

RAPTOR: Real-Time High-Resolution UAV Video Prediction with Efficient Video Attention

Zhan Chen^{1,2,3}, Zile Guo^{1,2,3}, Enze Zhu^{1,2,3}, Peirong Zhang^{1,2,3},
Xiaoxuan Liu^{1,2}, Lei Wang^{1,2}, Yidan Zhang^{1,2,3*}

¹Aerospace Information Research Institute, Chinese Academy of Sciences

²Key Laboratory of Target Cognition and Application Technology (TCAT)

³School of Electronic, Electrical and Communication Engineering, University of Chinese Academy of Sciences
{chenzhan21, guozile25, zhuzhenze23, zhangpeirong22, zhangyidan19}@mails.ucas.ac.cn,
{liuxiaoxuan, wanglei002931}@aircas.ac.cn

Abstract

Video prediction is plagued by a fundamental trilemma: achieving high-resolution and perceptual quality typically comes at the cost of real-time speed, hindering its use in latency-critical applications. This challenge is most acute for autonomous UAVs in dense urban environments, where foreseeing events from high-resolution imagery is non-negotiable for safety. Existing methods, reliant on iterative generation (diffusion, autoregressive models) or quadratic-complexity attention, fail to meet these stringent demands on edge hardware. To break this long-standing trade-off, we introduce RAPTOR, a video prediction architecture that achieves real-time, high-resolution performance. RAPTOR’s single-pass design avoids the error accumulation and latency of iterative approaches. Its core innovation is Efficient Video Attention (EVA), a novel translator module that factorizes spatiotemporal modeling. Instead of processing flattened spacetime tokens with $O((ST)^2)$ or $O(ST)$ complexity, EVA alternates operations along the spatial (S) and temporal (T) axes. This factorization reduces the time complexity to $O(S + T)$ and memory complexity to $O(\max(S, T))$, enabling global context modeling at 512^2 resolution and beyond, operating directly on dense feature maps with a patch-free design. Complementing this architecture is a 3-stage training curriculum that progressively refines predictions from coarse structure to sharp, temporally coherent details. Experiments show RAPTOR is the first predictor to exceed 30 FPS on a Jetson AGX Orin for 512^2 video, setting a new state-of-the-art on UAVid, KTH, and a custom high-resolution dataset in PSNR, SSIM, and LPIPS. Critically, RAPTOR boosts the mission success rate in a real-world UAV navigation task by 18%, paving the way for safer and more anticipatory embodied agents.

Code — <https://github.com/Thelegendzz/RAPTOR>

Introduction

The ability to anticipate the future is a hallmark of intelligence and a critical component for autonomous agents like UAVs navigating dynamic environments. Video prediction, the task of synthesizing future frames from past observations, offers this foresight. Yet, its practical application has

been hamstrung by a fundamental trilemma: a persistent trade-off among inference speed, input resolution, and prediction quality. For a UAV, this is not an abstract challenge; it is a direct barrier to safety. To avoid a suddenly opening door or a pedestrian stepping into its path, the agent must process high-resolution video and generate sharp predictions in milliseconds. This exposes three critical bottlenecks in prior art:

1. Prohibitive Latency. State-of-the-art quality is often achieved by iterative models like diffusion or autoregressive RNNs (Voleti, Jolicoeur-Martineau, and Pal 2022; Wang et al. 2023a). Requiring hundreds of sequential steps, their inference times are orders of magnitude too slow for the real-time control loops of UAVs on edge hardware.

2. Inability to Scale to High Resolution. Dominant attention-based predictors (Bertasius, Wang, and Torresani 2021) jointly model spacetime by treating video as a flat sequence of tokens. This leads to a computational complexity of $O((ST)^2)$, where S is the number of spatial locations and T is the number of frames. This quadratic scaling makes processing high-resolution video computationally infeasible, restricting models to low-resolution inputs (e.g., 64^2 – 256^2), thereby missing small but safety-critical details.


3. Blurry and Incoherent Motion. Many models rely on simple pixel-wise losses (e.g., MSE, L1) (Gao et al. 2022; Tan et al. 2023). These objectives average over possible futures, resulting in blurry predictions and a lack of temporal consistency. This is especially problematic at high resolutions, where motion artifacts can render predictions unreliable for navigation.

To break this trilemma, we propose RAPTOR. As laid out in our title, RAPTOR is an end-to-end, single-pass architecture designed for real-time, high-resolution prediction that directly generates all future frames at once, eliminating iterative latency and error accumulation.


The cornerstone of RAPTOR is our Efficient Video Attention (EVA) module. Instead of operating on a flattened ST-length sequence with its prohibitive $O((ST)^2)$ complexity, EVA factorizes the problem. It alternates between processing information along the temporal axis (a sequence of length T) and the spatial axis (a sequence of length S). This fundamental design choice reduces the time complexity to $O(S + T)$ and the memory complexity to $O(\max(S, T))$, enabling efficient global modeling on high-resolution fea-

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



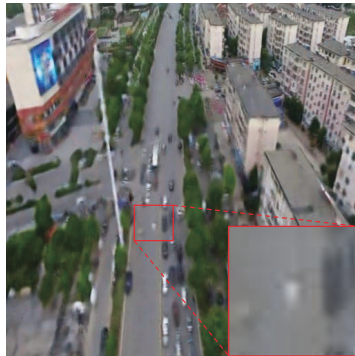
(a) Resolution-Based Computational Complexity Scaling




(b) Resolution-based Inference Memory Scaling



(c) Temporal Inconsistency



(d) Progressive Blurring and High-Frequency Detail Decay



(e) Ours

Figure 1: RAPTOR enables scalable, high-fidelity video prediction for demanding, high-resolution scenarios. (a, b) Benchmarked on an NVIDIA RTX 6000 Ada (48GB), RAPTOR’s linear complexity design maintains a low memory footprint, making it the **first framework to efficiently scale beyond 1024² resolution** as competitors hit memory ceilings (Out-of-Memory). (c, d) In contrast, conventional methods relying on pixel-wise losses often result in (c) temporal tearing and (d) edge blurring that obscures vital targets. (e) The proposed RAPTOR avoids these pitfalls, generating sharp and coherent predictions that preserve crucial details for reliable downstream planning.

ture maps for the first time. To ensure high-fidelity outputs, we complement our architecture with a three-stage curriculum learning strategy that shifts the focus from pixel reconstruction to edge sharpness and perceptual realism.

Our contributions are as follows:

- **RAPTOR:** The first video prediction framework to achieve both real-time performance (>30 FPS on Jetson AGX Orin) and high-resolution (512^2 and beyond) processing. It outperforms prior art on multiple benchmarks and improves UAV navigation success rate by 18%.
- **Efficient Video Attention (EVA):** A novel spatiotemporal modeling module that avoids the quadratic bottleneck of standard attention by using axis factorization. This reduces complexity to $O(S + T)$, enabling scalable, global receptive fields for high-resolution video without resorting to patch tokenization.
- **Three-Stage Curriculum Learning:** An effective training strategy that systematically improves prediction quality, yielding sharp edges and temporally coherent motion critical for downstream planning tasks.

Related Work

Video Prediction Paradigms

Video prediction methods fall into three main paradigms:

Autoregressive models. predict frames sequentially with recurrent units. Early ConvLSTM networks (Shi et al. 2015) and PredRNN (Wang et al. 2017), PredRNN++ (Wang et al. 2018) introduced spatiotemporal LSTMs and gradient highways. Subsequent works like MIM (Wang et al. 2019) and MotionRNN (Wu et al. 2021) add memory or motion modules. These excel at 64^2 – 256^2 but incur latency and memory increasing linearly with the horizon.

Diffusion-based models. adapt denoising diffusion to video. DDPMs (Ho, Jain, and Abbeel 2020) and DDIMs (Song, Meng, and Ermon 2021) achieve high fidelity but require hundreds of iterative steps. Conditional video diffusion models like MCVD (Voleti, Jolicœur-Martineau, and Pal 2022) and ExtDM (Zhang et al. 2024) improve efficiency via latent-space denoising or feature extrapolation, but remain unsuitable for real-time or ultra-high-resolution prediction.

End-to-end models. generate the entire sequence in one pass. SimVP (Gao et al. 2022) uses only an MSE loss

to match SOTA at moderate resolutions. Variants such as TAU (Tan et al. 2023) and DMVFN (Hu et al. 2023) incorporate flow consistency or temporal modules, yet still rely on monolithic pixel losses and struggle at 1024^2 .

Training Objectives and Curricula

Most predictors use a *single-stage* reconstruction loss (MSE or L1), often causing blurriness and temporal artifacts. Pre-DRNN++ applies a simple curriculum for stability (Wang et al. 2018), but does not separate coarse reconstruction from edge sharpening or perceptual refinement. In contrast, our three-stage curriculum explicitly stages objectives from pixel reconstruction to gradient-based edge losses to perceptual (VGG) fidelity, yielding sharper and more temporally coherent outputs.

Efficient Spatiotemporal Attention

Vanilla self-attention scales as $\mathcal{O}((ST)^2)$ for $L = ST$ tokens, infeasible for high resolutions. Sparse/local methods (Swin (Liu et al. 2021), Longformer (Beltagy, Peters, and Cohan 2020)) reduce interactions but miss global context; kernel/linear Transformers (Performer (Choromanski et al. 2022), RWKV (Peng et al. 2023), Mamba (Gu and Dao 2024)) achieve $\mathcal{O}(L)$ but still face large L . Factorized spatiotemporal variants (TimeSformer (Bertasius, Wang, and Torresani 2021), ViViT (Arnab et al. 2021)) decouple axes but spatial token counts remain prohibitive at 512^2 – 1024^2 . These limitations motivate our EVA, which natively scales as $\mathcal{O}(S + T)$ while preserving full global context and fine spatial details.

Method

Framework Overview

As illustrated in Fig. 2, RAPTOR adopts a classic Encoder–Translator–Decoder pipeline tailored for real-time, high-resolution video prediction.

Encoder/Decoder. Following the efficient convolutional strategy of SimVP (Gao et al. 2022), the encoder downsamples each frame and maps RGB inputs to a compact latent space. The decoder mirrors this architecture to reconstruct future frames in a single forward pass.

Translator. Our primary contribution lies within the Translator—the Efficient Video Attention (EVA) stack. Given latents $\mathbf{Z} \in \mathbb{R}^{B \times T \times C' \times H' \times W'}$ from the encoder, we adopt a **fine-grained, patch-free** tokenization strategy. Instead of grouping features into coarse patches, which loses spatial resolution, we treat each location on the feature map as an individual token by flattening the spatial and channel dimensions ($S = C' \times H' \times W'$). A linear projection then maps these tokens into a fixed-width embedding $\mathbf{E} \in \mathbb{R}^{B \times T \times S}$. This dense representation, which preserves the original feature map’s granularity, is then processed by the EVA stack.

EVA: Factorized Spatiotemporal Translator

Motivation. Vanilla attention over a flattened sequence of length $L=ST$ incurs $\mathcal{O}(L^2)$ complexity, which is prohibitive for high-resolution inputs. EVA circumvents this

by transposing and alternating the modeling axes. Instead of processing the full ST sequence, we apply a 1D sequence operator along the temporal (T) and spatial (S) dimensions separately. This reduces the dominant sequence length in any single pass to $\max\{T, S\}$, making the architecture highly scalable.

Unified Linear Gated Unit (LGU). A core principle of EVA’s design is simplicity and consistency.

Both TIMEMIX and SPACEMIX are built upon a unified Linear Gated Unit (LGU), inspired by Mamba (Gu and Dao 2024) and RWKV (Peng et al. 2023) architectures. For a given input sequence $\mathbf{U} \in \mathbb{R}^{B \times L \times S}$, the LGU operates as follows:

$$\mathbf{k} = \text{Mish}(\mathbf{U}\mathbf{W}_K), \quad \mathbf{v} = \mathbf{U}\mathbf{W}_V, \quad \mathbf{r} = \sigma(\mathbf{U}\mathbf{W}_R) \quad (1)$$

$$\text{LGU}(\mathbf{U}) = (\mathbf{r} \odot (\mathbf{k} \odot \mathbf{v}))\mathbf{W}_O$$

where $\mathbf{W}_K, \mathbf{W}_V, \mathbf{W}_R, \mathbf{W}_O$ are learnable projections, and \odot is the Hadamard product. This structure performs expressive, input-dependent gating with linear complexity in sequence length L .

EVA Block Structure. Each EVA block sequentially applies the unified LGU along the temporal and spatial dimensions, leveraging a pre-LayerNorm configuration and residual connections to enhance training stability.

First, the temporal mixing operation (TIMEMIX) is applied. Given input tensor $\mathbf{E} \in \mathbb{R}^{B \times T \times S}$, we treat each batch independently as a 1D sequence of length $L_T = T$. To ensure causality, we apply a parameter-free temporal shift (implemented via zero-padding and rolling by one step). This shifted sequence is then processed by the LGU along the temporal axis.

Following this, the spatial mixing operation (SPACEMIX) is performed. Starting from the output $\mathbf{E}' \in \mathbb{R}^{B \times T \times S}$ of TIMEMIX, we transpose the tensor to $\mathbf{E}'^T \in \mathbb{R}^{B \times S \times T}$, interpreting each batch again as a 1D sequence—this time of length $L_S = S$. The LGU is reused here to mix information across spatial locations. This spatial mixing is fully parallel and position-wise, enabling efficient interaction across spatial locations without introducing directional bias or path dependencies.

By stacking N such EVA blocks, the model progressively builds a global spatiotemporal receptive field, while maintaining per-block computational complexity that scales linearly with $\max\{T, S\}$.

Complexity and Edge-Cost Model

We evaluate per-layer inference cost in a UAV setting with $T = 10$ frames at 1024×1024 resolution. The encoder applies $8\times$ downsampling over three stages, yielding $S = 16384$ tokens per frame and $L = ST = 163840$ total tokens. The latent width is $D = 512$. On Jetson AGX Orin (10.65 TFLOPS FP16, 204.8 GB/s memory bandwidth), we estimate a lower bound on latency as $\text{lat}_{\text{theo}} = \max\{\text{FLOPs}/(10.65 \text{ TFLOPS}), \text{Bytes}/(204.8 \text{ GB/s})\}$.

We compare three architectures per translator layer. ViT incurs $2L^2D + 4LD^2$ FLOPs and $2L^2$ bytes for attention




Figure 2: Overview of RAPTOR architecture. The framework follows an Encoder–Translator–Decoder design. The EVA Translator employs a patch-free design, operating on full feature maps from the encoder and stacking N EVA blocks. Each block alternates a temporal TIMEMIX and a spatial SPACEMIX built from specialized, asymmetric operators, achieving global spatiotemporal receptive fields with linear-in-length complexity and enabling real-time high-resolution prediction.

logits. Linear attention (RWKV/Mamba) requires $8LD^2$ FLOPs and $\mathcal{O}(LD)$ memory. EVA reduces cost to $4(T + S)D^2$ FLOPs with $\mathcal{O}(D \max\{T, S\})$ memory, by decoupling temporal and spatial modeling.

As shown in Table 1, EVA achieves 1.72×10^{10} FLOPs and ~ 1.0 GB memory per layer, with a theoretical latency of 5.1 ms—over $1600\times$ faster than ViT and $20\times$ faster than linear baselines. While others become compute-bound, EVA is memory-bound and remains in the real-time regime, making it well-suited for UAV deployment.

Method	FLOPs	Bytes	Latency (ms)
ViT (quadratic)	2.75×10^{13}	~ 54 GB	~ 2600
Linear (RWKV)	3.44×10^{11}	~ 2.4 GB	~ 32
Linear (Mamba)	3.92×10^{11}	~ 1.8 GB	~ 37
EVA (ours)	1.72×10^{10}	~ 1.0 GB	~ 5.1

Table 1. Per-layer inference cost under an $8\times$ downsampling setting ($T=10$, $S=16384$, $D=512$) on Jetson AGX Orin. Latency is a theoretical lower bound. “Bytes/layer” refers to the estimated total memory footprint.

Three-Stage Curriculum Learning

Pixel-wise objectives alone tend to blur edges and ignore temporal coherence, whereas optimizing all advanced objectives from scratch often destabilizes training. We therefore adopt a **coarse-to-fine** curriculum.

Stage I: Foundational Reconstruction. We first fit coarse structures using the ℓ_1 loss:

$$\mathcal{L}_{S1} = \|\hat{\mathbf{X}} - \mathbf{X}\|_1. \quad (2)$$

Stage II: Edge and Temporal Consistency. We add a Gradient Difference Loss (GDL) and a frame-to-frame temporal smoothness term:

$$\mathcal{L}_{\text{gdl}} = \|\nabla_x(\hat{\mathbf{X}}) - \nabla_x(\mathbf{X})\|_1 + \|\nabla_y(\hat{\mathbf{X}}) - \nabla_y(\mathbf{X})\|_1 \quad (3)$$

$$\mathcal{L}_{\text{smooth}} = \frac{1}{T'-1} \sum_{t=1}^{T'-1} \|\hat{\mathbf{x}}_{t+1} - \hat{\mathbf{x}}_t\|_1. \quad (4)$$

$$\mathcal{L}_{S2} = \mathcal{L}_{S1} + \lambda_{\text{gdl}} \mathcal{L}_{\text{gdl}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}}. \quad (5)$$

Here ∇_x, ∇_y are first-order difference operators.

Stage III: Perceptual Refinement. Finally we add a perceptual loss in a pretrained VGG feature space to recover high-frequency details:

$$\mathcal{L}_{S3} = \mathcal{L}_{S2} + \lambda_{\text{perc}} \frac{1}{T'} \sum_{t=1}^{T'} \frac{1}{S} \|\phi(\hat{\mathbf{x}}_t) - \phi(\mathbf{x}_t)\|_2^2, \quad (6)$$

with $\phi(\cdot)$ the VGG feature extractor. We inherit weights across stages to stabilize optimization.

Implementation Notes. Our architecture incorporates several practical design choices for stability and efficiency. We use Pre-LayerNorm and residual connections in all EVA blocks. Causality in the TIMEMIX module is strictly enforced via a parameter-free time-shift operation on the input tensor before projection. The activation function in our SPACEMIX LGU is Mish, which we found to perform slightly better than GELU or ReLU in preliminary tests. To manage memory for high-resolution training, we employ mixed-precision (AMP) and gradient checkpointing. These engineering details are crucial for achieving the real-time performance on edge hardware reported in our experiments.

Experiments

We evaluate RAPTOR across four aspects: (i) comparative performance against strong state-of-the-art baselines, (ii) ablations isolating the contributions of EVA and the three-stage

Method	UAvid-128 ²			UAvid-512 ²				UAvid-1024 ²			
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	FPS↑	PSNR↑	SSIM↑	LPIPS↓	FPS↑
PredRNN	24.1	0.714	0.189	OOM				OOM			
MIM	24.6	0.746	0.174	OOM				OOM			
MotionRNN	26.6	0.757	0.177	OOM				OOM			
PredRNN-V2	25.9	0.761	0.141	OOM				OOM			
MCVD	28.7	0.788	0.115	OOM				OOM			
ExtDM	29.8	0.797	0.096	OOM				OOM			
SimVP	28.9	0.791	0.120	24.5	0.578	0.120	31.9	OOM			
TAU	29.4	0.805	0.109	19.8	0.431	0.499	31.7	OOM			
RAPTOR (ours)	30.7	0.820	0.076	28.4	0.784	0.095	145.6	19.3	0.455	0.582	59.6

Table 2. UAvid 10→10 at 128², 512², and 1024². FPS (batch=1, RTX 6000 Ada) is only reported for 512² and 1024² resolutions. *OOM* (Out Of Memory) indicates peak training/inference memory >48 GB. Best results in **bold**.

Method	PSNR↑	SSIM↑	LPIPS↓
PredRNN	27.6	0.839	0.204
MIM	27.9	0.843	0.177
MotionRNN	28.3	0.851	0.169
PredRNN-V2	28.4	0.838	0.139
MCVD	27.5	0.846	0.104
ExtDM	27.9	0.799	0.093
SimVP	33.7	0.905	0.078
TAU	30.6	0.893	0.086
RAPTOR (ours)	32.0	0.912	0.062

Table 3. KTH 10→30 at 64². Best results in **bold**.

curriculum, (iii) scalability to ultra-high resolutions where many methods fail, and (iv) downstream impact on a real-world UAV navigation task.

Experimental Setup

Datasets. Our evaluation is conducted on two standard benchmarks and a custom real-world dataset. **KTH** (Schuldt, Laptev, and Caputo 2004) is a classic low-resolution (64²) dataset for human action prediction. **UAvid** (Lyu et al. 2020) is a challenging high-resolution dataset from UAV footage, on which we perform extensive evaluation at 128², 512², and a frontier 1024² resolution to test the limits of scalability. A custom **Real-World Navigation** dataset was collected for downstream task validation.

Baselines. To ensure a thorough comparison, we select a strong and diverse set of state-of-the-art (SOTA) methods covering three dominant paradigms: **Autoregressive models** (PredRNN (Wang et al. 2017), MIM (Wang et al. 2019), MotionRNN (Wu et al. 2021), and PredRNN-V2 (Wang et al. 2023b)), which represent classic recurrent approaches. **Diffusion models** (MCVD (Voleti, Jolicoeur-Martineau, and Pal 2022) and ExtDM (Zhang et al. 2024)), known for their high generative fidelity. **End-to-end models** (SimVP (Gao et al. 2022) and TAU (Tan et al. 2023)), which are the most relevant competitors in terms of efficiency.

Metrics and Implementation Details. We evaluate prediction quality using PSNR, SSIM (Wang et al. 2004), and LPIPS (Zhang et al. 2018). Inference speed is measured in Frames Per Second (FPS). For the downstream task, we report mission Success Rate (SR, %). Our models were trained using PyTorch on 4* NVIDIA 6000 Ada GPUs. We leverage standard acceleration techniques including DDP, AMP, and gradient checkpointing. Further implementation details, including hyperparameters and specifics of the navigation task setup, are provided in the supplementary material.

Quantitative Comparison with State-of-the-Art

We present a comprehensive comparison against SOTA methods in Table 2 and Table 3. The results reveal a clear narrative about RAPTOR’s capabilities.

On standard low-resolution benchmarks like KTH, RAPTOR demonstrates highly competitive performance. While SimVP achieves a higher PSNR, our model obtains the best SSIM (0.912) and LPIPS (0.062), indicating superior structural and perceptual quality, which are often more critical for downstream tasks than raw pixel-level accuracy.

The advantage of RAPTOR becomes decisive on high-resolution UAV data. At 128², RAPTOR establishes a new SOTA across all metrics, significantly outperforming all baselines. As we scale to the challenging 512² setting, this gap widens dramatically. Most prior methods fail due to Out-of-Memory (OOM) errors, highlighting a fundamental architectural bottleneck. Among the methods that can run, RAPTOR surpasses the strong SimVP baseline by a large margin (+3.9 dB PSNR, +0.206 SSIM).

Crucially, RAPTOR is the only model capable of operating at 1024², establishing a new frontier for real-time video prediction. While performance metrics naturally decrease at this extreme resolution, RAPTOR’s ability to process such data at interactive frame rates (59.6 FPS) is a breakthrough that no other compared method achieves. This unique scalability is a direct result of our **patch-free**, linear-complexity EVA design.

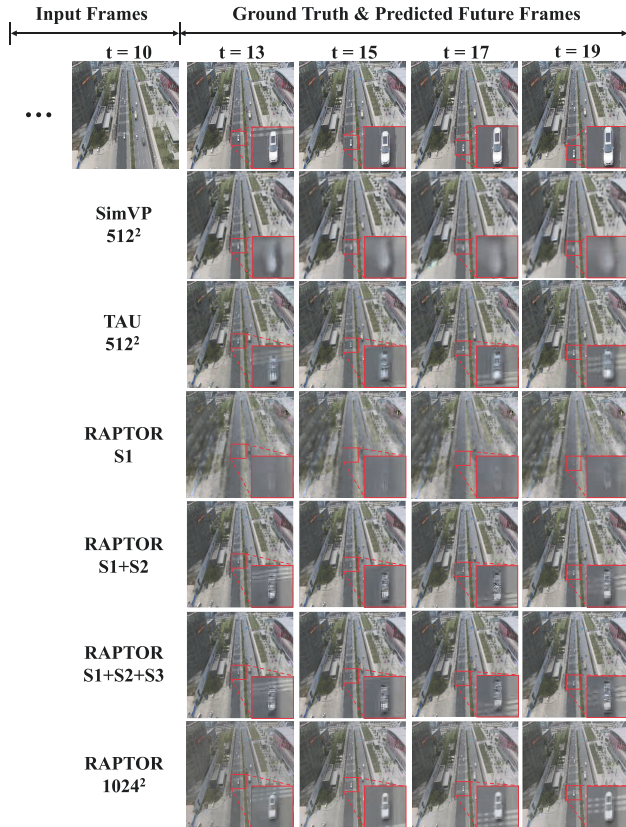


Figure 3: Qualitative comparison on a challenging UAVID scene. The magnified insets highlight how RAPTOR, particularly with the full S1+S2+S3 curriculum, generates sharper and more coherent predictions of moving vehicles compared to baselines. RAPTOR is also the only model capable of producing a prediction at 1024^2 resolution.

Ablation Studies

We conducted ablation studies to validate the individual contributions of RAPTOR’s core components on the UAVID-512² dataset.

Translator Variant	PSNR↑	SSIM↑	LPIPS↓
ViT-based	<i>Out Of Memory</i>		
Only TimeMix	21.2	0.555	0.198
Only SpaceMix	27.1	0.713	0.115
RAPTOR (full EVA)	28.4	0.784	0.095

Table 4. Ablation of the EVA components on UAVID-512².

Efficacy of the EVA Translator. To prove the necessity of our factorized design, we evaluate three variants of our translator: the full EVA module, and versions using only the TIMEMIX or SPACEMIX components. As shown in Table 4, using only TIMEMIX fails to capture spatial details, resulting in poor performance. Using only SPACEMIX performs better but still lags significantly behind the full model. This

Training Stage	PSNR↑	SSIM↑	LPIPS↓
RAPTOR (S1)	25.31	0.603	0.121
RAPTOR (S1+S2)	28.65	0.798	0.102
RAPTOR (S1+S2+S3)	28.40	0.784	0.095

Table 5. Ablation of the Three-Stage Curriculum on UAVID-512², highlighting the progressive benefits of each stage.

confirms that explicitly modeling both dimensions and, crucially, alternating between them, is essential for building a powerful spatiotemporal representation. We also note that a standard ViT-based translator runs Out-of-Memory, highlighting the efficiency of our approach.

Value of the Three-Stage Curriculum. The curriculum’s effectiveness is demonstrated in Table 5 and visualized in Figure 3. The S1 model, trained only on pixel loss, produces blurry results with low SSIM/LPIPS scores. Introducing edge and temporal losses in S2 dramatically improves structural metrics (PSNR and SSIM). The final S3, with the perceptual loss, achieves the best visual quality, as evidenced by the lowest LPIPS score. Interestingly, S3 shows a marginal drop in PSNR/SSIM compared to S2. This reflects a well-known trade-off: optimizing for perceptual realism (LPIPS) can sometimes slightly reduce pixel-level accuracy. The visual results in Figure 3 confirm that the S3 model is perceptually superior, justifying the full curriculum for generating visually plausible predictions.

High-Resolution Prediction and Evaluation

The true measure of our work lies in its ability to tackle previously intractable problems and deliver real-world value. We test RAPTOR on this frontier, first through qualitative analysis to form a hypothesis, and then through a real-world downstream task to validate it.

Qualitative Analysis. Quantitative metrics, while essential, do not capture the full picture of predictive quality. As shown in Figure 3, for instance, while SimVP often scores higher than TAU on PSNR, TAU’s predictions can appear more temporally consistent, highlighting the limitations of purely pixel-based evaluation.

Our qualitative results for RAPTOR tell a compelling story that aligns with our curriculum ablation study. The S1 model produces blurry outputs. The S1+S2 model, benefiting from edge and temporal losses, already generates remarkably sharp structures, especially visible in the distinct outlines of the moving vehicles. The final S1+S2+S3 model, refined by the perceptual loss, further enhances realism and textural details, resulting in the most visually plausible predictions. This progression validates our curriculum’s design: S2 significantly improves structural integrity (boosting SSIM), while S3 enhances perceptual fidelity (boosting LPIPS), as confirmed by the data in Table 5.

Most critically, the last row of Figure 3 showcases RAPTOR’s breakthrough scalability. While all other methods fail at 1024^2 , RAPTOR maintains the scene’s key structural and dynamic information. This leads to an insight: although our

1024² predictions have lower raw metric scores than their 512² counterparts (Table 2), they seem to preserve the fine-grained details of small, dynamic objects that are vital for navigation. This suggests a critical hypothesis: for an embodied agent, the practical value of high-resolution prediction may far outweigh its raw pixel-metric scores.

Method	PSNR↑	SSIM↑	LPIPS↓	FPS↑	SR↑
GPT-4o w/o predict	-	-	-	-	33%
SimVP 512 ²	24.1	0.54	0.22	5.6	+1%
TAU 512 ²	18.2	0.42	0.55	4.9	+4%
RAPTOR 512²	27.5	0.74	0.11	30.2	+7%
RAPTOR 1024²	16.3	0.49	0.65	18.6	+18%

Table 6. Downstream navigation performance comparison between different types of video processing methods. SR denotes success rate (absolute improvement over the 33% baseline). The *GPT-4o w/o predict* baseline represents a reactive planner where the Vision-Language Model, GPT-4o, makes decisions based only on the current frame, without any future prediction.

Real-World Navigation Task. To rigorously test this hypothesis, we integrated RAPTOR into a real UAV’s control loop for a challenging object-goal navigation task. The experiment was conducted on a university campus during nighttime to pose significant challenges for visual perception and prediction. The UAV was tasked with navigating towards periodically moving targets, including pedestrians and patrol vehicles. Due to the dynamic nature of these targets, which makes path-based metrics unreliable, we use Success Rate (SR) as the primary performance indicator, evaluated over 100 trials for each method.

The results, shown in Table 6, are compelling. While baselines offer minimal improvement, RAPTOR provides a substantial boost. As hypothesized, by leveraging native 1024² predictions, RAPTOR achieves a remarkable **+18%** absolute improvement in success rate. This is because the high-resolution foresight allows the planner to perceive and react to small, distant hazards that are simply lost in down-sampled inputs. This finding directly translates our model’s superior predictive fidelity into tangible gains in robotic safety. Figure 4 further provides a qualitative visualization of the real-world comparison, showing RAPTOR can predict the moving pedestrian with a clarity at both 512² and 1024² resolutions that the baselines fail to achieve.

Conclusion

In this work, we address the challenge of real-time, high-resolution video prediction for autonomous UAVs. We introduce RAPTOR, an end-to-end, single-pass framework designed to overcome the latency and scalability limitations of prior iterative methods. RAPTOR’s efficacy stems from two core innovations. The first is our Efficient Video Attention (EVA) module, which tackles the quadratic complexity and memory bottleneck of processing long spatiotemporal sequences where the effective length is $L = ST$. Through

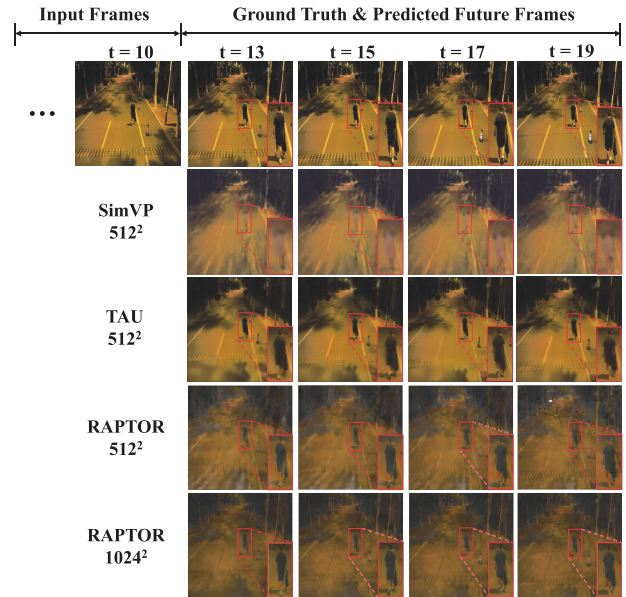


Figure 4: Qualitative comparison from the real-world UAV navigation task. Magnified insets highlight that RAPTOR (with full Three-Stage Curriculum) outperforms the baselines by generating sharp and coherent predictions of the moving pedestrian, even in a nighttime scenario.

spatiotemporal factorization and feature transposition, EVA applies linear gated units alternately along the spatial (S) and temporal (T) axes, which slashes the computational complexity from $O((ST)^2)$ to $O(S + T)$ and reduces the peak memory complexity to $O(\max(S, T))$. This fundamental reduction is what makes high-resolution processing tractable. Furthermore, its patch-free nature preserves fine-grained spatial details, while the alternating process ensures robust spatiotemporal interaction to maintain high modeling quality. The second is a three-stage curriculum learning strategy that progressively enhances the prediction’s perceptual fidelity. Our extensive experiments demonstrate that RAPTOR is the first model to achieve real-time performance at 1024² resolution. This capability directly translates into a significant 18% improvement in success rate on a real-world UAV navigation task. By breaking the long-standing trade-off between resolution, speed, and quality, RAPTOR represents a significant step towards creating safer, more intelligent, and truly anticipatory aerial embodied agents.

Acknowledgments

This work was supported by the Key Program of Chinese Academy of Sciences under Grant RCJJ-145-24-13 and KGFZD-145-25-38, and supported by the Science and Disruptive Technology Program under Grant AIRCAS2024-AIRCAS-SDTP-03.

References

Arnab, A.; Dehghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. Vivit: A video vision transformer. In

- 2021 *IEEE/CVF international conference on computer vision (ICCV)*, 6816–6826.
- Beltagy, I.; Peters, M. E.; and Cohan, A. 2020. Longformer: The Long-Document Transformer. arXiv:2004.05150.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is Space-Time Attention All You Need for Video Understanding? In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 813–824.
- Chromanski, K.; Likhoshesterov, V.; Dohan, D.; Song, X.; Gane, A.; Sarlos, T.; Hawkins, P.; Davis, J.; Mohiuddin, A.; Kaiser, L.; Belanger, D.; Colwell, L.; and Weller, A. 2022. Rethinking attention with performers. arXiv:2009.14794.
- Gao, Z.; Tan, C.; Wu, L.; and Li, S. Z. 2022. SimVP: Simpler Yet Better Video Prediction. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3160–3170.
- Gu, A.; and Dao, T. 2024. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. arXiv:2312.00752.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, 6840–6851.
- Hu, X.; Huang, Z.; Huang, A.; Xu, J.; and Zhou, S. 2023. A dynamic multi-scale voxel flow network for video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6121–6131.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9992–10002.
- Lyu, Y.; Vosselman, G.; Xia, G.-S.; Yilmaz, A.; and Yang, M. Y. 2020. UAVid: A semantic segmentation dataset for UAV imagery. *ISPRS journal of photogrammetry and remote sensing*, 165: 108–119.
- Peng, B.; Alcaide, E.; Anthony, Q.; Albalak, A.; Arcadinho, S.; Cao, H.; Cheng, X.; Chung, M.; Grella, M.; GV, K. K.; He, X.; Hou, H.; Kazienko, P.; Kocon, J.; Kong, J.; Koptyra, B.; Lau, H.; Mantri, K. S. I.; Mom, F.; Saito, A.; Tang, X.; Wang, B.; Wind, J. S.; Wozniak, S.; Zhang, R.; Zhang, Z.; Zhao, Q.; Zhang, P.; Zhu, J.; and Zhu, R.-J. 2023. RWKV: Reinventing RNNs for the Transformer Era. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 14048–14077.
- Schuldt, C.; Laptev, I.; and Caputo, B. 2004. Recognizing human actions: a local SVM approach. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 3, 32–36.
- Shi, X.; Chen, Z.; Wang, H.; Yeung, D.; Wong, W.; and Woo, W. 2015. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. In *Advances in Neural Information Processing Systems*, volume 28, 802–810.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *9th International Conference on Learning Representations (ICLR)*.
- Tan, C.; Gao, Z.; Wu, L.; Xu, Y.; Xia, J.; Li, S.; and Li, S. Z. 2023. Temporal Attention Unit: Towards Efficient Spatiotemporal Predictive Learning. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18770–18782.
- Voleti, V.; Jolicoeur-Martineau, A.; and Pal, C. 2022. MCVD - Masked Conditional Video Diffusion for Prediction, Generation, and Interpolation. In *Advances in Neural Information Processing Systems*, volume 35, 23371–23385.
- Wang, Y.; Gao, Z.; Long, M.; Wang, J.; and Yu, P. S. 2018. PredRNN++: Towards A Resolution of the Deep-in-Time Dilemma in Spatiotemporal Predictive Learning. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 5123–5132.
- Wang, Y.; Long, M.; Wang, J.; Gao, Z.; and Yu, P. S. 2017. PredRNN: Recurrent Neural Networks for Predictive Learning using Spatiotemporal LSTMs. In *Advances in Neural Information Processing Systems*, volume 30.
- Wang, Y.; Wu, H.; Zhang, J.; Gao, Z.; Wang, J.; Yu, P. S.; and Long, M. 2023a. PredRNN: A Recurrent Neural Network for Spatiotemporal Predictive Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2208–2225.
- Wang, Y.; Wu, H.; Zhang, J.; Gao, Z.; Wang, J.; Yu, P. S.; and Long, M. 2023b. PredRNN: A Recurrent Neural Network for Spatiotemporal Predictive Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2208–2225.
- Wang, Y.; Zhang, J.; Zhu, H.; Long, M.; Wang, J.; and Yu, P. S. 2019. Memory In Memory: A Predictive Neural Network for Learning Higher-Order Non-Stationarity from Spatiotemporal Dynamics. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9146–9154.
- Wang, Z.; Bovik, A.; Sheikh, H.; and Simoncelli, E. 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612.
- Wu, H.; Yao, Z.; Wang, J.; and Long, M. 2021. MotionRNN: A Flexible Model for Video Prediction with Spacetime-Varying Motions. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15430–15439.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 586–595.
- Zhang, Z.; Hu, J.; Cheng, W.; Paudel, D. P.; and Yang, J. 2024. ExtDM: Distribution Extrapolation Diffusion Model for Video Prediction. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19310–19320.