VideoMaMa++: Temporally Consistent Video Matting via Preserve-and-Refine Tokens

Sangbeom Lim[1] **Seoung Wug Oh**[2] **Heeji Yoon**[3]
Seungryong Kim[3†] **Joon-Young Lee**[2†]

[1]Korea University [2]Adobe Research [3]KAIST AI

https://cvlab-kaist.github.io/VideoMaMa-PlusPlus

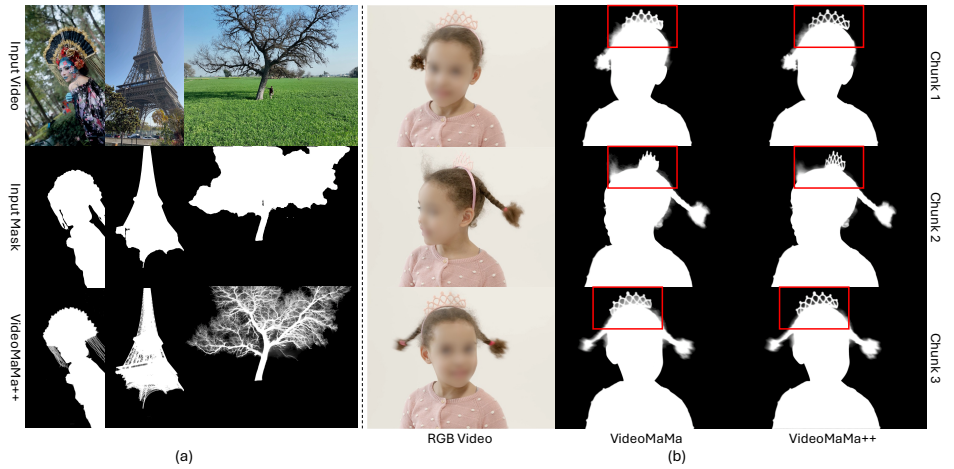


Figure 1: We introduce **VideoMaMa++**, a mask-guided video matting framework that recovers fine-grained details (e.g., hair strands) and scales reliably to long videos. (a) Per-frame qualitative results on diverse in-the-wild videos. (b) Long-video matting results: VideoMaMa [27], which processes videos in chunks, fails to maintain consistency across chunk boundaries (red boxes), whereas VideoMaMa++ produces mattes that are both finely detailed and temporally consistent.

Abstract

Diffusion-based video matting has recently emerged as a leading approach, leveraging video diffusion priors from internet-scale pretraining to achieve strong zero-shot generalization from limited synthetic data. However, the heavy compute of these diffusion backbones limits inference to a short chunk of just 6–24 frames, so even ordinary videos span multiple chunks. Because each chunk is matted independently from coarse mask guidance, existing methods fail to maintain consistency across chunk boundaries, producing flickering edges and discontinuous transparency. We address this with a mask-guided diffusion-matting framework built on two ideas. The RGB and mask conditioning are represented as separate token streams that interact through full spatio-temporal self-attention, letting the model adaptively combine fine RGB texture with coarse mask layout rather than committing to a fixed channel-wise fusion. On top of this, we propose a novel pair of learnable embeddings, the *preserve token* and the *refine token*, that act as a per-frame conditioning interface and enable a sliding-window inference scheme in which already-generated mattes are propagated across chunk boundaries for local tem-

†Co-corresponding authors.

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

poral consistency, complemented by a globally shared *reference anchor* at every chunk to prevent error accumulation in long sequences without any additional training. Across standard video matting benchmarks and in-the-wild long footage, our framework outperforms prior diffusion-based methods in both fine-detail accuracy and temporal consistency, and uniquely maintains consistency across chunk boundaries on long videos.

1 Introduction

Video matting, the task of extracting foreground objects with pixel-level precision from video, is a fundamental component of video editing applications such as background replacement [57, 12, 37], visual composition [10, 52], and relighting [4, 37, 47, 8]. Producing convincing composites in these settings requires alpha mattes that capture fine structures such as hair strands and motion-blurred edges, and that remain temporally consistent across the entire video.

Recent video matting methods leverage priors from video object segmentation models [19, 54]. However, these methods still rely on dedicated matting annotations and offer limited zero-shot generalization, with most existing methods restricted to human portraits due to the scarcity of matting labels for other categories. A more recent line of work instead leverages video diffusion priors [27, 13], transferring generative knowledge obtained from internet-scale pretraining to close the train–test gap that arises when models are trained only on synthetic matting data. In particular, VideoMaMa [27] demonstrates that a mask-to-matte diffusion model trained on synthetic data alone can generalize to non-human categories and to in-the-wild video footage.

However, applying diffusion-based matting to real videos exposes a fundamental limitation that hinders general video processing. Diffusion backbones are computationally heavy and cannot process arbitrarily long sequences in a single forward pass; existing methods are therefore chunk-based, handling only 6–24 frames at a time. As a result, even ordinary videos must be split into multiple chunks and processed sequentially. When each chunk is matted independently from its coarse mask guidance, the resulting mattes drift across chunk boundaries, producing flickering edges and discontinuous transparency, as illustrated in Figure 1 (b). The root cause is the inherent ambiguity of coarse mask guidance: a single binary mask admits many plausible fine-grained refinements at object boundaries, and independently processed chunks settle on different ones.

In this work, we present a novel mask-guided diffusion-matting framework designed for long, high-resolution video. At its core, we represent the two guidance modalities, the RGB frames and the coarse mask, as separate token streams alongside the noisy matte tokens, and let them interact through full spatio-temporal self-attention. This design replaces the channel concatenation used in prior U-Net-based diffusion matting [27] and preserves each modality's distinct structure. The high-frequency RGB texture carries hair and edge cues, while the binary mask provides a coarse but spatially anchored foreground extent, and attention adaptively composes the two at every layer and spatial location. As a result, the model leans on RGB at object boundaries where the mask is unreliable, and on the mask in the interior where it confidently localizes the foreground. We instantiate this on a Diffusion Transformer (DiT) [35, 46] with Rotary Position Embedding (RoPE) [43], which naturally accommodates diverse resolutions and frame counts.

To resolve the chunk-boundary inconsistency described above, we introduce a novel pair of learnable token embeddings, the *preserve token* and the *refine token*, a new conditioning interface for consistent video matting. The preserve token marks a frame whose guidance is already a fine matte, previously generated or ground truth, and tells the model to keep it as a reference. The refine token marks a frame with only a coarse mask and asks the model to refine it into a fine matte consistent with the preserve-tagged references via DiT's full self-attention. We employ these tokens in a sliding-window inference scheme in which each chunk overlaps a few preserve-tagged frames with the previous chunk, so that already-generated mattes propagate across chunk boundaries and maintain local temporal consistency. To further guard against error accumulation over long sequences, we additionally supply a globally reliable matte as a *reference anchor* at every chunk through the same preserve token, without any additional training.

We validate our framework through extensive experiments on standard video matting benchmarks and in-the-wild footage. Our method outperforms existing video matting approaches on VideoMatte [28], YoutubeMatte [54], and V-HIM60 [19] in both fine-detail accuracy and temporal consistency, and

achieves state-of-the-art performance on long-video matting where prior diffusion-based methods fail to maintain consistency across chunk boundaries. The same preserve mechanism also gives users direct control at inference time, for example to interactively correct erroneous outputs by re-tagging specific frames as refine while keeping trusted neighbors as preserve.

Our contributions can be summarized as follows.

- A mask-guided diffusion-matting framework with *token-level guidance representation*, where the RGB frames and the coarse mask are kept as separate token streams adaptively composed by full spatio-temporal self-attention, rather than fused via channel concatenation as in prior works.
- A novel *preserve* and *refine* token mechanism that enables chunk-consistent matting through sliding-window inference with overlap-based propagation, complemented by a globally shared *reference anchor* that prevents error accumulation in long sequences without any additional training.
- State-of-the-art performance on standard video matting benchmarks and in-the-wild long-video footage, with substantial gains in fine-detail accuracy and temporal consistency over prior diffusion-based methods that process video chunk-wise.

2 Related Work

Video matting. Video matting methods can be broadly categorized into auxiliary-free and auxiliary-guided approaches [34, 16]. Auxiliary-free methods [29, 21, 23] typically focus on portrait matting, restricting their applicability to human-centric scenarios. Auxiliary-guided approaches relax this restriction by accepting additional conditioning at the cost of an extra annotation. Trimap-guided methods [17, 41] require manual trimap annotations and cannot easily leverage modern segmentation models such as the SAM family [22, 36, 3]. More recent mask-guided methods reduce this burden by accepting only a coarse binary mask. MaGGIe [19] decouples target tracking from matting but struggles to generalize to in-the-wild videos, while MatAnyone [54, 53] propagates target-assigned mattes through a memory-augmented architecture trained on dedicated matting datasets such as VM800 [54] and VMReal [53]. Despite these advances, all of these methods remain constrained by predominantly synthetic and human-centric training data, leaving in-the-wild generalization to non-human categories largely unsolved.

Diffusion priors for perception. Diffusion priors from internet-scale pretraining have been transferred to dense perception tasks. Marigold [20] pioneered this direction by adapting Stable Diffusion [39] for depth estimation, and the same recipe extends to a range of video perception tasks [1, 42, 15, 51], with more recent works [50, 58] extending it to DiT-based backbones.

Diffusion long-video inference. Diffusion backbones, however, are computationally heavy and limited to short chunks, so applying them to long videos requires a chunk-wise inference scheme. ChronoDepth [42] and DepthCrafter [15] adopt sliding-window inference with overlapping frames between adjacent chunks for local temporal smoothness, but offer no global anchor against drift. A parallel line of work tackles this through auto-regressive generation [18, 5, 56, 31, 9, 55, 33], where each chunk is conditioned on the previous chunk via causal attention with KV-cache reuse. While effective for free-form video generation, this hidden-state propagation does not naturally compose with explicit external conditioning such as a guide mask. Our work instead propagates fully-formed reference mattes between chunks through the preserve token, keeping each chunk fully conditioned on its own user-supplied mask.

Diffusion-based video matting. In the matting domain, recent diffusion-based methods transfer video diffusion priors to video matting using only synthetic data. The auxiliary-free GVM [13] and the mask-guided VideoMaMa [27] both demonstrate that this prior generalizes to non-human categories and in-the-wild footage, but both inherit the chunk-wise inference limitation of prior diffusion-based perception models. VideoMaMa, the closest mask-guided baseline, additionally fuses RGB, mask, and noise via channel concatenation, entangling the guidance modalities and limiting fine-detail recovery. Our work departs from this design by representing the two guidance modalities as separate token streams, and by introducing the preserve and refine tokens that maintain chunk-wise consistency through globally shared reference anchors and overlap-based local smoothing.

3 Method

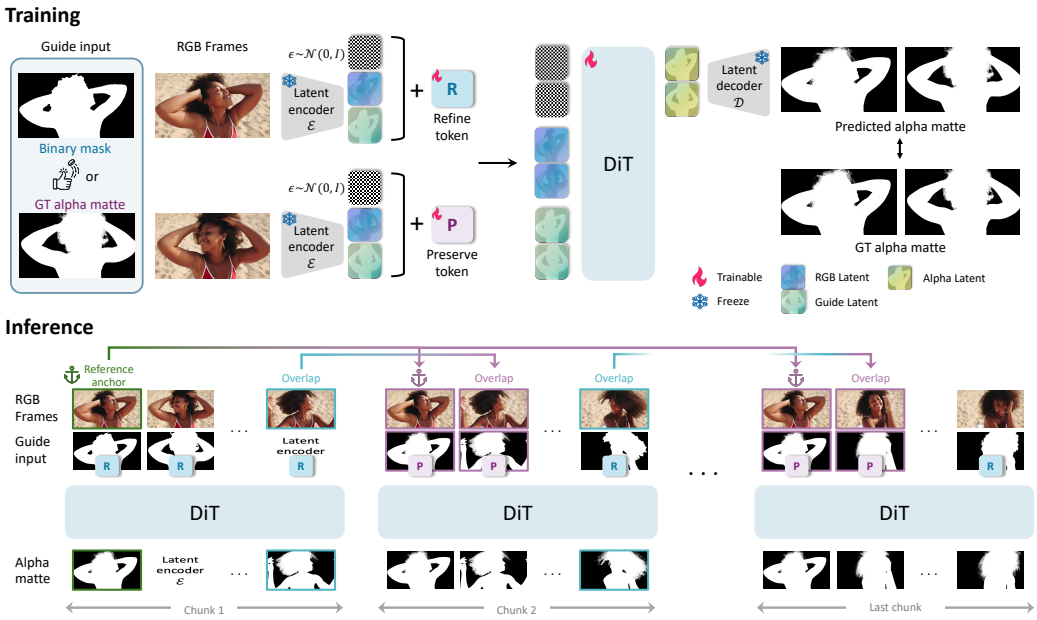


Figure 2: **Overview of the VideoMaMa++ training and inference. (Top)** During training, RGB frames are paired with either a binary mask or the ground-truth alpha matte as guide input, encoded into separate token streams by a frozen VAE encoder $\mathcal{E}$, and tagged with a refine token (**R**) or preserve token (**P**). The DiT processes both modes jointly within a single clip, supervised against the ground-truth matte. **(Bottom)** At inference, long videos are processed in a sliding-window manner. Each chunk receives the refine token on its non-overlapping frames and the preserve token on its overlap frames, which carry the previous chunk's predicted matte. A globally selected reference anchor frame is additionally supplied as preserve input to every chunk, providing a stable anchor that prevents error accumulation across the sequence.

3.1 Problem Formulation and Overview

Video matting is mathematically defined by the alpha compositing equation,

$$I_t = \alpha_t F_t + (1 - \alpha_t) B_t, \quad t = 1, \ldots, T, \tag{1}$$

where t denotes the frame index, T is the total number of frames in the video, I_t is the observed image at frame t, F_t is the foreground, B_t is the background, and $\alpha_t \in [0, 1]$ is the alpha matte representing pixel-level opacity. Unlike binary segmentation masks $M \in \{0, 1\}$, alpha mattes capture fine-grained details such as hair strands, motion blur, and semi-transparent regions, which are critical for realistic video compositing.

We adopt a mask-guided formulation that decomposes video matting into two subproblems: foreground tracking and boundary refinement. Tracking a target object across a video is itself a demanding task, and entangling it with matte generation forces a single model to simultaneously localize objects and resolve pixel-level matting. By conditioning on a binary mask, we decouple these two challenges and allow our model to focus on generating fine-grained matting details such as hair strands and motion blur, while inheriting localization capabilities from off-the-shelf segmentation models [36, 3, 7, 6].

We design an end-to-end pipeline for long-video matting that takes a video and a single user prompt on the first frame as input, and produces a temporally consistent alpha matte across the entire sequence. The pipeline operates in three stages, illustrated in Figure 2. First, we obtain a per-frame foreground signal by propagating the user prompt across all frames using an off-the-shelf video segmentation model such as Segment Anything [36, 3]. Second, the resulting binary mask track is passed to our Diffusion Transformer (DiT) backbone together with the RGB frames to generate fine-grained alpha

mattes (Section 3.2). Finally, long sequences are handled by a chunk-wise inference scheme equipped with preserve and refine tokens (Section 3.3), which propagates information across chunk boundaries to maintain temporal consistency throughout the video.

3.2 Repurposing Diffusion Transformer for Video Matting

Video matting annotations are scarce and limited in quantity and category coverage. Generalizing to real-world footage therefore requires a strong prior over natural video content. We build our framework on a video DiT [35, 46] backbone, which provides a rich generative prior learned from internet-scale video pretraining and offers several architectural properties well-suited to video matting.

The input video V, the guide mask track M, and, during training, the ground-truth alpha matte α are each encoded into the latent space by a shared, frozen VAE encoder $\mathcal{E}$, producing z_V, z_M, and z_α respectively. We sample Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$ with the same shape as z_α, and concatenate z_V, z_M, and ϵ along the token dimension as input to the DiT backbone:

$$\hat{z}_\alpha = \mathcal{F}_{\text{DiT}}\big(\text{concat}(z_V,\, z_M,\, \epsilon)\big), \tag{2}$$

where $\mathcal{F}_{\text{DiT}}(\cdot)$ denotes our adapted DiT and $\text{concat}(\cdot)$ denotes concatenation. Unlike UNet-based matting models [27, 2] that fuse the three modalities along the channel dimension, our token-level concatenation lets each modality occupy its own token stream and preserves modality-specific information throughout the network. The final alpha matte video is obtained by decoding the predicted latents through the VAE decoder, $\hat{\alpha} = \mathcal{D}(\hat{z}_\alpha)$. We train the model with the flow-matching [30] objective of the pretrained DiT [46]. Although training follows the standard multi-step schedule, we find a single denoising step at inference sufficient to produce high-quality mattes, and adopt it as our default for efficiency.

3.3 Preserve and Refine Tokens

The formulation above operates on a single clip, while video matting in practice requires handling long sequences. DiT backbones are computationally heavy and cannot process long videos in a single forward pass. A natural solution is chunk-wise inference, where the video is split into fixed-length clips and each clip is processed independently. However, when each chunk is refined from its guidance in isolation, the resulting mattes drift across chunk boundaries, producing flickering edges and discontinuous transparency.

To address this, we introduce two learnable token embeddings, a *preserve token* and a *refine token*, which assign each frame a per-frame mode that determines whether its matte is carried through unchanged or generated from coarse guidance. The preserve token marks anchor frames whose alpha mattes should be retained, while the refine token marks frames that should be converted from coarse guidance into fine-grained alpha mattes. DiT's full self-attention propagates information from preserve frames to refine frames, yielding mattes that remain consistent across chunks.

Why not auto-regressive diffusion? We chose to propagate fully-formed reference mattes through the preserve token rather than reuse the previous chunk's key-value (KV) cache as in auto-regressive diffusion approaches [18, 56]. Hidden-state propagation implicitly turns the model into a propagation-based matting framework, where each chunk's prediction is determined more by the previous chunk's internal state than by its own mask guidance. This re-entangles localization and fine-detail recovery, the two subproblems we explicitly decoupled by adopting mask-guided conditioning. The preserve token instead exposes the previous chunk's contribution as an explicit alpha matte, so each chunk remains fully grounded in its own mask guidance while still receiving long-range temporal context through self-attention.

Training. We train both tokens jointly within a single clip. For each training clip, we randomly select a subset of frames as preserve frames, replace their guide input with the ground-truth alpha matte, and add the preserve token embedding to those frames. The remaining frames receive the refine token together with the standard binary mask guide input. The model is supervised against the ground-truth matte at all frames through the pixel-level supervision, learning to reconstruct the supplied matte at preserve frames and to generate fine-grained mattes from coarse guidance at refine frames. Joint training within a single clip is essential, as it forces the DiT to learn a unified representation in which preserve and refine frames interact through self-attention, rather than treating them as two disjoint tasks.

Table 1: **Quantitative comparison of video matting methods on various benchmarks.**

Dataset	Metrics	SAM3	MGM	MGM-IW	MaGGIe	MaGGIe-FT	MatAnyone	MatAnyone 2	VideoMaMa	Ours
VideoMatte (1920 × 1080)	MAD ↓	5.05	11.28	5.01	4.26	4.44	4.17	4.04	4.34	**3.86**
	MSE ↓	0.94	3.23	0.35	0.27	0.33	0.27	0.24	0.27	**0.12**
	Grad ↓	12.02	6.33	5.01	3.61	4.53	3.59	3.13	4.06	**1.79**
	dtSSD ↓	3.11	2.11	1.38	1.23	1.35	1.17	1.13	1.21	**0.82**
YoutubeMatte (1920 × 1080)	MAD ↓	2.35	3.59	2.41	1.67	1.52	1.74	1.52	1.41	**1.15**
	MSE ↓	1.11	0.53	0.38	0.50	0.37	0.53	0.44	0.30	**0.23**
	Grad ↓	14.63	4.27	4.25	5.11	5.36	7.37	6.46	4.08	**3.27**
	dtSSD ↓	3.28	1.57	1.37	1.46	1.47	1.56	1.48	1.31	**1.17**
V-HIM60 (Hard)	MAD ↓	3.91	66.31	2.89	2.72	2.81	4.04	4.73	2.30	**1.69**
	MSE ↓	2.09	49.02	0.84	1.03	1.03	2.29	2.93	0.75	**0.45**
	Grad ↓	18.37	17.53	4.86	6.05	6.53	7.04	6.69	5.04	**2.92**
	dtSSD ↓	5.08	12.73	2.74	3.18	3.12	4.08	3.70	2.51	**1.98**
	Conn ↓	6.97	114.70	4.08	4.72	4.84	7.10	8.06	3.83	**2.72**

3.4 Temporally Consistent Inference

We deploy the preserve and refine tokens at inference through two complementary mechanisms: a sliding-window scheme that propagates mattes between adjacent chunks for local consistency, and a reference anchor that supplies a globally reliable matte at every chunk to prevent long-range drift.

Sliding window. At inference, we apply a sliding-window scheme in which consecutive chunks share an overlap region. The overlapping frames receive the previous chunk's predicted matte as input together with the preserve token, so that their mattes are passed through unchanged, while the remaining frames receive the guide signal together with the refine token. Although the preserve token is applied only to the overlap frames, DiT's full attention allows them to condition the generation of every other frame in the chunk, transferring temporal context across the chunk boundary and producing chunk-consistent mattes.

Reference anchor. The sliding-window scheme alone is vulnerable to error accumulation, since a corrupted overlap frame, e.g., due to occlusion, tracking drift, or other video artifacts, propagates through the preserve token to all subsequent chunks and cascades local failures through the rest of the video. To prevent this, we introduce a *reference anchor*, a globally reliable frame whose matte is supplied as preserve input alongside the local overlap frames at every chunk, providing a stable, error-free signal that the model can attend to throughout inference while the overlap frames continue to carry short-range temporal context. Crucially, this capability emerges directly from the preserve token without any additional training: because the preserve token already learns to treat its inputs as anchors, any user-supplied matte can serve as an anchor for the rest of the video.

This design gives users explicit control over which frame anchors the inference. By default, we reuse the matte predicted at the first chunk as the reference anchor throughout the video, which we find sufficient for most sequences. When the first frame is unreliable due to heavy occlusion, motion blur, or partial visibility, the user can instead select a later frame with cleaner foreground appearance as the reference, allowing the entire sequence to be conditioned on the most informative anchor available.

4 Experiments

4.1 Implementation Details

We build VideoMaMa++ on top of the Wan 2.1 (1.3B) [46] model, training on a combination of image and video matting datasets. To preserve the generative prior, we adapt the DiT to the video matting task using LoRA [14] rather than fine-tuning the full backbone. Training uses 3-frame clips at 1024×1024 resolution, which provides sufficient capacity for both spatial detail and temporal modeling. At inference, we process the video in chunks of 6 frames with a 2-frame overlap between consecutive chunks, where one of the overlap frames is the reference anchor. Although the training and inference settings differ in both frame count and chunking behavior, the model generalizes across them due to RoPE's [43] relative position encoding, which removes the dependence on absolute resolution and clip length present in fixed positional embeddings.

4.2 Quantitative Evaluations

We evaluate VideoMaMa++ on three benchmarks: VideoMatte240K [28], YouTubeMatte [54], and V-HIM60 (Hard) [19]. All mask-guided methods, including ours, receive identical guidance masks

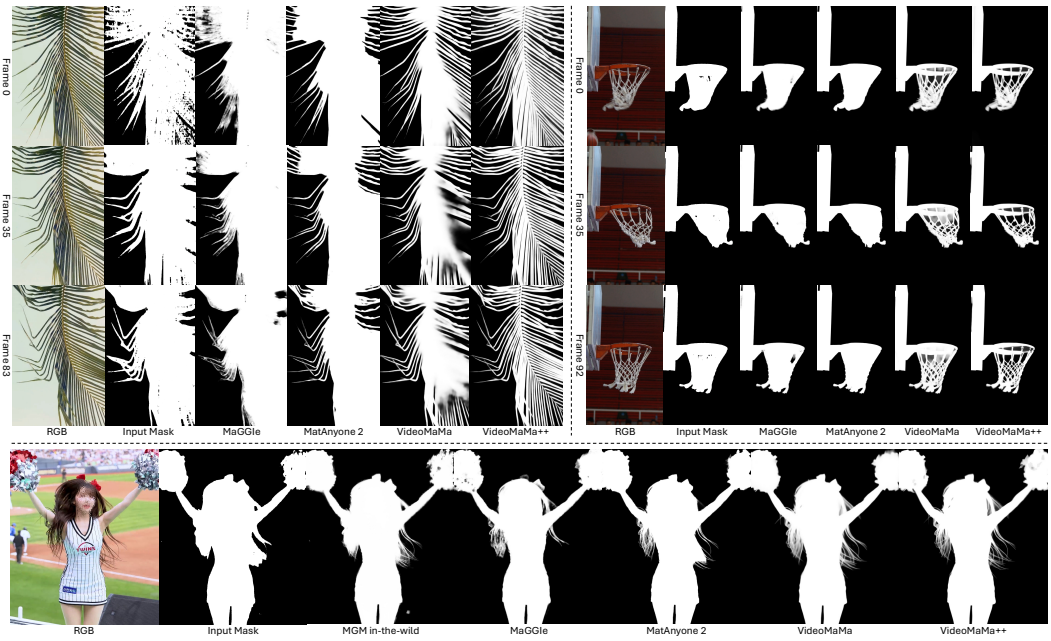


Figure 3: **Long-video temporal consistency.** Qualitative comparison between VideoMaMa++ and other methods. The comparison methods exhibits visible flickering at chunk boundaries or struggles to capture fine-grained details, while VideoMaMa++ produces smooth transitions across chunks while still capturing fine-grained details that emerge in later chunks.

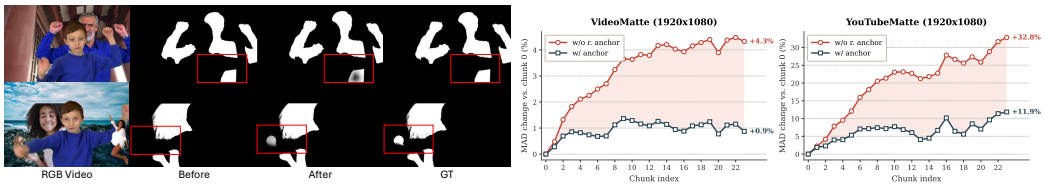
produced by a single procedure (Algorithm 1). The binarized first-frame ground-truth matte prompts SAM3 [3] on a duplicated first frame (frames $\{0, 0, 1, 2, \dots\}$), and SAM3 propagates it across the video, so every guidance mask, including that of the first frame, is a SAM3 prediction rather than the binarized ground truth, and no method receives a privileged mask. For one YouTubeMatte sequence, a sheet of paper that belongs to the foreground enters mid-video, so all methods share a track initialized at a later frame in which the paper is visible; VideoMaMa++ remains the best method under first-frame initialization as well (Appendix E).

As reported in Table 1, VideoMaMa++ achieves the best performance across all metrics on every benchmark. The improvements are particularly pronounced on Gradient [38] and dtSSD [11], which respectively measure spatial boundary sharpness and temporal smoothness. The Gradient gains confirm that the DiT-based generative prior recovers fine-grained boundary detail more accurately than prior mask-guided methods, while the dtSSD improvements indicate that the preserve and refine token mechanism produces temporally consistent mattes despite chunk-wise inference. We provide a further analysis of discontinuities at chunk boundaries in Appendix F.

4.3 Qualitative Evaluations

We compare VideoMaMa++ against representative video matting baselines, including MGM in-the-wild [34], MaGGIe [19], MatAnyone 2 [53], and VideoMaMa [27], on in-the-wild videos. Figure 3 shows long videos that diffusion models (VideoMaMa and ours) cannot process in a single chunk. As shown in the figure, the comparison methods either fail to capture fine-grained details or produce flickering results across chunks that are inconsistent over the full video. In contrast, VideoMaMa++ produces mattes that retain fine-grained boundary detail consistently throughout the sequence and remain stable across distant frames, demonstrating that our DiT-based generative prior, combined with the preserve and refine token mechanism, generalizes effectively to long-form footage well beyond the length seen during training. We refer the reader to the supplementary material for additional comparisons that more clearly demonstrate temporal consistency in motion.

We further quantify this observation with an annotation-free test for chunk-boundary discontinuities (Appendix F), which detects spikes in the frame-to-frame matte change that recur at chunk seams. With both models run at identical seams, VideoMaMa's matte change spikes at seams in-the-wild clips, whereas VideoMaMa++ stays at chance level. On YouTubeMatte, where ground truth is available, the same test detects VideoMaMa's discontinuity and finds none in the ground-truth mattes,



(a) **Interactive refinement examples.** (b) **MAD across chunk index (relative to chunk 0).**

Figure 4: **(Left)** Interactive refinement examples. **(Right)** MAD across chunk index.

Table 3: **Reference anchor analysis.** We evaluate (a) the anchor frame within the first chunk and re-anchoring every 3 chunks (Refresh), and (b) a corrupted anchor frame: dilated (D) or eroded (E) guidance mask, motion blur (B), and full occlusion (Occ.), with kernel sizes in px. In (b), the corrupted first frame is excluded from evaluation. All evaluations are performed on the **V-Him60 (Hard)** benchmark. The **best** results in (a) are highlighted.

	(a) Anchor choice				**(b) Corrupted anchor frame**								
						Poor mask				Motion blur			
Metrics	First	Random	Last	Refresh	Clean	D5	D21	E5	E21	B7	B15	B25	Occ.
MAD ↓	**1.69**	1.70	1.72	1.70	1.71	1.71	1.73	1.73	1.74	1.74	1.74	1.75	1.76
MSE ↓	**0.45**	**0.45**	**0.45**	0.46	0.46	0.46	0.46	0.46	0.48	0.48	0.48	0.48	0.48
Grad ↓	**2.92**	2.94	2.95	2.94	2.95	2.94	2.94	2.97	2.95	3.04	3.05	3.05	3.06
dtSSD ↓	**1.98**	1.99	2.00	1.99	1.99	1.99	1.98	2.00	1.99	2.02	2.02	2.03	2.03
Conn ↓	**2.72**	2.74	2.76	–	2.75	2.75	2.78	2.78	2.81	2.81	2.81	2.83	2.84

supporting its validity. Per clip, VideoMaMa's seam contrast exceeds that of VideoMaMa++ at identical boundaries, and these conclusions are stable across difference statistics and detrending windows. Since the test targets only chunk-boundary discontinuities, it applies only to chunk-wise methods, which is why we compare against VideoMaMa.

Interactive refinement. The preserve and refine tokens also enable selective refinement of individual frames within a chunk. Given an initial prediction, we identify the frame with the largest error against the ground-truth matte, mark it with the refine token, and mark the remaining frames with the preserve token. Re-running the model under this configuration regenerates the marked frame while attending to the others as anchors, producing a refined matte that draws on temporal context from the rest of the chunk. As illustrated in Figure 4a, this interactive procedure consistently improves matting quality, even though no new guide signal is provided. Quantitative results for interactive refinement are reported in the appendix.

4.4 Ablation Studies

Reference anchor. We ablate the reference anchor by comparing two sliding-window inference settings: *w/o r. anchor*, which uses previous chunk's matte on the local overlap frames with preserve tokens, and *w/ r. anchor*, our full setting that additionally supplies the first-chunk matte as a reference anchor providing a global preserve input. As reported in Table 2, adding the reference anchor yields consistent improvements across all six metrics on both benchmarks, with the largest gains on boundary-sensitive measures such as Gradient. This confirms that propagating preserve information through local overlaps alone leaves the model vulnerable to error accumulation, and that the global anchor effectively stabilizes long-range inference while preserving the cross-chunk information transfer enabled by the preserve token.

Table 2: **Ablation study on reference anchor.**

Setting	**MAD ↓**	**MAD-T ↓**	**MSE ↓**	**GRAD ↓**	**CONN ↓**	**DTSSD ↓**
VideoMatte (1920 × 1080)						
w/o r. anchor	3.9505	4.7366	0.1518	2.2398	1.2817	0.8576
w/ r. anchor	3.8633	4.6354	0.1240	1.7887	1.0755	0.8206
YoutubeMatte (1920 × 1080)						
w/o r. anchor	1.2691	31.6900	0.2861	4.4151	2.3294	1.1803
w/ r. anchor	1.1526	29.4655	0.2288	3.2741	2.0606	1.1663

The aggregate metrics in Table 2 measure average matting quality but do not reveal how that quality evolves as the sliding window advances along the sequence. To probe error accumulation directly, we re-aggregate MAD by chunk index, and report the relative change with respect to the first chunk. As shown in Figure 4b, the *w/o r. anchor* variant exhibits a steadily rising curve on both benchmarks, with MAD growing markedly by the final chunk. In contrast, the *w/ r. anchor* variant remains substantially flatter throughout the sequence, confirming that the reference anchor effectively suppresses error propagation across long videos.

Choice and quality of the reference anchor. The first frame may not represent later frames well, for example when the subject is partially occluded or only enters the frame later. Using a random frame or the last frame of the first chunk as the anchor instead yields nearly unchanged performance across all metrics (Table 3a), indicating that any frame of the first chunk serves as an effective anchor. Re-anchoring to a recent frame every few chunks likewise matches the default setting, so a first-frame anchor does not become stale on V-HIM60 despite its substantial pose, scale, and appearance variation; refreshing nonetheless remains available for footage with more drastic changes.

We further corrupt the input of the anchor frame (Table 3b) with a poor guidance mask (erosion or dilation), horizontal motion blur on the RGB frame, or a real occluder composited over nearly the entire subject, whose region is also removed from the guidance mask (Appendix H). Because the anchor is the model's own first-frame prediction, the degraded anchor matte propagates to all later chunks. Nevertheless, performance degrades only marginally under every corruption. Full occlusion is the most severe case, followed by motion blur and poor guidance masks, and even a fully occluded anchor leaves the model close to its clean-anchor performance.

Effect of preserve token at chunk boundaries. We design a targeted experiment to verify that the preserve and refine tokens produce consistent mattes across chunk boundaries. Using a chunk size of 6 with an overlap of 2 frames, each overlap frame is generated twice: once as part of the preceding chunk and once as part of the following chunk. We measure the discrepancy between these two predictions for the same frame, which directly reflects how much the matte drifts when the model crosses a chunk boundary. The two settings differ only in how the overlap frames are treated by the second chunk: *w/o preserve* regenerates them independently from the coarse mask guidance, whereas *w/ preserve* marks them with the preserve token and conditions on the matte produced by the preceding chunk. To isolate the effect of the preserve token, the reference anchor is disabled in both settings. As reported in Table 4, *w/ preserve* reduces both MAD and Gradient discrepancy by a large margin, confirming that the preserve token effectively transfers cross-chunk information and yields mattes that align across chunk boundaries rather than drifting into independently chosen refinements.

Table 4: **Chunk boundary consistency.**

Method	MAD ↓	Gradient ↓
w/o preserve	0.7128	0.6811
w/ preserve	**0.2583**	**0.1235**

Robustness to mask generators. To evaluate the robustness of our model to diverse binary mask sources, we measure performance on V-HIM60 [19] using XMem [7], Cutie [6], SAM2 [36], and SAM3 [3] as binary mask generators in Table 5. The results show that the model maintains consistent performance across all mask generators, demonstrating its robustness to the choice of binary mask input. Beyond the choice of generator, the model also tolerates local mask errors: when the guidance mask is wrong on individual frames of a chunk, such as missing thin structures or drift under motion blur, cross-frame attention to neighboring frames recovers the matte. If the tracker instead locks onto a wrong object for an entire sequence, the matte follows that track, a limitation shared by all mask-guided methods.

Table 5: **Robustness to mask generators.**

	MAD ↓		Gradient ↓	
Input Mask	Input	Refined	Input	Refined
XMem	5.5587	3.1119	17.7265	3.6025
Cutie	3.6998	1.9430	13.5395	**2.9448**
SAM2	4.0844	1.8994	19.3595	3.2741
SAM3	3.9106	**1.6924**	18.3659	2.9196

Continuous guidance via SAM logits. We investigate replacing the binary guide mask with SAM's raw logit map, which preserves the soft confidence information discarded by binarization (Figure 5). The structures most affected by this collapse, such as hair strands, thin wires, and motion-blurred edges, exhibit below-threshold logits that are removed entirely once the mask is binarized. We replace the guide mask track M in Section 3.2 with a SAM logit map, encoded by the same VAE into a latent that substitutes for z_M in Eq. 2, and retrain the model end-to-end with this logit conditioning rather than swapping inputs only at inference. As reported in Table 6a, the logit-conditioned variant outperforms its binary counterpart across all four metrics, confirming that soft confidence carries signals critical for fine-detail recovery. However, this gain does not transfer to in-the-wild videos, where the model overfits to SAM's specific logit distribution. We therefore adopt binary masks as the default guidance.

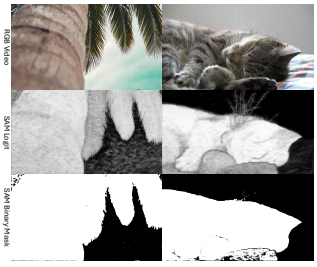


Figure 5: **SAM logit vs. Binary mask.**

Channel vs token concatenation. We compare two strategies for combining the RGB, mask, and noise inputs: channel concatenation, which stacks all three modalities along the channel axis into

Table 6: **Ablation studies on video matting design choices.** We evaluate the impact of our proposed components across different configurations. The ablations analyze (a) the advantage of incorporating SAM logits, and (b) channel-wise versus token-wise concatenation. All evaluations are performed on the **V-Him60 (Hard)** benchmark. The **best** results are highlighted.

(a) **Advantage of SAM logit**

Method	MAD	MAD-T	MSE	GRAD	CONN	DTSSD
with logit	**1.6272**	**40.8256**	**0.4247**	**2.9048**	**2.5843**	**1.9372**
Ours	1.6942	41.2279	0.4487	2.9196	2.7194	1.9796

(b) **Channel vs. Token concat**

Method	MAD	MAD-T	MSE	GRAD	CONN	DTSSD
Channel	2.2253	54.7369	0.6233	4.7119	3.8287	2.4183
Token	**1.6924**	**41.9167**	**0.4234**	**2.9813**	**2.7228**	**1.9758**

a single input tensor, and token concatenation, which appends each modality as a separate token stream. The channel-concat variant uses an expanded patch embedding to ingest the stacked input while keeping the rest of the DiT backbone identical to our model. Channel concatenation reduces the token sequence length and therefore the compute cost of self-attention, but as shown in Table 6b, it consistently underperforms token concatenation across all metrics. We attribute this to the loss of modality-specific information when the inputs are collapsed into a single tensor at the input stage.

Number of denoising steps. We evaluate how the number of denoising steps at inference affects matting quality. Following the observation of GenPercept [49] in monocular depth estimation, we find that a single denoising step is sufficient for the diffusion model to convert a binary mask into a high-quality alpha matte. As reported in Table 7, additional denoising steps lead the model to over-refine the matte, degrading both boundary sharpness and overall accuracy. We therefore adopt single-step inference as the default, which also yields a substantial efficiency gain for large-scale video processing (Appendix G).

Table 7: **Ablation on denoising steps.**

Steps	MAD ↓	MAD-T ↓	MSE ↓	GRAD ↓	CONN ↓	DTSSD ↓
1	**1.6680**	**41.7354**	**0.4027**	**2.9674**	**2.6769**	1.9546
2	1.7221	42.0193	0.4138	2.9708	2.7389	**1.9479**
3	2.7264	50.3829	0.5093	3.8081	3.3320	2.1253
5	10.7571	74.6095	2.0941	5.1494	17.5623	2.6155
10	16.6343	101.5099	3.9117	6.1175	31.1540	3.1936

Chunk size. Table 8 evaluates chunk sizes on V-HIM60 [19], with the overlap fixed at 2 frames across all settings. A 6-frame chunk outperforms both smaller (3) and larger (12) settings. Small chunks restrict the temporal context available for DiT's self-attention, whereas large chunks dilute the preserve frames' influence, reducing stability. A 6-frame chunk optimally balances temporal context with an effective preserve signal, serving as our default.

Table 8: **Ablation on chunk size.**

Chunk	MAD ↓	MAD-T ↓	GRAD ↓	CONN ↓	DTSSD ↓
3	2.1206	46.2131	3.5778	3.5776	2.2120
6	**1.6942**	**41.2279**	**2.9196**	**2.7194**	**1.9796**
12	1.7906	41.6290	3.0313	2.8858	2.0693

LoRA rank. We ablate the LoRA rank used to adapt the pretrained DiT for video matting. As reported in Table 9, performance improves monotonically as the rank increases, with rank 512 achieving the best results across all six metrics. The lower-rank variants (64 and 128) underperform on boundary-sensitive measures such as Gradient and MAD-T, indicating that mask-to-matte refinement requires sufficient adapter capacity to learn the fine-grained boundary structure absent from the input mask. We therefore adopt rank 512 as our default, which provides the best matting quality while remaining substantially lighter than full fine-tuning of the DiT backbone to preserve generative prior.

Table 9: **Ablation on LoRA rank.**

Rank	MAD ↓	MAD-T ↓	MSE ↓	GRAD ↓	CONN ↓	DTSSD ↓
64	2.4294	56.5292	0.8798	4.2180	4.2180	2.6284
128	2.3018	53.3100	0.8108	4.7875	3.9483	2.5732
512	**1.6942**	**41.2279**	**0.4487**	**2.9196**	**2.7194**	**1.9796**

5 Conclusion

We introduce VideoMaMa++, a diffusion transformer model that leverages generative priors for robust video matting, overcoming data scarcity. VideoMaMa++ introduces preserve and refine tokens that enable temporally consistent matting on long videos, addressing a fundamental limitation of diffusion-based approaches confined to short clips. We demonstrate that VideoMaMa++ achieves state-of-the-art performance across diverse benchmarks while producing fine-grained, temporally consistent results on long in-the-wild videos. An annotation-free analysis further shows that our chunk-wise inference exhibits no measurable discontinuity at chunk boundaries, unlike VideoMaMa processing chunks independently. Nevertheless, VideoMaMa++ is more computationally demanding than feed-forward models and, like all mask-guided methods, depends on the upstream mask track; reducing the DiT inference cost and benchmarking temporal smoothness on real footage are promising future directions. We believe VideoMaMa++ provides a scalable framework bridging the gap between short-clip generation and real-world long-video applications.

Acknowledgments

References

[1] Yanrui Bin, Wenbo Hu, Haoyuan Wang, Xinya Chen, and Bing Wang. Normalcrafter: Learning temporally consistent normals from video diffusion priors. *arXiv preprint arXiv:2504.11427*, 2025.

[2] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.

[3] Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, et al. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*, 2025.

[4] Sumit Chaturvedi, Mengwei Ren, Yannick Hold-Geoffroy, Jingyuan Liu, Julie Dorsey, and Zhixin Shu. Synthlight: Portrait relighting with diffusion model by learning to re-render synthetic faces. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 369–379, 2025.

[5] Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural Information Processing Systems*, 37:24081–24125, 2024.

[6] Ho Kei Cheng, Seoung Wug Oh, Brian Price, Joon-Young Lee, and Alexander Schwing. Putting the object back into video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3151–3161, 2024.

[7] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European conference on computer vision*, pages 640–658. Springer, 2022.

[8] Jun Myeong Choi, Annie Wang, Pieter Peers, Anand Bhattad, and Roni Sengupta. Scribblelight: Single image indoor relighting with scribbles. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5720–5731, 2025.

[9] Justin Cui, Jie Wu, Ming Li, Tao Yang, Xiaojie Li, Rui Wang, Andrew Bai, Yuanhao Ban, and Cho-Jui Hsieh. Self-forcing++: Towards minute-scale high-quality video generation. *arXiv preprint arXiv:2510.02283*, 2025.

[10] Yusuf Dalva, Yijun Li, Qing Liu, Nanxuan Zhao, Jianming Zhang, Zhe Lin, and Pinar Yanardag. Layerfusion: Harmonized multi-layer text-to-image generation with generative priors. *arXiv preprint arXiv:2412.04460*, 2024.

[11] Mikhail Erofeev, Yury Gitman, Dmitriy S Vatolin, Alexey Fedorov, and Jue Wang. Perceptually motivated benchmark for video matting. In *BMVC*, volume 1, page 2, 2015.

[12] Wenshuo Gao, Xicheng Lan, and Shuai Yang. Anyportal: Zero-shot consistent video background replacement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18990–18999, 2025.

[13] Yongtao Ge, Kangyang Xie, Guangkai Xu, Li Ke, Mingyu Liu, Longtao Huang, Hui Xue, Hao Chen, and Chunhua Shen. Generative video matting. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–10, 2025.

[14] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.

[15] Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiaodong Cun, Yong Zhang, Long Quan, and Ying Shan. Depthcrafter: Generating consistent long depth sequences for open-world videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2005–2015, 2025.

[16] Longfei Huang, Yu Liang, Hao Zhang, Jinwei Chen, Wei Dong, Lunde Chen, Wanyu Liu, Bo Li, and Peng-Tao Jiang. Sdmatte: Grafting diffusion models for interactive matting. *arXiv preprint arXiv:2508.00443*, 2025.

[17] Wei-Lun Huang and Ming-Sui Lee. End-to-end video matting with trimap propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14337–14347, 2023.

[18] Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. Self forcing: Bridging the train-test gap in autoregressive video diffusion. *arXiv preprint arXiv:2506.08009*, 2025.

[19] Chuong Huynh, Seoung Wug Oh, Abhinav Shrivastava, and Joon-Young Lee. Maggie: Masked guided gradual human instance matting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3870–3879, 2024.

[20] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9492–9502, 2024.

[21] Zhanghan Ke, Jiayu Sun, Kaican Li, Qiong Yan, and Rynson WH Lau. Modnet: Real-time trimap-free portrait matting via objective decomposition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1140–1147, 2022.

[22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026, 2023.

[23] Jiachen Li, Vidit Goel, Marianna Ohanyan, Shant Navasardyan, Yunchao Wei, and Humphrey Shi. Vmformer: End-to-end video matting with transformer. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6678–6687, 2024.

[24] Jizhizi Li, Sihan Ma, Jing Zhang, and Dacheng Tao. Privacy-preserving portrait matting. In *Proceedings of the 29th ACM international conference on multimedia*, pages 3501–3509, 2021.

[25] Jizhizi Li, Jing Zhang, Stephen J Maybank, and Dacheng Tao. Bridging composite and real: towards end-to-end deep image matting. *International Journal of Computer Vision*, 130(2):246–266, 2022.

[26] Jizhizi Li, Jing Zhang, and Dacheng Tao. Referring image matting. In *Proceedings of the IEEE Computer Vision and Pattern Recognition*, 2023.

[27] Sangbeom Lim, Seoung Wug Oh, Jiahui Huang, Heeji Yoon, Seungryong Kim, and Joon-Young Lee. Videomama: Mask-guided video matting via generative prior. *arXiv preprint arXiv:2601.14255*, 2026.

[28] Shanchuan Lin, Andrey Ryabtsev, Soumyadip Sengupta, Brian L Curless, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Real-time high-resolution background matting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8762–8771, 2021.

[29] Shanchuan Lin, Linjie Yang, Imran Saleemi, and Soumyadip Sengupta. Robust high-resolution video matting with temporal guidance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 238–247, 2022.

[30] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

[31] Kunhao Liu, Wenbo Hu, Jiale Xu, Ying Shan, and Shijian Lu. Rolling forcing: Autoregressive long video diffusion in real time. *arXiv preprint arXiv:2509.25161*, 2025.

[32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[33] Yuexiao Ma, Xuzhe Zheng, Jing Xu, Xiwei Xu, Feng Ling, Xiawu Zheng, Huafeng Kuang, Huixia Li, Xing Wang, Xuefeng Xiao, et al. Flow caching for autoregressive video generation. *arXiv preprint arXiv:2602.10825*, 2026.

[34] Kwanyong Park, Sanghyun Woo, Seoung Wug Oh, In So Kweon, and Joon-Young Lee. Mask-guided matting in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023.

[35] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023.

[36] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.

[37] Mengwei Ren, Wei Xiong, Jae Shin Yoon, Zhixin Shu, Jianming Zhang, HyunJoon Jung, Guido Gerig, and He Zhang. Relightful harmonization: Lighting-aware portrait background replacement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6452–6462, 2024.

[38] Christoph Rhemann, Carsten Rother, Jue Wang, Margrit Gelautz, Pushmeet Kohli, and Pamela Rott. A perceptually motivated online benchmark for image matting. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1826–1833. IEEE, 2009.

[39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.

[40] Fatemeh Saleh, Sadegh Aliakbarian, Charlie Hewitt, Lohit Petikam, Antonio Criminisi, Thomas J Cashman, Tadas Baltrušaitis, et al. David: Data-efficient and accurate vision models from synthetic data. *arXiv preprint arXiv:2507.15365*, 2025.

[41] Hongje Seong, Seoung Wug Oh, Brian Price, Euntai Kim, and Joon-Young Lee. One-trimap video matting. In *European Conference on Computer Vision*, pages 430–448. Springer, 2022.

[42] Jiahao Shao, Yuanbo Yang, Hongyu Zhou, Youmin Zhang, Yujun Shen, Vitor Guizilini, Yue Wang, Matteo Poggi, and Yiyi Liao. Learning temporally consistent video depth from video diffusion priors. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22841–22852, 2025.

[43] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

[44] Yanan Sun, Chi-Keung Tang, and Yu-Wing Tai. Ultrahigh resolution image/video matting with spatio-temporal sparsity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14112–14121, June 2023.

[45] Yanan Sun, Guanzhi Wang, Qiao Gu, Chi-Keung Tang, and Yu-Wing Tai. Deep video matting via spatio-temporal alignment and aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6975–6984, 2021.

[46] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.

[47] Junying Wang, Jingyuan Liu, Xin Sun, Krishna Kumar Singh, Zhixin Shu, He Zhang, Jimei Yang, Nanxuan Zhao, Tuanfeng Y Wang, Simon S Chen, et al. Comprehensive relighting: Generalizable and consistent monocular human relighting and harmonization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 380–390, 2025.

[48] Tiantian Wang, Sifei Liu, Yapeng Tian, Kai Li, and Ming-Hsuan Yang. Video matting via consistency-regularized graph neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4902–4911, 2021.

[49] Guangkai Xu, Yongtao Ge, Mingyu Liu, Chengxiang Fan, Kangyang Xie, Zhiyue Zhao, Hao Chen, and Chunhua Shen. What matters when repurposing diffusion models for general dense perception tasks? *arXiv preprint arXiv:2403.06090*, 2024.

[50] Shaocong Xu, Songlin Wei, Qizhe Wei, Zheng Geng, Hong Li, Licheng Shen, Qianpu Sun, Shu Han, Bin Ma, Bohan Li, et al. Diffusion knows transparency: Repurposing video diffusion for transparent object depth and normal estimation. *arXiv preprint arXiv:2512.23705*, 2025.

[51] Tian-Xing Xu, Xiangjun Gao, Wenbo Hu, Xiaoyu Li, Song-Hai Zhang, and Ying Shan. Geometrycrafter: Consistent geometry estimation for open-world videos with diffusion priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6632–6644, 2025.

[52] Jinrui Yang, Qing Liu, Yijun Li, Soo Ye Kim, Daniil Pakhomov, Mengwei Ren, Jianming Zhang, Zhe Lin, Cihang Xie, and Yuyin Zhou. Generative image layer decomposition with visual effects. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7643–7653, 2025.

[53] Peiqing Yang, Shangchen Zhou, Kai Hao, and Qingyi Tao. MatAnyone 2: Scaling video matting via a learned quality evaluator. In *CVPR*, 2026.

[54] Peiqing Yang, Shangchen Zhou, Jixin Zhao, Qingyi Tao, and Chen Change Loy. Matanyone: Stable video matting with consistent memory propagation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7299–7308, 2025.

[55] Shuai Yang, Wei Huang, Ruihang Chu, Yicheng Xiao, Yuyang Zhao, Xianbang Wang, Muyang Li, Enze Xie, Yingcong Chen, Yao Lu, et al. Longlive: Real-time interactive long video generation. *arXiv preprint arXiv:2509.22622*, 2025.

[56] Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and Xun Huang. From slow bidirectional to fast autoregressive video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22963–22974, 2025.

[57] Jianshu Zeng, Yuxuan Liu, Yutong Feng, Chenxuan Miao, Zixiang Gao, Jiwang Qu, Jianzhang Zhang, Bin Wang, and Kun Yuan. Lumen: Consistent video relighting and harmonious background replacement with video generative models. *arXiv preprint arXiv:2508.12945*, 2025.

[58] Hongfei Zhang, Harold Haodong Chen, Chenfei Liao, Jing He, Zixin Zhang, Haodong Li, Yihao Liang, Kanghao Chen, Bin Ren, Xu Zheng, et al. Dvd: Deterministic video depth estimation with generative priors. *arXiv preprint arXiv:2603.12250*, 2026.

[59] Yunke Zhang, Chi Wang, Miaomiao Cui, Peiran Ren, Xuansong Xie, Xian-Sheng Hua, Hujun Bao, Qixing Huang, and Weiwei Xu. Attention-guided temporally coherent video object matting. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 5128–5137, 2021.

Appendix Overview

This supplementary material provides additional details on implementation, evaluation, and qualitative results that could not be included in the main paper due to space constraints.

The document is organized as follows:

- **Section A** provides additional implementation details, including the training configuration and the list of datasets used to train VideoMaMa++.
- **Section B** presents qualitative results for the interactive refinement experiment, illustrating how VideoMaMa++ progressively improves its own predictions across iterations through the preserve and refine token mechanism.
- **Section C** reports zero-shot evaluation of VideoMaMa++ on image matting benchmarks, demonstrating that our model generalizes effectively beyond video to single-image settings.
- **Section D** discusses annotation noise in the AM-2K benchmark and explains why MAD-T provides a more reliable measure of matting quality than MAD on this dataset.
- **Section E** describes the guidance-mask generation protocol shared by all mask-guided methods and reports a single evaluation exception on the YouTubeMatte static_0017 sequence, where SAM3 [3] must be re-initialized at a later frame to capture an object that enters the scene partway through the video, together with results under first-frame initialization
- **Section F** introduces an annotation-free test for chunk-boundary discontinuities, which quantifies temporal consistency on in-the-wild videos without ground truth, and validates it on YouTubeMatte.
- **Section G** reports the computational cost of VideoMaMa++ and prior methods, including a per-stage runtime breakdown and end-to-end throughput with mask generation.
- **Section H** details the reference-anchor experiments, including the anchor-refresh setting and the corruption protocol used to test robustness to a degraded anchor.
- **Section I** discusses the limitations of this work: its dependence on the upstream mask track, and the absence of a video matting benchmark that can rigorously evaluate temporal smoothness on real-world footage.
- **Section J** discusses the positive and negative societal impacts of mask-guided video matting.

A Additional Implementation Details

Training details. We train VideoMaMa++ with an effective batch size of 16 and a learning rate of 1×10^{-4} using the AdamW [32] optimizer. The training objective combines an L1 loss for pixel-wise accuracy and a gradient loss for preserving edge sharpness and boundary details. Training is conducted on 8 NVIDIA H100 GPUs for approximately 2.5 days. A detailed list of the training datasets is provided in Table 10.

Table 10: **Datasets used to train VideoMaMa and SAM2-Matte.**

Category	Datasets
Image	DVM [45], AM-2k [25], DAViD [40], P3M-10k [24], RefMatte [26]
Video	CRGNN [48], VideoMatte240K [28], DVM [45]
Background	VideoMatting108 [59], DVM [45]

B Interactive Refinement Qualitative Results

We provide qualitative results for the interactive refinement experiment in Figure 4a, complementing the quantitative results reported in Table 11. Without modifying the input mask guidance, VideoMaMa++ progressively improves its own predictions across iterations by alternating preserve and refine assignments within each chunk. As shown in the figure, interactive refinement recovers regions missed in the initial prediction and sharpens fine-grained details along the foreground boundary, drawing on context from adjacent frames through the DiT's self-attention. This confirms that the gains in Table 11 stem from genuine cross-frame reasoning

Table 11: **Ablation on interactive refinement.**

Iter.	MAD ↓	MAD-T ↓	GRAD ↓	CONN ↓	DTSSD ↓	MSE ↓
1	1.6924	41.9167	2.9813	2.7228	1.9758	0.4234
2	1.6688	41.8064	2.9677	2.6761	1.9589	0.3992
3	**1.6634**	**41.7927**	**2.9633**	**2.6655**	**1.9545**	**0.3949**

rather than simple smoothing, and demonstrates that the preserve and refine tokens enable the model to serve as its own corrector at inference time.

C Image benchmark

We further evaluate VideoMaMa++ in a zero-shot setting on P3M-10K [24] and AM-2K [25], a portrait image matting benchmark and animal matting benchmark, to assess its generalization beyond video. Guide masks are obtained from SAM3 [3] following the same procedure as in Section 4.2. Image matting benchmarks emphasize fine-grained detail recovery at high resolution, posing a stringent test for pixel-level accuracy. As reported in Table 12, VideoMaMa++ achieves state-of-the-art performance against prior portrait matting methods [44, 29, 13], despite not being specifically designed for portrait matting. We note that AM-2K contains ground-truth annotation errors; our method shows inflated MAD scores on this benchmark despite producing visually superior mattes. Additional discussion on Section D.

Table 12: **Quantitative evaluation on P3M-500 and AM-2k datasets.**

Method	P3M-500-NP				AM-2k		
	MAD ↓	MSE ↓	Grad ↓	Conn ↓	MAD ↓	MAD-T ↓	Grad ↓
Mask free methods							
RVM	10.89	6.57	14.97	18.81	-	-	-
SparseMat	11.14	6.36	13.47	18.13	-	-	-
GVM	8.39	4.89	12.89	14.21	-	-	-
Mask guided methods							
SAM3	5.22	2.31	21.79	8.37	6.8729	93.2540	27.8649
MaGGIe	4.68	1.54	12.07	7.78	**6.7721**	73.5696	11.2396
Ours	**4.23**	**1.23**	**10.00**	**6.90**	9.1688	**61.0324**	**8.3373**

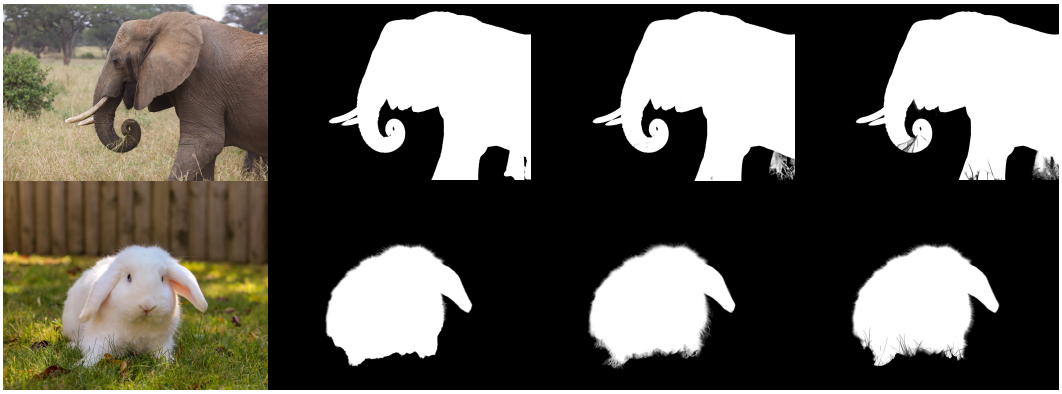


Figure 6: **Annotation noise in AM-2K [25].** The ground-truth alpha mattes in AM-2K often miss fine-grained internal details, that VideoMaMa++ correctly recovers. Because MAD averages pixel-wise error over the entire image, these missing details penalize accurate predictions and inflate the reported error. MAD-T, computed only within the transition region, isolates the matting quality that actually matters and avoids the bias introduced by such annotation gaps.

D AM-2K Justification

We observe that the AM-2K [25] ground-truth annotations contain systematic errors that disproportionately penalize fine-detail predictions under the MAD metric. As illustrated in Figure 6, the AM-2K ground truth often lacks detail within foreground regions, despite the predictions correctly recovering them. Because MAD averages absolute differences over the entire image, these annotation gaps inflate the reported error even when the predicted matte is visually superior. MAD-T, by contrast, restricts the evaluation to the transition region defined by the ground-truth trimap, isolating the boundary detail where matting quality matters most. Under this more reliable metric, VideoMaMa++ achieves state-of-the-art performance against all comparison methods, indicating that the apparent MAD gap on AM-2K reflects annotation noise rather than a true performance deficit.

E Guidance Mask Protocol and YouTubeMatte SAM3 Mask

Mask generation protocol. All mask-guided methods, including ours, receive the identical guidance mask track for every sequence, generated by Algorithm 1. We binarize the ground-truth alpha

Algorithm 1 Guidance mask generation, shared by all mask-guided methods.

Require: Frames $\{I_t\}_{t=0}^{T-1}$, ground-truth mattes $\{\alpha_t^{\mathrm{gt}}\}_{t=0}^{T-1}$, initialization frame s (s=0 except static_0017)
Ensure: Binary mask track $\{M_t\}_{t=0}^{T-1}$

$m \leftarrow \mathbf{1}\big[\alpha_s^{\mathrm{gt}} > \tau\big]$ ▷ binarize the GT matte, $\tau = 127$
$\mathcal{J} \leftarrow (I_s, I_0, I_1, \ldots, I_{T-1})$ ▷ prepend a duplicate; s=0 gives $\{0, 0, 1, 2, \ldots\}$
Initialize the SAM3 video predictor on $\mathcal{J}$ with mask prompt m at index 0
Propagate forward over $\mathcal{J}$ to obtain mask logits $\ell_0, \ell_1, \ldots, \ell_T$
for $t = 0$ **to** $T - 1$ **do**
    $M_t \leftarrow \mathbf{1}\big[\ell_{t+1} > 0\big]$ ▷ discard ℓ_0, the output on the duplicate
end for
return $\{M_t\}_{t=0}^{T-1}$

matte of the initialization frame and prepend a duplicate of that frame to the video; for the default first-frame initialization, SAM3 [3] thus processes the frame sequence $\{0, 0, 1, 2, \ldots\}$. The binarized matte is given as a mask prompt on the prepended duplicate, SAM3 propagates it forward over the whole sequence, and we binarize the resulting per-frame mask logits while discarding the output on the duplicate. Every frame's guidance mask, including that of the initialization frame itself, is therefore a SAM3 prediction rather than the binarized ground truth. The initialization frame is the first frame for every sequence, with the single exception described below. No method therefore receives a privileged mask.

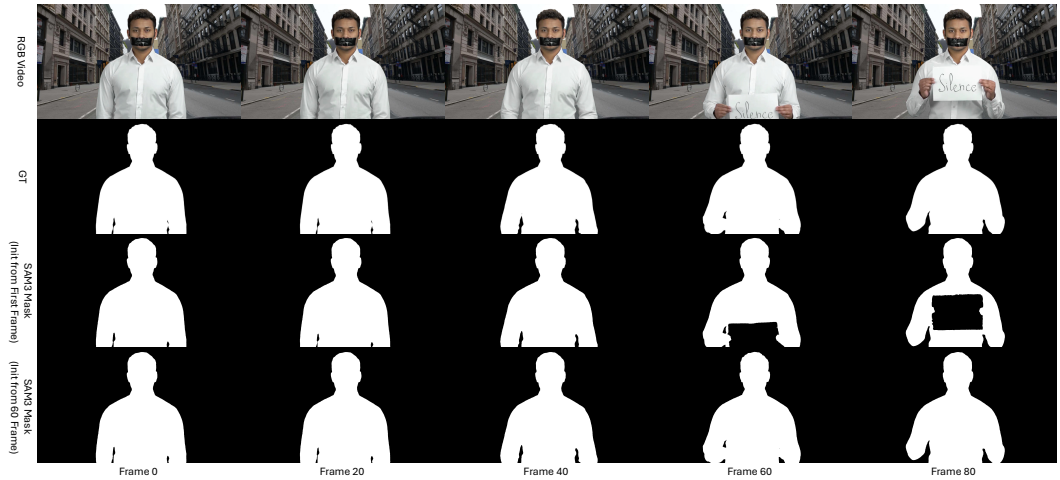


Figure 7: **Re-initialization of SAM3 on YouTubeMatte static_0017.** Initializing SAM3 [3] on the first frame fails to track a sheet of paper that the subject begins holding mid-sequence, since the object is not yet visible at initialization. Initializing on a later frame where both are simultaneously visible recovers the correct mask track.

Re-initialized sequence. For all evaluations, the input masks are obtained automatically without human intervention, with one exception. The static_0017 sequence from YouTubeMatte [54] presents a granularity ambiguity at the first frame: a sheet of paper that becomes part of the foreground only enters the scene partway through the video, and is therefore not visible at the moment of mask initialization. As a result, when SAM3 [3] is prompted on the first frame, the resulting mask track captures only the subject and cannot recover the paper once it appears, as shown in Figure 7. Because mask-guided matting models, including MaGGIe [19], VideoMaMa [27], and ours, faithfully follow the granularity defined by the input mask, the paper is consequently excluded from the predicted alpha matte. To resolve this without manual annotation, we initialize SAM3 on a later frame (frame 60) where both the subject and the paper are simultaneously visible, and propagate the prompt across the entire sequence. This re-initialized track is used for all mask-guided methods. We note that the

Table 13: **Effect of re-initializing YouTubeMatte static_0017.** "Re-init": SAM3 re-initialized at frame 60 for static_0017, as used in Table 1. "First-frame": SAM3 initialized on the first frame for every sequence. All metrics ↓.

Metric	Init.	MGM	MGM-IW	MaGGIe	MaGGIe-FT	VideoMaMa	Ours
MAD	First-frame	3.87	2.87	2.29	2.30	1.88	**1.61**
	Re-init	3.59	2.41	1.67	1.52	1.41	**1.15**
MSE	First-frame	0.78	0.83	0.93	0.95	0.76	**0.69**
	Re-init	0.53	0.38	0.50	0.37	0.30	**0.23**
Grad	First-frame	4.37	4.37	7.02	6.42	4.20	**3.31**
	Re-init	4.27	4.25	5.11	5.36	4.08	**3.27**
dtSSD	First-frame	1.61	1.40	1.64	1.58	1.33	**1.20**
	Re-init	1.57	1.37	1.46	1.47	1.31	**1.17**

MatAnyone [54, 53] series is not affected by this issue, as it does not rely on an external mask track. We report this exception in the interest of full transparency.

Results under first-frame initialization. Table 13 compares two settings on YouTubeMatte: the re-initialized track used in the main paper, and first-frame initialization for every sequence. All mask-guided methods degrade under first-frame initialization, since the sheet of paper is then missing from their guidance, and VideoMaMa++ remains the best method in both settings. The degradation itself reflects the behavior described above: the predicted matte follows the granularity of its guidance mask. Because a guidance mask is cheap to inspect and correct while a matte is not, we regard this faithful mask following as a desirable property for practical pipelines.

F Annotation-free Temporal Consistency

Ground-truth alpha mattes are unavailable for long real-world footage, so reference-based temporal metrics such as dtSSD cannot be computed on in-the-wild videos (Section I). Chunk-wise inference, however, has a characteristic failure mode that can be measured without annotation: a discontinuity where separately processed chunks meet. We therefore introduce a ground-truth-free test for such discontinuities.

Seam signal. Given a predicted matte sequence $\{\hat{\alpha}_t\}$, we compute the frame-to-frame change

$$d_t = \mathcal{A}\big(|\hat{\alpha}_{t+1} - \hat{\alpha}_t|\big), \qquad \tilde{d}_t = d_t - \operatorname{med}\big(d_{t-w}, \ldots, d_{t+w}\big), \tag{3}$$

where $\mathcal{A}$ aggregates over pixels, either as the mean over foreground pixels (fg-mean) or as the maximum (MAX). Subtracting a running median removes change caused by genuine motion, so $\tilde{d}_t$ retains only spikes above the local neighborhood.

Chunk cadence. To remove any confound from the two models tiling time differently, we run VideoMaMa [27] and VideoMaMa++ with the same chunk cadence: a 6-frame first chunk followed by 4 new frames per chunk. With frames indexed from 0 and transition t denoting the change from frame t to frame $t{+}1$, seams fall at transitions $\mathcal{S} = \{5, 9, 13, \ldots\}$ for both models. A 100-frame YouTubeMatte clip thus has 99 transitions, of which 24 are seams $(5, 9, \ldots, 97)$ and 75 are interior. VideoMaMa processes each chunk without overlap, whereas VideoMaMa++ additionally conditions each chunk on its preserve-token frames.

Statistics. We test whether the spikes in $\tilde{d}_t$ recur at the chunk period with two statistics. First, we compute the autocorrelation of $\tilde{d}_t$ per clip at the seam period (lag 4) and at twice the seam period (lag 8), and average it over clips. Second, we call a clip *seam-elevated* if the mean of $\tilde{d}_t$ over $t \in \mathcal{S}$ exceeds its mean over interior transitions, and refer to the gap between the two as its *seam contrast*. Under the null hypothesis of no chunk-boundary discontinuity, a clip is seam-elevated with probability $1/2$, which the ground-truth control confirms empirically (Table 14). For the seam-elevated counts, we report p-values from an exact one-sided sign test against this chance level.

Table 14: **Annotation-free seam test.** Both models run at identical chunk seams. *Left:* ground-truth validation on YouTubeMatte (64 clips; 24 seam vs. 75 interior transitions per clip). Entries count clips whose mean seam transition exceeds their mean interior transition (chance: 32/64). *Right:* 16 in-the-wild clips (fg-mean), for which no ground truth exists. Higher autocorrelation at the seam lag and more seam-elevated clips indicate a stronger chunk-boundary discontinuity. p-values for seam-elevated clips are from an exact one-sided sign test against chance.

	YouTubeMatte		In-the-wild (no GT)		
	Seam-elevated clips		Autocorr.		Seam-elevated
Model	MAX	fg-mean	lag 4	lag 8	clips
VideoMaMa [27]	62/64 $p{=}1.1{\times}10^{-16}$	63/64 $p{=}3.5{\times}10^{-18}$	$+0.35$ $p{=}5{\times}10^{-3}$	$+0.34$ $p{=}6{\times}10^{-3}$	14/16 $p{=}2.1{\times}10^{-3}$
GT (no chunking)	34/64 $p{=}0.35$	32/64 $p{=}0.55$	–	–	–
VideoMaMa++ (Ours)	18/64 $p{>}0.99$	18/64 $p{>}0.99$	$+0.03$ $p{=}0.66$	$+0.11$ $p{=}0.12$	7/16 $p{=}0.77$

Table 15: **Computational cost** on a single H100 at resolutions 256^2 to 2048^2 ($1\text{K} = 1024^2$, $2\text{K} = 2048^2$), excluding mask generation for all methods. For VideoMaMa++, the cost of each 6-frame window is amortized over its 4 new frames, so FPS accounts for the recomputed overlap frames. MatAnyone/MatAnyone 2 and MGM/MGM-ITW share architectures and are listed once.

	FPS ↑					Peak memory (GB) ↓					TFLOPs ↓				
Method	256	512	768	1K	2K	256	512	768	1K	2K	256	512	768	1K	2K
MGMatting	66.71	66.23	63.56	60.44	18.05	0.58	0.84	1.37	1.98	5.41	0.036	0.115	0.240	0.411	1.55
MaGGIe	23.2	19.3	20.0	15.0	7.5	0.52	1.18	2.52	4.88	16.01	0.021	0.083	0.187	0.332	1.33
MatAnyone	43.6	42.0	38.9	34.6	7.5	0.33	0.67	1.32	3.43	22.40	0.030	0.124	0.289	0.542	2.94
VideoMaMa	27.56	13.64	6.06	3.41	0.71	4.75	5.36	6.41	7.88	17.05	1.60	6.55	15.4	28.9	159
VideoMaMa++	16.4	6.9	2.6	1.17	0.11	3.73	5.17	7.86	11.33	35.59	4.25	24.8	85.4	225	2920

Validation on YouTubeMatte. We first validate the test on YouTubeMatte, where ground truth is available (Table 14, left). VideoMaMa's seam transitions exceed its interior transitions in nearly all clips under both summaries. As a control, we apply the identical test to the ground-truth mattes at the same transition positions. Since the ground truth involves no chunking, any elevation found there would indicate that the test responds to content rather than to chunk boundaries; it detects nothing, with counts at the chance level. VideoMaMa++ shows no seam elevation: its seams exceed its interior transitions in fewer clips than expected by chance under both summaries. Per clip, VideoMaMa's seam contrast exceeds that of VideoMaMa++ in 62–64 of 64 clips (sign test, $p < 10^{-15}$).

In-the-wild results. On the 16 in-the-wild videos (Table 14, right), the detrended signal of VideoMaMa autocorrelates at the chunk period (lags 4 and 8), and its seam transitions exceed its interior transitions in most clips. At identical seams, VideoMaMa++ shows none of this structure: neither autocorrelation is significant, and the number of seam-elevated clips is consistent with chance. These results quantitatively corroborate Figure 3: at matched boundaries, VideoMaMa exhibits a chunk-boundary discontinuity that VideoMaMa++ does not.

Robustness and scope. These conclusions are unchanged across alternative difference statistics and detrending windows of 3 to 11 frames. The test targets only discontinuities at chunk boundaries: it measures neither per-frame accuracy nor flicker within a chunk, and it applies only to methods that process video in chunks, which is why the comparison is restricted to VideoMaMa.

G Computational Cost

Table 15 reports throughput, peak memory, and compute on a single H100 from 256^2 to 2048^2. VideoMaMa++ is the most expensive of the compared methods: at 1024^2 it runs at 1.17 FPS with 11.3 GB of peak memory, compared to 3.41 FPS and 7.9 GB for VideoMaMa, while the feed-forward

Table 16: **Per-output-frame cost of VideoMaMa++** (ms) on a single H100, including SAM3 mask propagation. Matting-stage costs are measured per 6-frame window and divided by its 4 new frames. E2E FPS $= 1000 \,/$ (SAM3 + matting total).

Res.	SAM3	VAE Enc.	DiT (1-step)	VAE Dec.	Matting total	Matting FPS	E2E FPS
256^2	45.4	24.9	20.1	15.9	60.9	16.4	9.4
512^2	45.4	37.9	79.2	27.5	144.6	6.9	5.3
768^2	45.4	78.3	254.3	55.8	388.4	2.6	2.3
1024^2	45.4	139.0	616.1	97.7	852.8	1.17	1.11
2048^2	45.4	681.0	7602.1	448.6	8731.7	0.11	0.11

baselines run at 15–60 FPS. Table 16 breaks down this cost by stage. At 1024^2, the single DiT forward pass accounts for 616 ms of the 853 ms per output frame, with the VAE encoder and decoder taking 139 ms and 98 ms. SAM3 mask propagation runs once per frame at a fixed resolution and adds 45.4 ms per frame regardless of the matting resolution, giving an end-to-end rate of 1.11 FPS at 1024^2 and 9.4 FPS at 256^2. Since the DiT pass dominates, single-step inference (Table 7) is essential for keeping this cost practical, as each additional denoising step would add a full DiT evaluation per window.

H Reference Anchor Analysis Details

This section details the reference-anchor experiments of Table 3, all conducted on V-HIM60 [19].

Anchor choice and refresh. The first frame may not be representative of later frames. We therefore replace it as the anchor with a random frame or the last frame of the first chunk, which leaves performance nearly unchanged across all metrics. Because the foreground may also change in appearance over long sequences, we further re-anchor to a recent frame every 3 chunks. V-HIM60 sequences contain substantial pose, scale, and appearance variation, yet refreshing yields nearly identical quality to the default setting, so a first-frame anchor does not become stale on this benchmark. Refreshing nevertheless remains available at negligible cost for footage whose appearance changes more drastically.

Corruption protocol. We corrupt the input of the first frame, which serves as the anchor, rather than the propagated matte. Because the anchor is the model's own first-frame prediction, a corrupted input yields a degraded anchor matte that is propagated to all later chunks. We consider three corruptions:

- **Poor guidance mask.** We erode (under-segmentation, E) or dilate (over-segmentation, D) the first-frame guidance mask using grayscale morphology with a square kernel of $k \in \{5, 21\}$ px.
- **Motion blur.** We apply a horizontal motion-blur kernel of $k \in \{7, 15, 25\}$ px to the first-frame RGB (B), so the model must matte a motion-blurred appearance.
- **Occlusion.** We composite a real animal foreground from AM-2K [25] over the subject using its ground-truth matte, enlarging it until at least 99% of the subject is covered, and zero the occluded region of the guidance mask so that it remains consistent with the now-invisible subject (Occ.).

We exclude the corrupted first frame from evaluation, so the metrics measure how the degraded anchor propagates to the rest of the video. This matters most under occlusion, where the input RGB and the ground-truth matte of the first frame disagree. It is also why the clean baseline in Table 3b differs slightly from the main results: the only difference is the excluded clean first frame.

Results. Across all corruptions, performance degrades only marginally relative to the clean baseline. Full occlusion is the most severe case, followed by motion blur and poor guidance masks, and larger kernels generally lead to slightly larger degradation. Even with a fully occluded anchor, the model remains close to its clean-anchor performance, indicating that a degraded anchor has only a limited effect on the rest of the video.

I Limitation

Dependence on the mask track. VideoMaMa++ deliberately separates tracking, delegated to a mask generator such as SAM3 [3], from matting, so its robustness is bounded by the quality of the mask track. Errors confined to individual frames of a chunk, such as missing thin structures or drift under motion blur, can be corrected through cross-frame attention to neighboring frames or through interactive refinement (Section B). However, if the tracker locks onto the wrong object for an entire sequence, the matte follows that wrong track. This limitation is shared by all mask-guided matting methods.

Evaluating temporal smoothness on real video. A key limitation of this work, and of video matting research more broadly, is the absence of a benchmark that can rigorously evaluate temporal smoothness on real video. Existing benchmarks [28, 54, 27] are constructed by compositing per-frame foregrounds onto random backgrounds, which provides accurate per-frame ground truth but does not capture the temporal dynamics of naturally captured footage. As a result, temporal consistency on real-world videos is still largely assessed qualitatively. The annotation-free seam test of Section F quantifies one specific failure mode, discontinuities at chunk boundaries, without ground truth. However, it measures neither per-frame accuracy nor flicker within a chunk, and it applies only to methods that process video in chunks, so it cannot compare the overall temporal behavior of arbitrary methods. Developing a benchmark with reliable ground-truth temporal smoothness, ideally captured rather than composited, remains an open problem for the community.

J Societal Impact

Positive Societal Impact. Our mask-guided video matting framework can substantially benefit creative industries, education, and communication. By reducing the manual effort required for high-quality video matting, the method lowers the barrier to professional-grade video production, enabling independent creators, educators, and small studios to produce polished content without specialized expertise or costly proprietary software. In film and visual effects, it can accelerate labor-intensive rotoscoping workflows, allowing artists to focus on creative decisions rather than frame-by-frame manual labeling. The technology also improves the quality of background replacement in video conferencing and virtual classrooms, supporting remote collaboration and accessibility-oriented applications such as isolating speakers or interpreters in educational content.

Negative Societal Impact. Despite these benefits, high-quality video matting can be misused. As a core component in video manipulation pipelines, improved matting may facilitate the creation of more convincing deepfakes and manipulated media, contributing to misinformation and the erosion of public trust in video evidence.

NeurIPS Paper Checklist

1. **Claims**

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

   Answer: [Yes]

   Justification: We wrote abstract and introduction the accurately reflect paper's contributions.

   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: We discuss the limitation of the work performed by us in supplementary section.

   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: We do not have any theoretical results or assumptions to proof.

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: 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)?

   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: Our paper does not introduce new data, but all experiments were conducted using publicly available, open-source materials.

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: We shared our specific setting for the experiments on the main paper and supplementary file.

   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: [No]

   Justification: We do not report error bars or statistical significant on the experiments. We show qualitative comparison and quantitative comparison on video matting metrics.

   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: We reported our GPU setting and training time on main and supplementary file.

   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: We do follow the NeurIPS code of Ethics on our research.

   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: We discussed the societal impacts on supplementary section.

    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: We do not release any dataset or generate new contents using our model. Our model is video matting model which generate continuous value regarding the input videos.

    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: Our model is based on Wan 1.3B model, which is widely use in the research community. Dataset used the train our model is publicly available for research.

    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: [Yes]

    Justification: We do not made any dataset but trained a model using Wan 1.3B model. Training recipe are well documented on the paper.

    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: We do not include any crowdsourcing experiments and research with human subjects.

    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: Our paper does not involve crowdsourcing nor research with human subjects.

    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: On our reserach, LLM is only used for writing, formatting purposes and did not impact the core 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.