari 100k (Bellemare et al. 2013), a benchmark that is widely used for testing reinforcement learning algorithms. We first introduce the benchmark and the RL methods that are used for comparison, then analyze the comprehensive results, and finally present our ablation study of GLAM.

Baselines

We compare GLAM with the following model-based DRL methods: SimPLe (Kaiser et al. 2019), TWM, IRIS, DreamerV3, Hieros (Mattes, Schlosser, and Herbrich 2023), and STORM. SimPLe and DreamerV3 uses Recurrent State Space Models (RSSM) as the core module for world model inference. Hieros builds a S5-based world model. IRIS, TWM, and STORM uses Transformers as their backbones. IRIS employs long sequences containing all time steps for imagination reasoning, with the sequence length growing as the inference steps progress. TWM performs inference based on fixed-length sequences, while STORM relies only on historical features from the previous time step. Similar to previous work, we aim to validate the effectiveness of GLAM, so we do not compare with look-ahead search methods like MuZero (Schrittwieser et al. 2020) and EfficientZero (Ye et al. 2021).

Same with recent works (Zhang et al. 2024), we train

Game	Random	Human	SimPLe	TWM	IRIS	DreamerV3	Hieros	Storm	GLAM(Ours)
Alien	228	7128	617	675	420	959	828	984	1065
Amidar	6	1720	74	122	143	139	127	205	202
Assault	222	742	527	683	1524	706	1764	801	646
Asterix	210	8503	1128	1116	854	932	899	1028	504
Bank Heist	14	753	34	467	53	649	177	641	477
Battle Zone	2360	37188	4031	5068	13074	12250	15140	13540	13190
Boxing	0	12	8	78	70	78	65	80	76
Breakout	2	30	16	20	84	31	10	16	11
Chopper Command	811	7388	979	1697	1565	420	1475	1888	1683
Crazy Climber	10780	35829	62584	71820	59234	97190	50857	66776	71872
Demon Attack	152	1971	208	350	2034	303	1480	165	209
Freeway	0	30	17	24	31	0	31	0	23
Frostbite	65	4335	237	1476	259	909	2901	1316	2792
Gopher	258	2413	597	1675	2236	3730	1473	8240	3149
Hero	1027	30826	2657	7254	7037	11161	7890	11044	7599
James Bond	29	303	101	362	463	445	939	509	738
Kangaroo	52	3035	51	1240	838	4098	6590	4208	1900
Krull	1598	2666	2204	6349	6616	7782	8130	8413	12649
Kung Fu Master	256	22736	14862	24555	21760	21420	18793	26182	32680
Ms Pacman	307	6952	1480	1588	999	1327	1771	2673	2094
Pong	-21	15	13	19	15	18	5	11	18
Private Eye	25	69571	35	87	100	882	1507	7781	100
Qbert	164	13455	1289	3331	746	3405	770	4522	2434
Road Runner	12	7845	5641	9109	9615	15565	16950	17564	20030
Seaquest	68	42055	683	774	661	618	560	525	511
Up N Down	533	11693	3350	15982	3546	7667	—	7985	7417
Human Mean	0%	100%	33%	96%	105%	112%	120%	122.3%	130.6%
Human Median	0%	100%	13%	51%	29%	49%	56%	58.4%	61.8%

Table 1: Mean scores of models on the 26 games of Atari 100k benchmark, along with the overall human-normalized scores. As per the convention, model results with score fluctuations within 5% are indicated in bold font.

Game	Pong	Boxing	Kung.	Battle.
Dbl.Transformer	4	68	24030	7600
Dbl.Mamba	3	49	23790	6150
Ours w/o G&L	8	63	26984	8320
Ours w/o G	14	71	29475	9570
Ours w/o \mathcal{L}_{VAR}	11	75	26210	8600
Ours	18	76	32680	13190

Table 2: Ablation study results on inference module.

GLAM using 100k samples in each game. Considering a skip step of 4 frames, the samples correspond to 400k actual game frames, which is about 1.85 hours (Zhang et al. 2024) of real-time game time. The final results of the agent are quantified using the human-normalized score: $S_{Norm} = (S_{Agent} - S_{Random}) / (S_{Human} - S_{Random})$.

Results on Atari 100k

In our experiment we train GLAM with 5 different seeds in each game. To evaluate the agent, we perform 20 evaluations for final checkpoints and compute the average scores as the

results. The results in Tab. 1 are the average scores obtained using the final checkpoints.

As shown in Tab. 1, our method surpasses all baselines in terms of human-normalized scores in mean and median. Notably, GLAM excels in games such as *KungFuMaster*, *Krull*, *RoadRunner*, and *Alien*, which are strongly associated with variation. This can be attributed to GMamba and LMamba, enabling the world model to improve reasoning quality through variation awareness. We present a comparison of the training results between GLAM and STORM, the state-of-the-art Transformer-based world model framework, on partial games in Fig. 3.

Ablation Studies

Inference module in GLAM We conduct ablation studies on the design of the inference modules on *Pong*, *Boxing*, *KungFuMaster* and *BattleZone*. **Dbl.** indicates the combination of two identical modules in direct parallel. To validate the effectiveness of GMamba and LMamba, we choose STORM as the baseline, labeled as **ours w/o G&L**. We replace the Transformer in STORM with an equivalent layer number of LMamba, and labeled the new world model as **ours w/o G**. We also conduct an ablation study on the loss function \mathcal{L}_{VAR} in GMamba. The average scores of the final




Figure 3: Comparison of the training results between GLAM and STORM.




Figure 4: Ablation study on variable imagine steps (left) and number of layers in the Mamba (right).




Figure 5: Loss value in the training process of ours w/o G and ours w/o G&L.

checkpoints are shown in Tab. 2 and suggest that our design markedly enhances the model’s performance.

In Fig. 5, we show the fluctuation of loss values in the training process of **ours w/o G** and **ours w/o G&L**. We observe that the fluctuation amplitude of the convergence curve for the **ours w/o G** is smaller and more stable than that of **ours w/o G&L**. This stability is advantageous for the agent, as it aids in maintaining a consistent direction of convergence during imagined training. We attribute these results to the local awareness capability of LMamba, which minimizes the interference of low value information in inference.

Variable Number of Interaction We explore the impact of different maximum number of interaction on the learning effectiveness of the agent. We design three different parameters of interaction, $n_{max} = 16, 24, 32$, labeled as $N - 16, N - 24, N - 32$ respectively. The variable parameter n_t in training is calculated by Eq. (11), and the parameters $n_0 = 16$ and $\Delta n = 8$ are fixed. Fig. 4 shows the performance of GLAM in *KungFuMaster* under these three conditions. We confirm that during the training process, as the world model continuously converges on environment prediction, adopting a gradually increasing imagination training steps can improve the training efficiency of the agent.

Number of layers in the Mamba As an earlier world model framework that utilizes Mamba and parallelly employs two inference models, we explore the impact of different numbers of Mamba layers on training outcomes. We set up four different configurations, labeled $G1 - L1, G1 - L2, G2 - L1$, and $G2 - L2$, where G and L represent the number of layers in GMamba and LMamba, respectively. For example, $G1 - L1$ indicates using 1 layer of Mamba in both GMamba and LMamba. The training results in *Kung-*

FuMaster are shown in Fig. 4. Similar to STORM, the performance of model deteriorates in larger-scale model. This decline may be attributed to two factors: 1) The increased complexity of parallel inference modules makes it difficult to achieve synchronized improvement during training, leading to interference in predictions. 2) Minor variation in states can gradually distort as they propagate through multiple layers of the inference module.

Conclusion

In this study, we introduce GLAM, a world model that incorporates difference awareness in the inference process. GLAM leverages GMamba and LMamba to focus on the global variation patterns and locally variation-related unknowns in the input sequence during inference, respectively. Experimental results demonstrate that GLAM achieves state-of-the-art performance on the Atari 100K benchmark. Furthermore, GLAM is the first world model to implement multi-module parallel inference, offering a novel design paradigm for future world model.

Limitations Mamba provides computational efficiency and effective state representation in capturing long-range dependencies. But our model may still face challenges in handling uncertainty and non-temporality in nonlinear or chaotic systems. Specific limitations include: 1) difficulty in identifying the relationship between state variation and chaotic future change, 2) complex variation patterns that may impede the world model’s ability to produce stable imagine training, 3) agent struggles to explore valuable information.

Acknowledgments

We thank all reviewers for their insightful comments and valuable suggestions. This work is supported by the National Natural Science Foundation of China (No.U2013210).

References

- Alonso, E.; Jelley, A.; Micheli, V.; Kanervisto, A.; Storkey, A.; Pearce, T.; and Fleuret, F. 2024. Diffusion for World Modeling: Visual Details Matter in Atari. *arXiv preprint arXiv:2405.12399*.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279.
- Chen, C.; Wu, Y.-F.; Yoon, J.; and Ahn, S. 2022. Transdreamer: Reinforcement learning with transformer world models. *arXiv preprint arXiv:2202.09481*.
- Deng, F.; Park, J.; and Ahn, S. 2024. Facing off world model backbones: Rnns, transformers, and S4. 36.
- Ding, Z.; Zhang, A.; Tian, Y.; and Zheng, Q. 2024. Diffusion world model. *arXiv preprint arXiv:2402.03570*.
- Dong, J.; Liang, W.; Cong, Y.; and Sun, G. 2023. Heterogeneous forgetting compensation for class-incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11742–11751.
- Dosovitskiy, A. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Elfwing, S.; Uchibe, E.; and Doya, K. 2018. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks*, 107: 3–11.
- Gao, Z.; Mu, Y.; Chen, C.; Duan, J.; Luo, P.; Lu, Y.; and Li, S. E. 2024. Enhance Sample Efficiency and Robustness of End-to-End Urban Autonomous Driving via Semantic Masked World Model. *IEEE Transactions on Intelligent Transportation Systems*, 1–13.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Goel, K.; and Ré, C. 2021. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*.
- Ha, D.; and Schmidhuber, J. 2018. Recurrent world models facilitate policy evolution. 31.
- Ha, J.; Kim, K.; and Kim, Y. 2023. Dream to generalize: zero-shot model-based reinforcement learning for unseen visual distractions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 7802–7810.
- Hafner, D.; Lillicrap, T.; Ba, J.; and Norouzi, M. 2019a. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*.
- Hafner, D.; Lillicrap, T.; Fischer, I.; Villegas, R.; Ha, D.; Lee, H.; and Davidson, J. 2019b. Learning latent dynamics for planning from pixels. In *International conference on machine learning*, 2555–2565. PMLR.
- Hafner, D.; Lillicrap, T.; Norouzi, M.; and Ba, J. 2020. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*.
- Hafner, D.; Pasukonis, J.; Ba, J.; and Lillicrap, T. 2023. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*.
- Islam, M. M.; and Bertasius, G. 2022. Long movie clip classification with state-space video models. In *European Conference on Computer Vision*, 87–104. Springer.
- Kaiser, L.; Babaeizadeh, M.; Milos, P.; Osinski, B.; Campbell, R. H.; Czechowski, K.; Erhan, D.; Finn, C.; Koza-kowski, P.; Levine, S.; et al. 2019. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*.
- Kalman, R. E.; et al. 1960. A new approach to linear filtering and prediction problems [J]. *Journal of basic Engineering*, 82(1): 35–45.
- LeCun, Y.; Boser, B.; Denker, J. S.; Henderson, D.; Howard, R. E.; Hubbard, W.; and Jackel, L. D. 1989. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4): 541–551.
- Liang, W.; Sun, G.; He, Q.; Ren, Y.; Dong, J.; and Cong, Y. 2024. Never-Ending Embodied Robot Learning. *arXiv preprint arXiv:2403.00336*.
- Mattes, P.; Schlosser, R.; and Herbrich, R. 2023. Hieros: Hierarchical Imagination on Structured State Space Sequence World Models. *arXiv preprint arXiv:2310.05167*.
- Micheli, V.; Alonso, E.; and Fleuret, F. 2022. Transformers are sample-efficient world models. *arXiv preprint arXiv:2209.00588*.
- Min, C.; Zhao, D.; Xiao, L.; Zhao, J.; Xu, X.; Zhu, Z.; Jin, L.; Li, J.; Guo, Y.; Xing, J.; Jing, L.; Nie, Y.; and Dai, B. 2024. DriveWorld: 4D Pre-trained Scene Understanding via World Models for Autonomous Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15522–15533.
- Nguyen, E.; Goel, K.; Gu, A.; Downs, G.; Shah, P.; Dao, T.; Baccus, S.; and Ré, C. 2022. S4nd: Modeling images and videos as multidimensional signals with state spaces. 35: 2846–2861.
- Pan, M.; Zhu, X.; Zheng, Y.; Wang, Y.; and Yang, X. 2023. Model-Based Reinforcement Learning with Isolated Imaginations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Robine, J.; Höftmann, M.; Uelwer, T.; and Harmeling, S. 2023. Transformer-based world models are happy with 100k interactions. *arXiv preprint arXiv:2303.07109*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Schrittwieser, J.; Antonoglou, I.; Hubert, T.; Simonyan, K.; Sifre, L.; Schmitt, S.; Guez, A.; Lockhart, E.; Hassabis, D.; Graepel, T.; et al. 2020. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839): 604–609.

Seo, Y.; Lee, K.; James, S. L.; and Abbeel, P. 2022. Reinforcement learning with action-free pre-training from videos. In *International Conference on Machine Learning*, 19561–19579. PMLR.

Shi, Y.; Dong, M.; and Xu, C. 2024. Multi-Scale VMamba: Hierarchy in Hierarchy Visual State Space Model. *arXiv preprint arXiv:2405.14174*.

Smith, J. T.; Warrington, A.; and Linderman, S. W. 2022. Simplified state space layers for sequence modeling. *arXiv preprint arXiv:2208.04933*.

Sun, G.; Liang, W.; Dong, J.; Li, J.; Ding, Z.; and Cong, Y. 2024. Create your world: Lifelong text-to-image diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Sutton, R. S. 1991. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4): 160–163.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. 30.

Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.

Wang, J.; Zhu, W.; Wang, P.; Yu, X.; Liu, L.; Omar, M.; and Hamid, R. 2023. Selective Structured State-Spaces for Long-Form Video Understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6387–6397.

Wang, Y.; He, J.; Fan, L.; Li, H.; Chen, Y.; and Zhang, Z. 2024. Driving into the Future: Multiview Visual Forecasting and Planning with World Model for Autonomous Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14749–14759.

Wu, J.; Ma, H.; Deng, C.; and Long, M. 2024. Pre-training contextualized world models with in-the-wild videos for reinforcement learning. 36.

Yan, J. N.; Gu, J.; and Rush, A. M. 2024. Diffusion Models Without Attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8239–8249.

Ye, W.; Liu, S.; Kurutach, T.; Abbeel, P.; and Gao, Y. 2021. Mastering atari games with limited data. 34: 25476–25488.

Zhang, W.; Wang, G.; Sun, J.; Yuan, Y.; and Huang, G. 2024. STORM: Efficient stochastic transformer based world models for reinforcement learning. 36.

Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*.