

FrescoDiffusion: 4K Image-to-Video with Prior-Regularized Tiled Diffusion

Hugo Caselles-Dupré¹^{*}, Mathis Koroglu^{1,2}^{*}, Guillaume Jeanneret²^{*},
Arnaud Dapogny², and Matthieu Cord^{2,3}

¹ Obvious Research, Paris, France

² Sorbonne Université, CNRS, ISIR, F-75005 Paris, France

³ Valeo.ai, France

Abstract. Diffusion-based image-to-video (I2V) models have improved substantially, but they still struggle with ultra-high-resolution inputs such as 4K images. Native-resolution generation often loses fine details, while tiled denoising preserves local structure but disrupts global layout consistency. This issue is especially challenging for fresco animation, where large artworks contain many characters, objects, and semantically distinct regions that must remain spatially coherent over time. We introduce FrescoDiffusion, a training-free method for coherent large-format I2V generation from a single complex image. Our approach augments tiled denoising with a precomputed latent prior: a low-resolution video is first generated at the model’s native resolution, then its latent trajectory is upsampled to provide a global spatiotemporal reference. During 4K generation, per-tile velocity predictions are fused with this reference at every diffusion step through a weighted least-squares objective, producing a closed-form update that improves global coherence while preserving fine detail. We also introduce a spatial regularization variable for region-level motion control. Experiments on VBench-I2V and our fresco I2V dataset show that FrescoDiffusion improves consistency, fidelity, and controllability over tiled baselines while remaining computationally efficient.

Keywords: 4K Image-to-Video · Prior-Regularized Training-Free Generation · Scheduled-Gated Regularization

1 Introduction

Media is shifting toward ultra high-definition video, with resolutions around 4K (3840×2160). In many applications, especially the creative industries such as cinema, animation, projections and art, there is a need for tools that can process and generate 4K content with sufficient detail. One example, and the focus of this work, is fresco animation. The same structural challenge also appears in projection mapping, animated films and motion comics, and urban scene simulation, although we evaluate only frescoes and VBench-I2V here. In contrast to standard images, we define a *fresco* as a large-scale image containing multiple

^{*} These authors contributed equally.




Fig. 1: From a single ultra-high-definition image (3500×3500), FrescoDiffusion animates it at the same resolution. We show three frames from the generated video. The red box marks a fixed spatial region tracked across time, illustrating motion temporal consistency and fine-detail preservation.

scenes (*i.e.*, up to hundreds) that blend seamlessly into a coherent visual. To achieve animation of any picture, diffusion models [3, 5, 18, 30, 33] have become the dominant approach for video synthesis, producing strong results in both text-to-video (T2V) and image-to-video (I2V) applications. While T2V models generate videos that faithfully follow an input prompt, I2V models follow the same paradigm but start with a first frame provided by the user.

Despite the progress in I2V, the most advanced models [15, 34, 36] operate under constraints that do not gracefully allow large-scale I2V. On the one hand, using regular video diffusion models on high-resolution images resized to their native spatial regime yields insufficient detail, as the input is too large and complex to be represented by the standard resolution of video diffusion models. On the other hand, memory and compute grow steeply with spatial and temporal dimensions, preventing the use of those models on 4K images.

Scaling image-to-video to very large videos containing many distinct regions remains largely unexplored in the literature. Existing work tries to address both issues, but the fresco setting makes it especially hard: simple tiling [4, 26] introduces cross-tile drift and visible seams, and post-hoc video super-resolution [16, 25] cannot create scene-level content that never existed at the model’s native scale. A key difficulty specific to fresco animation is that different regions of the image play fundamentally different roles over time. Frescoes typically contain numerous loosely coupled scenes and characters: some regions are visually static and define architectural context, while others are semantically active and expected to exhibit motion. Rather than a simplified 4K-I2V setting, this exposes a distinct failure mode (adding local motion without breaking global structure) and motivates a region-aware treatment of global coherence.

In this paper, we introduce FrescoDiffusion to achieve 4K-I2V up to the fresco setting, as illustrated in Fig. 1. Although animating the full-resolution fresco is

not feasible directly, video models can animate a resized thumbnail of the same image at the model’s native scale. So, our core idea is to take advantage of this feature and use the resized thumbnail as a prior to guide the high-resolution denoising of the initial fresco with a novel loss. Our loss admits a closed-form solution that allows control of the balance between creativity (tiled denoising) and prior alignment. Moreover, we introduce a tool that adapts the generation process in active zones and the background differently with a mask to deal with frescoes specifically. It allows the generative model enough flexibility to animate the active zones while keeping the background close to the prior. In our experiments, we show the superiority, both in terms of quality and computation time, of FrescoDiffusion compared to other tiled-denoising methods through extensive qualitative and quantitative evaluation, and a user preference study, on our novel dataset FrescoArchive, and standard VBench 4K dataset [20]. Videos generated with FrescoDiffusion are provided in the supplementary, with an interactive webpage for visualizing them conveniently.

To summarize, our contributions are as follows:

1. We introduce FrescoDiffusion, a training-free 4K I2V method with a prior-regularized tile fusion applied at *every* diffusion step. Our method leverages a closed-form solution of a newly introduced prior-based loss controlled with one interpretable strength λ .
2. To demonstrate its use in an artistic domain, we propose a dynamic prior-strength schedule with a spatial gating mechanism that controls the trade-off between creativity and prior similarity, especially for fresco-to-video.
3. We propose FrescoArchive, a new I2V dataset composed of complex multi-scene images for fresco-to-video evaluation.

To contribute further to 4K-I2V, we will make our algorithm and dataset public with the paper release. Project page : <https://obvious-research.github.io/frescodiffusion/>.

2 Related Work

Diffusion-based video generation. Early text-to-video systems such as Make-A-Video [33] and Imagen Video [18] generate short clips by cascading a base video diffusion model with spatial and temporal super-resolution modules. Subsequent work extends latent diffusion models [30] to videos, e.g. Video Latent Diffusion Models (Video LDM) [6], which map videos to a compressed latent space for efficient high-resolution text-to-video generation, and Lumiere [3], whose space-time U-Net jointly denoises all frames. Recent foundation-scale models such as Wan [36], Mochi [34] and LTX-Video [15] push video latent diffusion to higher quality and longer durations. Image-to-video methods (I2V) are usually supported by all these recent models. All these methods, including commercial systems (such as Kling, Sora2, Gen4.5 or Veo3), operate uniquely on their native spatial resolutions (typically 480–1080p) and video lengths of a few seconds. In

contrast, we formulate a training-free approach that takes an existing image-to-video diffusion model and extends it beyond its native resolution, bringing additional creative control.

Video super-resolution. Video super-resolution (VSR) seeks to reconstruct high-resolution videos from low-resolution inputs, typically with strong temporal coherence. While these methods [9–11, 16, 24, 38, 42], including recent streaming diffusion VSR such as FlashVSR [43], excel at producing globally coherent, temporally consistent videos, they are fundamentally constrained to staying close to the information present in the low-resolution input. Their objectives encourage fidelity to the input video under distortion and perceptual metrics, and any hallucinated detail is limited to local texture refinement. Starting from a 480p or 720p animation of a fresco, VSR can only upscale and slightly enrich this coarse representation; it cannot create hundreds of semantically distinct scenes that were never visible in the low-resolution video. In contrast, FrescoDiffusion performs tiled denoising directly on a large latent canvas and uses a thumbnail animation purely as a prior in latent space, allowing each tile to carry as much semantic content as a native-resolution video while maintaining global coherence.

Training-free high-resolution tiled denoising. A line of work studies generating large images and videos from pre-trained diffusion models without additional training. MultiDiffusion [4] fuses overlapping diffusion trajectories via a weighted least-squares objective, enabling large-scale image generation with many scenes, but without explicit global coherence. Several approaches build upon this idea to improve spatial consistency. Mixture of Diffusers [22] runs multiple regional diffusions on a shared canvas to control high-resolution composition, while Diff-Collage [40] formulates generation as a factor graph over patches and overlaps. SpotDiffusion [14] further reduces memory by denoising disjoint windows over time, trading overlap for efficiency. Recent training-free methods instead modify the sampling procedure of a single diffusion model to scale resolution. Scale-Crafter [17] and DemoFusion [13] progressively enlarge the effective receptive field through re-dilation, dispersed convolutions, or staged upscaling, treating the full canvas as a single sample rather than independent tiles.

Closer to our work, DynamicScaler [26] (DS) proposes an offset-shifting denoising strategy for panoramic or 360° video generation, where spatial windows are shifted across denoising steps to synchronize content and motion across fields of view.

It also relies on training-free tiling with global guidance from a low-resolution video, but differs from FrescoDiffusion in three ways. **(i)** DS injects the low-resolution prior *once*, as a noisy initialization, whereas we regularize *every* denoising step through a single quadratic objective. **(ii)** Our objective admits a single closed-form fused velocity via a per-step least-squares solution, whereas DS relies on heuristic window shifting and blending across steps rather than an explicit fusion objective. **(iii)** DS targets 360° panoramas with a single shared global motion, while we allow for regional controllability of where motion should appear.




Fig. 2: Overview of FrescoDiffusion. Starting from a 4K fresco image, we first build a global latent prior, by resizing the image to the native input size of the image-to-video backbone. Next, we upsample the prior latents x_{prior} to fit the 4K image size. We then apply tiled denoising to the large latent canvas, x_t^{4K} , obtaining per-tile flow predictions, $\{y_i\}$. We then use $\{y_i\}$ and x_{prior} to compute the optimal output velocity field (Eq. (6)) according to our loss ℓ_{FD} (Eq. (5)). This updated field is then used to update the large latent canvas, x_t^{4K} , with the flow-matching scheduler.

More broadly, FrescoDiffusion is related to latent-guidance methods such as Blended Latent Diffusion [1], which steer denoising toward a reference constraint. We instead regularize the fused velocity of an entire tiled video canvas and obtain a spatially varying closed-form update. Concurrent PixelWizard [23] learns a compact spatiotemporal anchor and a shortcut-trained native 2K/4K generator; FrescoDiffusion is training-free and reuses an existing I2V backbone.

3 FrescoDiffusion Method

In this section, we introduce FrescoDiffusion, a method tailored for multi-scene 4K-I2V. Our proposed approach consists of two steps, as shown in Fig. 2. First, we compute a prior thumbnail to guide the diffusion process. Second, during the denoising process, we analytically minimize the energy loss to produce the optimal fused output (Sec. 3.2). This output then guides the diffusion process. During this step, we optionally employ our novel masking strategy to direct the diffusion toward the prior by explicitly indicating which regions should be modified and which should converge to the prior (Sec. 3.3).

3.1 Ultra high-definition tiled denoising baseline framework

We start by introducing our baseline and notation. Let $t \in [0, 1]$ be the timestep, c be the conditioning tuple (input image and prompt) for the I2V flow-matching model f_θ , and $x_t \in \mathbb{R}^{C \times T \times H \times W}$ be the latent state. At each t , the model predicts a velocity field

$$y(x_t) = f_\theta(x_t, t, c). \quad (1)$$

Starting from $x_{t=0} \sim \mathcal{N}(0, I)$, a scheduler integrates these velocities up to $t = 1$. The resulting latent x_1 is then decoded to produce the video.

Next, we introduce MultiDiffusion [4] (MD), adapted to I2V as in prior work [26]. MD generates high-resolution latent codes, *i.e.* $x_1^{4K} \in \mathbb{R}^{C \times T \times H^{4K} \times W^{4K}}$ with $H^{4K} \gg H$ and $W^{4K} \gg W$, by running f_θ on overlapping tiles of a large “canvas” x_t^{4K} and merging the tile-wise predictions.

More specifically, let x_0^{4K} be the initial 4K latent canvas. Let $\mathcal{C}_p$ crop a window of shape (C, T, H, W) at position p , and let $\mathcal{P}_p$ zero-pad a tile back to shape (C, T, H^{4K}, W^{4K}) at coordinates p . For each tile i , define the tile prediction

$$y_i(x_t^{4K}) = \mathcal{P}_{p_i}(y(\mathcal{C}_{p_i}(x_t^{4K}))), \quad (2)$$

where p_i is the coordinate of the i^{th} tile. Given x_t^{4K} , and the tiled-velocity $\{y_i(x_t^{4K})\}_{i=1}^n$ predictions, MD solves for a single merged velocity y^* that best matches these overlapping tile predictions by minimizing the loss

$$\ell_{\text{MD}}(y^*; t) = \sum_{i=1}^n \left\| \sqrt{w_i} \odot (y^* - y_i(x_t^{4K})) \right\|_2^2, \quad (3)$$

where n is the number of windows, p_i are the coordinates and w_i are weight maps of window i used to reduce seams between tiles. This loss admits the closed-form solution

$$y_{\text{MD}}(x_t^{4K}) = \frac{\sum_{i=1}^n w_i \odot y_i(x_t^{4K})}{\sum_{i=1}^n w_i}. \quad (4)$$

Finally, y_{MD} is used to update the canvas, x_t^{4K} , using the iterative standard flow-matching sampling process.

3.2 FrescoDiffusion: Prior-Regularized Tile Fusion

MD provides a solution for merging overlapping windows using a weighted sum. However, MD lacks the ability to regularize window merging with an existing prior, such as the initial frame, to create a cohesive scene. To this end, we propose to extend ℓ_{MD} with a novel regularization term ℓ_{prior} . Our new FrescoDiffusion loss is

$$\ell_{\text{FD}}(y^*; t) = \underbrace{\left\| \sqrt{\lambda} \odot [(x_t^{4K} - \sigma_t y^*) - x_{\text{prior}}] \right\|_2^2}_{\ell_{\text{prior}}(y^*; t, x_{\text{prior}})} + \ell_{\text{MD}}(y^*; t) \quad (5)$$

Our loss has two terms: ℓ_{MD} reduces the disparity between the shifting-window outputs, while ℓ_{prior} minimizes the dissimilarity between the current-step prediction of the clean latent, *i.e.* $(x_t^{4K} - \sigma_t y)$ in the flow matching formulation, and the prior $x_{\text{prior}} \in \mathbb{R}^{C \times T \times H^{4K} \times W^{4K}}$. Here, λ is a regularization variable that can

be either a constant ($\lambda \in \mathbb{R}$) or a tensor ($\lambda \in \mathbb{R}^{C \times T \times H^{4K} \times W^{4K}}$), whose design is discussed in the next section. Additionally, σ_t denotes the scheduler’s discrete noise standard deviation at step t . Note that in other diffusion formulations, we just have to adapt the corresponding one-step prediction, see Sec. G.

Equation (5) is separable across canvas coordinates and strictly convex. Thus, the unique minimizer can be found in closed form by setting the derivative to zero. Therefore, the prior-regularized fused velocity is

$$y_{\text{FD}}(x_t^{4K}) = \frac{\sigma_t \cdot \lambda \odot (x_t^{4K} - x_{\text{prior}}) + \sum_{i=1}^n w_i \odot y_i(x_t^{4K})}{\sigma_t^2 \cdot \lambda + \sum_{i=1}^n w_i}. \quad (6)$$

When $\lambda = 0$, the closed-form solution in Eq. (6) reduces to the MultiDiffusion fusion in Eq. (4). Prior errors enter the update proportionally to λ : a large value can propagate incorrect structure or motion from an unreliable thumbnail, whereas the decreasing schedule progressively releases this constraint. Reducing λ_{base} or its cutoff moves the method toward MultiDiffusion when the prior is poor.

To create the global prior for x_t^{4K} , we resize the input fresco to the model’s native spatial size, generating a small image-to-video sequence. Then, we perform a per-frame trilinear upscale in latent space (to the large canvas size). Fig. 2 illustrates FrescoDiffusion’s generation process, on top of an algorithm in Sec. A.2.

3.3 Spatio-Temporal Prior Strength Scheduling

We will now discuss the design of the prior strength, λ , in Eq. (5). The parameter λ addresses two objectives: (i) Remain structurally close to the prior while allowing the creation of new details. (ii) Treat spatial regions differently in the case of frescoes. Background regions are supposed to remain structurally stable while other active regions should be animated.

To attain objective (i), we use a high λ in the initial diffusion stages, keeping the model close to the prior, and decay it in the later stages to add detail. We thus model λ as a global gated decreasing schedule of the diffusion step

$$\lambda_G(t, \tau) = \lambda_{\text{base}} \cdot \cos\left(t \frac{\pi}{2}\right) \cdot \mathbf{1}[t \leq \tau] \quad (7)$$

where τ is the gating and λ_{base} is the strength of the regularization. When $\lambda = \lambda_G \in \mathbb{R}$, we name our method FrescoDiffusion. This schedule follows a clear principle: stay aligned to the prior while coarse structure forms, then decay to free fine detail. It introduces no free parameters tuned per image; a single fixed $(\lambda_{\text{base}}, \tau)$ is reused across all images and datasets, which we find robust and interpretable. Both Fig. 3 Mean Square Error (MSE vs. t) and the component ablation in Tab. 3 support each design choice.

To reach objective (ii), we compute a spatial activity map $A(p)$ to differentiate active zones from the background. Let $A(p) \in \{0, 1\}$ be a binary map, in which $A(p) = 1$ denotes active regions on the position p (*e.g.*, characters or local scenes expected to move) and $A(p) = 0$ denotes structurally static regions. Also, let τ_{act} and τ_{bg} be two temporal cutoffs, with $\tau_{\text{act}} \leq \tau_{\text{bg}}$, which control the application of the prior to active and background regions, respectively. Hence, our prior strength factor becomes

$$\lambda_R(t, p) = \begin{cases} \lambda_G(t, \tau_{\text{act}}) & \text{if } A(p) = 1 \text{ (pixel } p \text{ in the foreground)} \\ \lambda_G(t, \tau_{\text{bg}}) & \text{if } A(p) = 0 \text{ (pixel } p \text{ in the background)} \end{cases} \quad (8)$$

Please note that λ_R is a tensor with shape $(C \times T \times H^{4K} \times W^{4K})$. We refer to this variant ($\lambda = \lambda_R$) as Regional-FrescoDiffusion (R-FrescoDiffusion).

This design is motivated by Fig. 3, which plots the MSE between the noised prior and the current latents x_t^{4K} . The gate enforces global coherence early ($t \leq \tau_{\text{act}}$), then relaxes the prior first in active regions ($\tau_{\text{act}} < t \leq \tau_{\text{bg}}$) to let motion and detail emerge, while background regions stay constrained longer. In late steps ($t > \tau_{\text{bg}}$), the prior is fully disabled everywhere, and sampling focuses purely on fine detail refinement. The end result is that the background MSE is much closer to the prior compared to the foreground. The coefficient $\lambda_{\text{base}} \geq 0$ controls how strongly we adhere to the prior versus letting tiled denoising add new detail. FrescoDiffusion (λ_G) is our general default. R-FrescoDiffusion helps most when localized actors move against a stable background; with globally coupled motion or no reliable activity map, we use the global variant. When $\tau_{\text{act}} = \tau_{\text{bg}}$, $\lambda_R = \lambda_G$ and the regional variant reduces exactly to FrescoDiffusion. R-FrescoDiffusion does not require a deliberately static prior: the background stability emerges from the longer cutoff τ_{bg} , and a dynamic prior remains compatible.




Fig. 3: Spatial prior schedule λ_R (Eq. (8)). (Top) MSE between the foreground / background regions and the prior: the background stays bound to the prior longer than the foreground. (Bottom) Schedule for both regions.

4 Dataset and Evaluation Protocols

Dataset We use the *Image Suite* of VBench [20], as our first 4K-I2V dataset. Such datasets focus on one or a few objects, and we are looking for frescoes, *i.e.* complex images with multiple intricate scenes, to evaluate our method thoroughly. Therefore, we propose a new dataset named *FrescoArchive* for I2V techniques on a fresco scale. Starting with the LAION-2B Aesthetic Subset [31], we filtered the images based on criteria such as pixel count, aesthetics, watermarks and NSFW scores. Next, we performed text-based filtering, followed by zero-shot classification to detect frescoes. Subsequently, we deduplicated [21] the dataset and generated captions using Qwen3-VL-32B [2] with both the image and the LAION caption. Ultimately, we manually selected 371 pairs to achieve the best possible image-caption match. We provide details on this process in Sec. B, along with statistics. This dataset will be used exclusively for validation.

Evaluation Metrics We employ a user study at full resolution to quantify the human preference over the baselines. We used Amazon Mechanical Turk to conduct the study. Participants were carefully filtered to avoid bots and lower quality evaluators: we required a masters status and a task approval rate superior to 85% to enroll. They were shown pairs of videos generated with two concurrent methods, from the same input image, drawn from both the *FrescoArchive* and VBench-I2V datasets. Videos were displayed at identical resolution and duration, with randomized ordering and no method identification. We report preference percentages with 95% confidence intervals. For each pair, participants answered two binary-choice questions:

- Animation Fidelity: Which video most closely resembles a fresco artwork that has been smoothly and naturally animated?
- Motion plausibility: Which video provides the most convincing animation of the input image, with appropriate and perceptible motion?

We also evaluate our method on the VBench protocols [19, 20, 41], following standard studies [26]. These metrics compute the similarity between the input image and each frame, as well as the similarity between consecutive frames. These protocols evaluate several criteria: Subject Consistency, Motion Smoothness, Aesthetic Score, and Imaging Quality. We complement it using VBench’s I2V metrics (Image-to-Video Subject Consistency and Image-to-Video Background Consistency), to measure the similarity between input image and video.

While VBench is the standard for 480p/1080p, it provides incomplete 4K-I2V evaluations because it downsizes videos to fit metric models’ requirement (DINOv2 [27] and CLIP [29]). We thus complement our testbed with three metrics that specifically target 4K-I2V. (i) We perform standard sharpness measures at full scale to quantify fine-detail generation using the Tenengrad [28] function. (ii) We use a simple yet efficient Temporal Consistency metric, measuring the mean square error between consecutive downsized frames to quantify differences between frames. (iii) We compute a prior similarity metric using DINOv3 [32], which is not limited to 1080p resolution compared to DINOv2 used in VBench.

5 Experiments

5.1 Implementation Details and Baselines

Video generation backbone. All experiments use the 14B-parameter Wan2.2-I2V [36] model. We generate 81 frames at 16 fps with six TurboDiffusion [39] steps and guidance scale 1.0. FrescoDiffusion and MultiDiffusion use 480×832 pixel tiles with 30% overlap. We select $\lambda_{\text{base}} = 1.5$ as a central operating point in the $[0.01, 3]$ sweep in Fig. 7, then freeze it and the τ cutoffs for every image and both datasets. Exact schedules, strides, fusion weights, and resolution handling are reported in Sec. A.1.

Baselines. We compare our method to three tiled-diffusion methods: MultiDiffusion [4], DemoFusion [13] (state-of-the-art tiled image diffusion method adapted to a video setup), and DynamicScaler [26] (state-of-the-art tiled video diffusion method). All implementation details are available in Sec. C.3. For a fair backbone comparison, we adapt every baseline to Wan2.2 and use identical prompts, sampling steps, guidance scales, seeds, output resolutions, and frame counts. Tile geometry is shared where each algorithm permits; method-specific settings follow the authors’ released implementations and are listed in Sec. C.3. We did not perform independent per-baseline hyperparameter searches after transfer to Wan2.2, so the results compare documented default adaptations rather than separately tuned optima. No method is tuned per test image.

Spatial activity maps. R-FrescoDiffusion uses a spatially gated prior schedule to differentiate active regions from structurally static background. For each input image, we compute the activity map $A(p)$ (see Sec. 3.3) using the Segment Anything Model 3 (SAM3) [8], as illustrated in Fig. 4. We apply SAM3 with a fixed set of prompts producing a binary activity map in $[0, 1]$, which is downsampled to latent resolution and used directly in Eq. (8). See details in Sec. C.4. These activity maps are computed once per input image and remain fixed throughout sampling, without additional learnable parameters.



Fig. 4: Overlay of the spatial activity map onto the input fresco.

5.2 Qualitative Evaluation

We begin with a qualitative study. In Fig. 5, we show the first, 40th, and last frame for each model, with fixed-coordinate central and corner crops that expose fine-grained temporal changes at full output resolution. DemoFusion preserves the global structure. Yet, some elements are modified, such as the path in the center crop and the hut in the corner crop. Visible wobbling is present when the video is playing. MultiDiffusion and DynamicScaler tend to introduce excessive




Fig. 5: A qualitative comparison of fresco-scale inputs. FrescoDiffusion generates coherent global scenes and animates details at a local level. By contrast, DemoFusion, DynamicScaler and MultiDiffusion only manage to produce either coherent scenes or high-quality details, but not both.

novel content across tiles or sampling stages, which results in accumulation of structural inconsistencies and loss of temporal coherence. In contrast, FrescoDiffusion produces high quality videos while conserving the general layout of the video. Later, in Sec. 5.4, we discuss the differences between FrescoDiffusion and its regional counterpart from a qualitative perspective. We also contrast with video super-resolution (VSR): methods such as FlashVSR [43] are tied to the trajectory of the base video at its native resolution, so they sharpen texture but cannot introduce the multi-scene content and motion that FrescoDiffusion synthesizes. VSR is therefore not adapted to our task: it can refine an already-determined animation but cannot animate a static fresco. Further full-resolution examples, a VSR comparison, and quantitative motion analysis are reported in Secs. E and F and Fig. 12.

5.3 State-of-the-Art quantitative comparison

High-Resolution Evaluation: User Study. We start with a user study to quantitatively measure human preference, run on both FrescoArchive (FA) and VBench-I2V. Table 1 shows that both of our methods are strongly preferred over DynamicScaler and MultiDiffusion on both datasets, with average preference rates of 76% to 84% on FA and 61% to 69% on VBench-I2V, confirming that these baselines produce noticeably lower quality animations. Against the stronger DemoFusion baseline, the two variants are complementary: on FA, FrescoDiffusion is significantly preferred (62% average), reflecting its more perceptible motion on frescoes, while on VBench-I2V R-FrescoDiffusion wins (62% average) on more structured content. This split mirrors their design: FrescoDiffusion is the general default, and the regional schedule of R-FrescoDiffusion helps

most when foreground and background play distinct roles. Both observations are consistent across the motion and fidelity questions.

Table 1: User study results: human preference rates (% of annotators preferring our method over each baseline) on both FrescoArchive (FA; 371 source images, 1344 ratings) and VBench-I2V (361 source images, 1465 ratings). Each pair is drawn from the same source image, shown at identical resolution and duration in randomized order. Green cells mark a statistically significant preference ($> 50\%$ at the 95% level); all rates use 95% confidence intervals of at most $\pm 6\%$.

	<i>FrescoArchive</i>						<i>VBench-I2V</i>					
	FrescoDiffusion			R-FrescoDiffusion			FrescoDiffusion			R-FrescoDiffusion		
	Motion	Fidelity	Avg.	Motion	Fidelity	Avg.	Motion	Fidelity	Avg.	Motion	Fidelity	Avg.
vs. DemoFusion	63%	60%	62%	54%	62%	58%	55%	46%	51%	66%	57%	62%
vs. DynamicScaler	78%	80%	79%	88%	80%	84%	57%	64%	61%	68%	64%	66%
vs. MultiDiffusion	78%	74%	76%	78%	74%	76%	61%	66%	64%	73%	64%	69%

Standard Low-Resolution I2V Metrics Next, as a sanity check, we compare our approach with modern methods (Tab. 2) using standard lower-resolution I2V metrics (VBench and VBench-I2V). These metrics are not designed for 4K videos as the evaluated videos are downsampled aggressively ($\sim 10 - 20\times$) to fit the backbones’ resolution. As a result, the reported scores provide only a coarse proxy for performance at the original 4K resolution. We perform this evaluation on the FrescoArchive and 4K *Image Suite* VBench dataset, using both regional and non-regional configurations. In the FrescoArchive dataset, our method slightly outperforms the baselines on average. On the *Image Suite* VBench dataset, our approach performs second best on average, slightly outperformed by DemoFusion. Full results for each metric are available in Sec. C.2. Note that for DynamicScaler, our re-implementation performs better, as the I2V backbone in their original implementation, VideoCrafter [12], is much older and clearly underperforms Wan2.2. Our conclusion on this benchmark is consistent with what we observed in our user study and qualitatively: we outperform all methods on the fresco-to-video task and are competitive with DemoFusion on the 4K-I2V task.

Table 2: Standard low-resolution I2V metrics. The table shows the average performance (higher is better) using the VBench evaluation suite on both FrescoArchive and VBench-I2V image sets. We also display the average generation time in minutes.

Method	FrescoArchive	VBench-I2V	Time (min)	Rel. Time
DynamicScaler (original)	0.857	0.862	18.45	$2.27\times$
DynamicScaler* (CVPR’25)	0.871	0.865	10.25	$1.26\times$
DemoFusion* (CVPR’24)	0.903	0.879	13.5	$1.66\times$
MultiDiffusion* (ICML’23)	0.876	0.860	8.15	$1.00\times$
FrescoDiffusion	<u>0.904</u>	0.875	<u>8.58</u>	<u>$1.05\times$</u>
R-FrescoDiffusion	0.907	<u>0.878</u>	9.08	$1.11\times$

Computational efficiency. We report average wall-clock time over all runs on both datasets on a single H100 GPU. FrescoDiffusion and R-FrescoDiffusion remain within 5% and 11% of MultiDiffusion runtime, respectively, while reducing generation time by 36% and 33% relative to DemoFusion (13.5 vs. 8.58/9.08 minutes). Our Wan2.2 DynamicScaler adaptation is faster than the original but remains slower than both variants.

5.4 Controlling FrescoDiffusion

Ablation of the prior strength schedule. We perform an ablation study of the spatio-temporal prior strength schedule’s design, $\lambda(t, p)$, using VBench metrics on FrescoArchive. We start with MultiDiffusion (no λ) and add a constant regularization $\lambda = \lambda_{\text{base}} = 1.5$. This actually worsens performance. Next, we add the cosine schedule $\lambda = \lambda_{\text{base}} \cos(t \frac{\pi}{2}) \in \mathbb{R}$, and obtain significant gains over MD. Then, we build FrescoDiffusion by setting $\lambda = \lambda_G$ (see Eq. (7)), and we finish with R-FrescoDiffusion by setting $\lambda = \lambda_R$ (see Eq. (8)). This results in the best measured performance. Crucially, the global prior-regularized fusion *alone, without any spatial mask* (λ_G) already lifts MultiDiffusion from 0.876 to 0.904 (+0.028 avg), whereas the regional SAM3 mask adds a smaller further +0.003 avg (+0.019 SC). The prior-regularization term, not the mask, therefore drives most of the gain; R-FrescoDiffusion is a complementary control axis, consistent with the user study, where its advantage concentrates on content that requires more creativity. Qualitatively, we show the difference between FrescoDiffusion and R-FrescoDiffusion in Fig. 6. When adding our regional regularization, R-FrescoDiffusion is more similar to the prior on background regions, as intended. We provide further visualization of that effect in Sec. F.

Table 3: Prior strength schedule ablation study on FrescoArchive. Best in **bold**. The results suggest that including the schedule, the gating, and the spatial regularization enhances the quantitative performance.

Method	λ function	SC	MS	A	I	ISC	IBC	Avg
MultiDiffusion	0	0.876	0.974	0.686	0.754	0.981	0.989	0.876
FrescoDiffusion	λ_{base}	0.942	0.991	0.645	0.598	0.979	0.983	0.856
FrescoDiffusion	$\lambda_{\text{base}} \cos(t \frac{\pi}{2})$	0.946	0.991	0.724	0.730	0.987	0.992	0.895
FrescoDiffusion	λ_G (Eq. (7))	0.958	0.990	0.738	0.753	0.991	0.995	0.904
R-FrescoDiffusion	λ_R (Eq. (8))	0.977	0.991	0.736	0.753	0.991	0.994	0.907

λ -Controlled Pareto Trade-off Between Creativity and Prior Similarity Creativity and prior similarity are essentially contrary objectives. One cannot improve one without hurting the other. This inherent tradeoff creates a Pareto frontier composed of the set of optimal compromises between the two objectives. To navigate this frontier using $\lambda_G(t, \tau)$ (Eq. (7)), we linearly modify the prior strength ($\lambda_{\text{base}} \in [0.01, 3]$), and the temporal gating ($\tau \in [0, 1]$). To




Fig. 6: Regional constraint. Prior shows an overlay of the activity map. The red/blue boxes represent background/foreground regions. Our regional loss pulls the background toward the prior while allowing new details in foreground regions.

represent the creativity objective we use the sharpness metric as a proxy, and for the prior similarity we use both temporal consistency and our prior similarity metric presented in Sec. 4. The results in Fig. 7 show two expected behaviors. (i) When λ_{base} and the temporal gating, τ , increase, the outputs equal those of the prior. (ii) On the contrary, when those parameters decrease, we reach the same performance as MultiDiffusion (full creativity, no prior). The curve formed between these two opposites creates the Pareto frontier which allows a trade-off between the two objectives. Thus, FrescoDiffusion allows full control over this crucial trade-off.




Fig. 7: Quantitative evaluation of the trade-off between creativity and prior similarity controlled by λ . (a) Temporal consistency versus sharpness, illustrating the Pareto frontier between preserving temporal coherence and maintaining high image sharpness. (b) Evolution of prior similarity over time, showing how increasing prior strength and temporal gating progressively aligns the generated outputs with the prior.

6 Limitations and Conclusion

Limitations. FrescoDiffusion relies on a meaningful low-resolution prior; an inaccurate thumbnail can propagate erroneous structure or motion in proportion to λ , although the decreasing schedule limits this effect. Content demanding globally coupled geometric motion (sustained camera motion or live sports) also falls outside its current scope. Our experiments use only Wan2.2; while Sec. G gives the ϵ -prediction extension, empirical cross-backbone validation remains future work. Finally, tiled denoising remains computationally expensive despite reduced-step sampling and low-precision optimizations.

Conclusion. FrescoDiffusion is a simple, training-free solution that uses existing video diffusion models to animate 4K, multi-scene images. It combines tiled denoising with a latent prior derived from a thumbnail animation to preserve global coherence while introducing local detail at large scales. In user studies, our variants are preferred over MultiDiffusion and DynamicScaler on both datasets, preferred over DemoFusion on FrescoArchive, and competitive with it on VBench-I2V, while using less wall-clock time than the stronger baselines. Our approach also exposes an interpretable balance between creativity and fidelity for large-scale image animation.

Acknowledgements

This work has been partially supported by the VISA DEEP chair (grant ANR-20-CHIA-0022), the PostGenAI@Paris cluster (grant ANR-23-IACL-0007; France 2030), and funded by the French National Research Agency (ANR) under the Renaissance project (grant ANR-23-CE23-0023; France 2030).

Computing and storage resources were provided by GENCI at IDRIS under grant 2025-AD011016538 on the Jean Zay H100 partition.

A CC BY 4.0 public copyright license has been applied by the authors to the present document and will be applied to all subsequent versions up to the Author Accepted Manuscript (AAM) arising from this submission, in accordance with the grant’s open access conditions.

References

1. Avrahami, O., Fried, O., Lischinski, D.: Blended latent diffusion. ACM TOG **42**(4), 1–11 (2023)
2. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., et al.: Qwen3-vl technical report. arXiv preprint arXiv:2511.21631 (2025)
3. Bar-Tal, O., Chefer, H., Tov, O., Herrmann, C., Paiss, R., Zada, S., Ephrat, A., Hur, J., Liu, G., Raj, A., et al.: Lumiere: A space-time diffusion model for video generation. In: SIGGRAPH Asia 2024 Conference Papers. pp. 1–11 (2024)
4. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: MultiDiffusion: Fusing diffusion paths for controlled image generation. In: ICML. vol. 202, pp. 1737–1752 (2023)

5. Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., Kulal, S., et al.: Flux. 1 kontekst: Flow matching for in-context image generation and editing in latent space. arXiv e-prints pp. arXiv-2506 (2025)
6. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
7. Bolya, D., Huang, P.Y., Sun, P., Cho, J.H., Madotto, A., Wei, C., Ma, T., Zhi, J., Rajasegaran, J., Rasheed, H.A., et al.: Perception encoder: The best visual embeddings are not at the output of the network. In: NeurIPS (2025)
8. Carion, N., Gustafson, L., Hu, Y.T., Debnath, S., Hu, R., Suris, D., Ryali, C., Alwala, K.V., Khedr, H., Huang, A., et al.: Sam 3: Segment anything with concepts. arXiv preprint arXiv:2511.16719 (2025)
9. Chan, K.C., Wang, X., Yu, K., Dong, C., Loy, C.C.: Basicvsr: The search for essential components in video super-resolution and beyond. In: CVPR. pp. 4947–4956 (2021)
10. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Basicvsr++: Improving video super-resolution with enhanced propagation and alignment. In: CVPR. pp. 5972–5981 (2022)
11. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Investigating tradeoffs in real-world video super-resolution. In: CVPR. pp. 5962–5971 (2022)
12. Chen, H., Xia, M., He, Y., Zhang, Y., Cun, X., Yang, S., Xing, J., Liu, Y., Chen, Q., Wang, X., et al.: Videocrafter1: Open diffusion models for high-quality video generation. arXiv preprint arXiv:2310.19512 (2023)
13. Du, R., Chang, D., Hospedales, T., Song, Y.Z., Ma, Z.: Demofusion: Democratising high-resolution image generation with no \$\$\$\$. In: CVPR. pp. 6159–6168 (2024)
14. Frolov, S., Moser, B.B., Dengel, A.: Spotdiffusion: A fast approach for seamless panorama generation over time. In: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 2073–2081 (2025)
15. HaCohen, Y., Chiprut, N., Brazowski, B., Shalem, D., Moshe, D., Richardson, E., Levin, E., Shiran, G., Zabari, N., Gordon, O., et al.: Ltx-video: Realtime video latent diffusion. arXiv preprint arXiv:2501.00103 (2024)
16. He, J., Xue, T., Liu, D., Lin, X., Gao, P., Lin, D., Qiao, Y., Ouyang, W., Liu, Z.: Venhancer: Generative space-time enhancement for video generation. arXiv preprint arXiv:2407.07667 (2024)
17. He, Y., Yang, S., Chen, H., Cun, X., Xia, M., Zhang, Y., Wang, X., He, R., Chen, Q., Shan, Y.: Scalecrafter: Tuning-free higher-resolution visual generation with diffusion models. In: ICLR (2023)
18. Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D.P., Poole, B., Norouzi, M., Fleet, D.J., et al.: Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303 (2022)
19. Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., Jiang, Y., Zhang, Y., Wu, T., Jin, Q., Chanpaisit, N., Wang, Y., Chen, X., Wang, L., Lin, D., Qiao, Y., Liu, Z.: VBench: Comprehensive benchmark suite for video generative models. In: CVPR (2024)
20. Huang, Z., Zhang, F., Xu, X., He, Y., Yu, J., Dong, Z., Ma, Q., Chanpaisit, N., Si, C., Jiang, Y., Wang, Y., Chen, X., Chen, Y.C., Wang, L., Lin, D., Qiao, Y., Liu, Z.: VBench++: Comprehensive and versatile benchmark suite for video generative models. IEEE TPAMI (2025). <https://doi.org/10.1109/TPAMI.2025.3633890>
21. Jain, T., Lennan, C., John, Z., Tran, D.: Imagededup. <https://github.com/idealo/imagededup> (2019)

22. Jiménez, Á.B.: Mixture of diffusers for scene composition and high resolution image generation. arXiv preprint arXiv:2302.02412 (2023)
23. Li, W., Ren, J., Zhang, P., Ye, T., Zhou, D., Luan, J., Zhu, L.: PixelWizard: Towards efficient high-fidelity video generation at ultra-large spatial resolution. arXiv preprint arXiv:2605.25801 (2026)
24. Liang, J., Cao, J., Fan, Y., Zhang, K., Ranjan, R., Li, Y., Timofte, R., Van Gool, L.: Vrt: A video restoration transformer. *IEEE TIP* **33**, 2171–2182 (2024)
25. Liu, H., Ruan, Z., Zhao, P., Dong, C., Shang, F., Liu, Y., Yang, L., Timofte, R.: Video super-resolution based on deep learning: a comprehensive survey. *Artificial Intelligence Review* **55**(8), 5981–6035 (2022)
26. Liu, J., Lin, S., Li, Y., Yang, M.H.: Dynamicscaler: Seamless and scalable video generation for panoramic scenes. In: *CVPR*. pp. 6144–6153 (2025)
27. Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Howes, R., Huang, P.Y., Xu, H., Sharma, V., Li, S.W., Galuba, W., Rabbat, M., Assran, M., Ballas, N., Synnaeve, G., Misra, I., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision (2023)
28. Pertuz, S., Puig, D., Garcia, M.A.: Analysis of focus measure operators for shape-from-focus. *PR* **46**(5), 1415–1432 (2013)
29. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *ICML*. pp. 8748–8763 (2021)
30. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *CVPR*. pp. 10684–10695 (2022)
31. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in neural information processing systems* **35**, 25278–25294 (2022)
32. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. arXiv preprint arXiv:2508.10104 (2025)
33. Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., Parikh, D., Gupta, S., Taigman, Y.: Make-a-video: Text-to-video generation without text-video data. In: *ICLR* (2023)
34. Team, G.: Mochi 1. <https://github.com/genmoai/models> (2024)
35. Teed, Z., Deng, J.: RAFT: Recurrent all-pairs field transforms for optical flow. In: *European conference on computer vision*. pp. 402–419. Springer (2020)
36. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025)
37. Wang, W., Gao, Z., Gu, L., Pu, H., Cui, L., Wei, X., Liu, Z., Jing, L., Ye, S., Shao, J., et al.: Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. arXiv preprint arXiv:2508.18265 (2025)
38. Xie, R., Liu, Y., Zhou, P., Zhao, C., Zhou, J., Zhang, K., Zhang, Z., Yang, J., Yang, Z., Tai, Y.: Star: Spatial-temporal augmentation with text-to-video models for real-world video super-resolution. In: *ICCV*. pp. 17108–17118 (2025)
39. Zhang, J., Zheng, K., Jiang, K., Wang, H., Stoica, I., Gonzalez, J.E., Chen, J., Zhu, J.: Turbodiffusion: Accelerating video diffusion models by 100-200 times (2025)
40. Zhang, Q., Song, J., Huang, X., Chen, Y., Liu, M.Y.: Diffcollage: Parallel generation of large content with diffusion models. In: *CVPR*. pp. 10188–10198 (2023)

41. Zheng, D., Huang, Z., Liu, H., Zou, K., He, Y., Zhang, F., Zhang, Y., He, J., Zheng, W.S., Qiao, Y., Liu, Z.: VBench-2.0: Advancing video generation benchmark suite for intrinsic faithfulness. arXiv preprint arXiv:2503.21755 (2025)
42. Zhou, S., Yang, P., Wang, J., Luo, Y., Loy, C.C.: Upscale-a-video: Temporal-consistent diffusion model for real-world video super-resolution. In: CVPR. pp. 2535–2545 (2024)
43. Zhuang, J., Guo, S., Cai, X., Li, X., Liu, Y., Yuan, C., Xue, T.: FlashVSR: Towards real-time diffusion-based streaming video super-resolution. arXiv preprint arXiv:2510.12747 (2025)

FrescoDiffusion: 4K Image-to-Video with Prior-Regularized Tiled Diffusion - Supplementary Material

A FrescoDiffusion Additional Details

A.1 Implementation Details

All experiments use the Wan2.2-I2V 14B backbone with the same accelerated 6-step TurboDiffusion setting as in the main paper. We generate 81 frames at 16 fps with guidance scale 1.0, i.e., without effective classifier-free guidance. For an input image of size (H, W) , we first generate the low-resolution prior at fixed target area $480 \times 832 = 399,360$ pixels while preserving aspect ratio:

$$\hat{H} = \text{round}\left(\sqrt{399360 H/W}\right), \quad \hat{W} = \text{round}\left(\sqrt{399360 W/H}\right),$$

and then snap both dimensions down to the nearest multiple of 16. The full-resolution pass uses $H^{4K} = \max(16, \lfloor H/16 \rfloor \cdot 16)$ and $W^{4K} = \max(16, \lfloor W/16 \rfloor \cdot 16)$, after an optional isotropic downscaling when the input exceeds ultra-HD resolution. This multiple-of-16 constraint comes from the latent video backbone: the Wan VAE reduces spatial resolution by a factor 8, and the downstream latent grid is processed in spatial patches of size 2, yielding an effective validity constraint of $8 \times 2 = 16$. With spatial and temporal compression factors 8 and 4, the latent tensor therefore has shape $16 \times 21 \times (H^{4K}/8) \times (W^{4K}/8)$ for the default 81-frame setting. The low-resolution latent prior is resized to the large latent canvas with endpoint-aligned trilinear interpolation in latent space, VAE tiling is enabled during inference, and the decoded output is resized back to the original image size only if snapping changed the resolution.

Tiled denoising and regularization. The high-resolution pass uses MultiDiffusion windows of size 480×832 pixels with 30% overlap, giving nominal pixel strides 336×582 . After conversion to latent coordinates and rounding to the valid latent grid, this becomes 60×104 latent windows with strides 42×72 ; extra final windows are added whenever needed so that the right and bottom boundaries are exactly covered. Tile fusion uses linear ramps with minimum border weight 0.1. Standard MultiDiffusion uses $\sum_i w_i y_i / \sum_i w_i$, whereas the prior-regularized implementation additionally accumulates $\sum_i w_i y_i$ and $\sum_i w_i$ for the one-shot closed-form update in model-output space. In all reported FrescoDiffusion runs we use a cosine prior schedule with $\lambda_{\text{base}} = 1.5$ and cutoff $\tau_{\text{end}} = 0.1$, where $\tau = i/(N-1)$ and $N = 6$ (number of steps); hence the prior is active only at the first denoising step. In R-FrescoDiffusion, the active regions cutoff is $\tau_{\text{act}} = 0.1$ and the inactive regions cutoff is $\tau_{\text{bg}} = 0.35$. Since the six normalized step positions are $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$, the foreground prior is active only at $i = 0$, while the background prior is active at $i = 0$ and $i = 1$. The active or inactive parts of the video are determined based on masks computed according to Sec. C.4.

A.2 Sampling Procedure

At each step, we run the transformer on each crop of the canvas latent x_0^{4K} to obtain a per-tile prediction y_i . Then, we accumulate two canvas-shaped tensors: $\sum_i w_i \odot y_i$ and $\sum_i w_i$. We then add the prior term and divide as in (6) to produce a single fused prediction on the full canvas. Finally, we invoke the scheduler exactly once to obtain $x_{t+\Delta t}$. This preserves the original sampler while adding only light overhead: one extra reduction per pixel, a single pointwise rational fuse, and no additional network passes beyond those already required by tiled denoising. The full step is summarized in Algorithm 1.

Algorithm 1: FrescoDiffusion: one sampler step at time t

Input: canvas latent $x_t^{4K} \in \mathbb{R}^{C \times T \times H^{4K} \times W^{4K}}$; tile positions $\{p_i\}_{i=1}^n$; weight maps $\{w_i\}_{i=1}^n$; upscaled prior $x_{\text{prior}} \in \mathbb{R}^{C \times T \times H^{4K} \times W^{4K}}$; noise level σ_t ; prior-strength schedule λ ; flow-matching step size Δt ; flow-matching model f_θ

Output: updated canvas latent $x_{t+\Delta t}^{4K}$

- 1 Initialize $num \leftarrow 0 \in \mathbb{R}^{C \times T \times H^{4K} \times W^{4K}}$, $den \leftarrow 0 \in \mathbb{R}^{1 \times 1 \times H^{4K} \times W^{4K}}$
- 2 **for** $i = 1, \dots, n$ **do**
- 3 $\tilde{x}_t \leftarrow \mathcal{C}_{p_i}(x_t^{4K})$ // Crop canvas at p_i
- 4 $\tilde{y}_i \leftarrow f_\theta(\tilde{x}_t, t, c)$ // Estimates crop flow
- 5 $y_i \leftarrow \mathcal{P}_{p_i}(\tilde{y}_i)$ // Zero-pads the output
- 6 $num \leftarrow num + w_i \odot y_i$ // Updates the numerator
- 7 $den \leftarrow den + w_i$ // Updates denominator
- 8 // Add prior regularization

$$y \leftarrow \frac{num + \lambda \sigma_t (x_t^{4K} - x_{\text{prior}})}{den + \lambda \sigma_t^2}$$

- 9 $x_{t+\Delta t}^{4K} \leftarrow \text{SCHEDULERSTEP}(x_t^{4K}, y, t)$ // Update noisy states

B Fresco Evaluation Dataset Construction

To facilitate the creation of UHD-I2V, we noticed that no dataset fits our requirements for UHD-I2V. Unlike VBench high-definition [20] set or panoramas [26] that focus on one object or less, we search for images to animate with multiple intricate scenes. Here, we propose a new set to generate and evaluate UHD-I2V techniques at a fresco-like scale. We name our dataset FrescoArchive. We start with the LAION-2B Aesthetic Subset [31]. The first step is then filtering the images based on several criteria: having more than a million pixels, an aesthetic score of at least 5.8, and a watermark and unsafe scores of less than 0.8 and 0.5, respectively. This initial filtering process yielded a total of two million images. Next, we performed a semantic filtering process. To do this, for each image, we compute the average cosine similarity between the target instance and the prompts

"A large detailed fresco"
 "A magnificent fresco with many different scenes"
 "A narrative composition"
 "A fresco with lots of details"
 "A large polyptych and composite image fresco"
 "A large metapicture with several compositions"
 "A fresco tableau"
 "A painting fresco"

using the PerceptionEncoder [7] G14-448 variant similarity model, and we finish by selecting the top 50,000 images. After, we used Intern-VL-3.5 [37] as a classifier to detect frescoes. To do so, we used the following prompt:

You are a visual classifier. Decide whether the image is a fresco-like composition.
 Definition (for this task): A qualifying image resembles a large, detailed, integrated scene (like historical frescoes). Modern photos or digital works count if they share these traits.
 Answer yes only if all are true:

- The image shows high apparent resolution / detail density (many fine, precise elements).
- There are multiple distinct sub-scenes or groups (from a few to dozens+) distributed across the same frame.
- These elements are blended into one coherent composition (no panel borders or obvious collage seams).

Answer no if any of the following:

- Single subject or minimal detail.
- The fresco is not the main content of the image (e.g. the photo shows a wall, room, or museum scene where the fresco only appears as a small part, rather than the fresco itself being the full image).
- Simple graphics, logos, posters, or text-only images.
- Comics/manga with separate panels, tiled grids, or collages with hard borders.
- Diagrams, charts, UI/screenshots, or patterns.

Unsure: answer no.
 Output format: Respond with exactly yes or no (lowercase, no punctuation, no extra words).

This filtering results in a total of 10,000 images. Finally, we deduplicated the dataset using both perceptual hashing and CNN-based deduplication techniques using the ImageDedup library [21]. Then, to generate UHD image captions, we used Qwen3-VL-32B [2] with both the image and LAION caption, resulting in 6,734 UHD image-caption pairs. To prompt Qwen3-VL-32B, we used the following text:

Using the existing caption below as context, write a long, highly detailed, precise, and fluent caption that thoroughly describes the image. Give some precise contextual information, relative positional information, subjects, objects and elements description and identification. Caution: the caption provided can be false or wrong; use the image as the only source of truth. The caption is only here to help you be more precise. Respond with the caption only (no preface, no metadata, no quotes). Existing caption: {caption}

Finally, we manually selected 371 pairs to get the best qualitative image-captions pairs.

Dataset	Samples	Words	Image Width	Image Height
FrescoArchive	371	355.79 ± 88.0	2265 ± 1257	1552 ± 864
VBench	361	14.70 ± 2.24	4592 ± 1305	3748 ± 1214

Table 4: Dataset statistics. We compare FrescoArchive and VBench datasets’ low-level statistics.

For the statistics, we compute several metrics (text and image-wise) and quantitatively compare with VBench I2V [20] to mark a reference point. First, we compute high-level statistics, seen in Tab. 4. We measure the number of samples, average number of words per prompt, and average width and height. As we can see, our set contains a similar number of images as VBench. Yet, our prompts contain an order of magnitude more than VBench, depicting more precise and detailed prompts. As for the average image shape, our set contains smaller images, but with more complex scenery. In Fig. 8, you can see the number of words and shape distribution of our dataset.




Fig. 8: Analysis of FrescoArchive dataset. Left: distribution of word in caption. Right: Resolution of images distribution.

Next, we qualitatively study FrescoArchive’s complexity with reference to VBench. To do this, we first encode all images using the DINOv2 [27] model to get a global view of each image. Then, we perform a PCA dimensionality

reduction to visualize their distribution. Finally, using the resulting reduction, we further visualize individual crops. As can be seen in Fig. 9, the resulting PCA shows that our dataset covers a wider span of the main axes, unlike VBench. This suggests that our dataset has more diversity in shared components using the global DINO features.




Fig. 9: PCA between FrescoArchive and VBench I2V dataset. We visualize the DINOv2 [27] feature PCA between our proposed dataset and VBench’s images (red and green points, respectively). Also, in light green and red, random crops of the images. This plot shows that our dataset contains images with different statistics than standard ones, found in the VBench I2V set.

C Experiment Details

C.1 Metrics

Sharpness metric. We assess per-frame sharpness using the Tenengrad [28] measure, a classical no-reference focus metric based on directional gradient energy. For each frame, the luminance channel is extracted and horizontal and vertical gradients are computed using a 3×3 Sobel operator. The Tenengrad score is defined as the mean squared gradient magnitude $T = \frac{1}{HW} \sum_{i,j} (G_x^2(i,j) + G_y^2(i,j))$, where G_x and G_y are the responses of the Sobel filters along each axis. This quantity is sensitive to high-frequency spatial structure and penalizes blurry outputs where gradient energy is suppressed. Per-frame scores are averaged over all frames of a video to yield a single video-level score. Higher values indicate sharper, more detail-preserving outputs.

Temporal consistency metric. Each frame is first converted to grayscale and then downsampled to 128×128 pixels using area-averaging interpolation, retaining only coarse spatial structure. The temporal inconsistency score for a video of T frames is defined as the mean squared error between consecutive downsampled frames: $\mathcal{C} = \frac{1}{T-1} \sum_{t=1}^{T-1} |\hat{f}_t - \hat{f}_{t-1}|_F^2 / (64^2)$, where $\hat{f}_t$ denotes the

downscaled grayscale frame at time t . By operating at this coarse resolution, the metric captures global consistency while remaining agnostic to legitimate scene motion.

Prior alignment metric. We measure how faithfully each generated video preserves the global semantic content of the prior using a frame-level feature alignment score based on DINOv3 [32]. For each frame pair $(f_t^{\text{prior}}, f_t^{\text{gen}})$, both frames are independently forwarded through a frozen DINOv3 ViT-S/16 encoder without any resizing (*i.e.*, at native resolution), and we extract the CLS token from each: the pooled global representation output by the transformer. Prior alignment is then defined as the mean cosine similarity between corresponding CLS tokens across all T frames of a video: $\mathcal{A} = \frac{1}{T} \sum_{t=1}^T \frac{\langle \mathbf{z}_t^{\text{prior}}, \mathbf{z}_t^{\text{gen}} \rangle}{\|\mathbf{z}_t^{\text{prior}}\| \|\mathbf{z}_t^{\text{gen}}\|}$, which equivalently measures the cosine of the angle between the two CLS token directions in feature space. We report the mean of this score across all videos in the benchmark; higher values indicate that the generated video remains semantically aligned to the prior along its temporal trajectory. A key advantage of DINOv3 over its predecessor DINOv2 is its ability to process high-resolution inputs, including 4K frames, without interpolating positional embeddings or resizing the input. This property is essential in our setting, where the generated videos are high-resolution and downscaling prior to feature extraction would discard fine-grained content that may be critical for alignment assessment.

C.2 Standard low-resolution I2V metrics

On Table 2 in the main paper, we presented the average results on both FrescoArchive and VBench-I2V datasets. Here, in Table 5, we detail each metric: Subject Consistency, Motion Smoothness, Aesthetic, Imaging, I2V Subject Consistency, and I2V Background Consistency.

C.3 Baseline Implementation Details

MultiDiffusion. Our implementation follows the original MultiDiffusion [4] procedure without additional heuristics. The only modifications are (i) replacing the base backbone with Wan2.2 I2V in place of the original denoiser, and (ii) applying a linear decay blending mask outside each tile to smoothly attenuate contributions near tile borders and reduce seam artifacts when merging overlapping predictions.

DemoFusion. DemoFusion [13] introduces three techniques for high-resolution image generation. We re-implemented (i) progressive phase upsampling and (ii) skip-residual global guidance, reusing the exact hyper-parameters from the authors’ official code to enable a faithful comparison. These parameters control how resolution increases across phases (number of phases, per-phase upsampling factor, and per-scale denoising schedule) and the strength of global-structure guidance during refinement. In contrast, we did not observe reliable gains from

Table 5: Quantitative evaluation on FrescoArchive and VBench-I2V. Higher is better for all metrics except time. Best in **bold** and second best underlined. * denotes methods adapted to the video prior setting. SC, MS, A, I, ISC, and IBC stand for Subject Consistency, Motion Smoothness, Aesthetic, Imaging, I2V Subject Consistency, and I2V Background Consistency, respectively.

Method	Quality Metrics				I2V Metrics		Avg	Time
	SC	MS	A	I	ISC	IBC		
FrescoArchive								
DynamicScaler (original)	0.945	0.975	0.693	0.706	0.893	0.930	0.857	18.45
DynamicScaler* (CVPR’25)	0.852	0.971	0.681	0.754	0.980	0.989	0.871	10.25
DemoFusion* (CVPR’24)	<u>0.960</u>	0.987	0.734	0.752	<u>0.990</u>	<u>0.994</u>	0.903	13.5
MultiDiffusion* (ICML’23)	0.876	0.974	0.686	0.754	0.981	0.989	0.876	8.15
FrescoDiffusion	0.958	<u>0.990</u>	0.738	<u>0.753</u>	0.991	0.995	<u>0.904</u>	<u>8.58</u>
R-FrescoDiffusion	0.977	0.991	<u>0.736</u>	<u>0.753</u>	0.991	<u>0.994</u>	0.907	9.08
VBench-I2V								
DynamicScaler (original)	0.949	0.976	0.707	0.713	0.893	0.932	0.862	18.45
DynamicScaler* (CVPR’25)	0.904	<u>0.985</u>	0.621	0.701	0.988	0.991	0.865	10.25
DemoFusion* (CVPR’24)	0.943	0.989	0.639	0.720	0.990	0.993	0.879	13.5
MultiDiffusion* (ICML’23)	0.893	0.984	0.611	0.698	0.987	0.989	0.860	8.15
FrescoDiffusion	<u>0.933</u>	0.989	0.632	<u>0.716</u>	<u>0.989</u>	<u>0.992</u>	0.875	<u>8.58</u>
R-FrescoDiffusion	0.956	0.989	<u>0.634</u>	0.712	0.988	0.991	<u>0.878</u>	9.08

DemoFusion-style dilated sampling when transferring it from SDXL’s UNet denoiser to Wan2.2’s DiT-based denoiser. We hypothesize this is because dilated sampling assumes updates are roughly separable across interleaved sub-lattices an assumption that fits UNets’ local, convolutional structure but breaks for DiT models with global self-attention. Evaluating Wan2.2 on sparse lattices changes the attention context and likely shifts inputs off-distribution, leading to inconsistent offsets that merge into visible artifacts (seams/checkerboards) rather than improved coherence.

DynamicScaler. We follow the authors’ official implementation of DynamicScaler [26] and reuse their released hyper-parameters for both the offset-shifting denoiser and the global motion-guidance module. We condition motion guidance on the same low-resolution video that we use for our FrescoDiffusion regularization prior, so both signals rely on an identical motion reference. For the sliding/rotating denoising window, we set the per-step offset (stride) to half the window size, *i.e.* a 50% overlap between consecutive windows, which stabilizes stitching across steps and mitigates boundary artifacts.

C.4 Spatial Activity Map Computation

Given an input frame, we compute a *spatial activity map* using SAM3 [8]. The region to animate can also be explicitly specified by the user. For automated

processing and ease of use, we employ a segmentation pipeline that identifies plausible dynamic entities and converts them into a spatial activity map.

The activity map is computed once per image, stored at latent resolution, clamped to $[0, 1]$, resized to the current latent size with standard trilinear interpolation, binarized using the test $A > 0$, and then kept fixed throughout sampling.

Prompt-based segmentation. Our default pipeline queries SAM3 using a fixed set of prompts corresponding to categories that commonly exhibit motion. Specifically, we provide the following textual prompts to the model:

- person
- vegetation
- vehicles

For each prompt, SAM3 predicts candidate spatial masks corresponding to instances of the queried category. The masks are then averaged. Visualizations of such masks are provided in Fig. 10.

Mask extraction and filtering. SAM3 predictions are filtered using a score threshold $\tau_s = 0.45$. We discard masks whose relative area exceeds 0.30 of the image area, as such regions typically correspond to overly coarse detections. We additionally remove masks with excessive boundary contact, defined as cases where more than 80% of the mask pixels lie within a 10-pixel margin of the image border.

Spatial support. For each retained mask, we construct a spatial support region by dilating the mask using a Euclidean distance transform with radius $r = 75$ pixels. The resulting regions are merged to produce the final spatial activity map.

Exploratory variants. We explored two extensions of this pipeline. First, we experimented with extending the spatial mask across time. Using SAM3 together with the generated prior video, the mask predicted on the input frame was propagated to cover the full temporal extent of the video, as seen in Fig. 11. Second, we evaluated a variant in which the prompts provided to SAM3 are generated automatically using a vision-language model. In this setup, Qwen3-VL-32B [2] analyzes both the input image and the generated prior video to produce textual prompts corresponding to entities that could plausibly support animation. These prompts are then used to condition SAM3, while mask extraction, filtering, and spatial support construction remain identical to the default pipeline. In practice, neither the temporal mask propagation nor the VLM-guided prompting produced measurable improvements over the fixed-prompt approach. As both variants introduce additional computational overhead, they were not used for the final results reported in the paper.



Fig. 10: Additional qualitative examples of spatial activity maps obtained with SAM3. Each image shows the overlay used to identify regions likely to contain dynamic content, which are then used to guide the active/inactive prior regularization in R-FrescoDiffusion.




Fig. 11: Two qualitative examples showing the temporal masks obtained with SAM3 at frames 1, 40, and 80. The video overlaid is the prior generated with Wan model.

D VSR Baseline Qualitative Comparison

To illustrate the VSR discussion in Sec. 2, Fig. 12 compares, at $t = 2s$ on the same FrescoArchive input, the low-resolution prior (Wan2.2-I2V at native resolution), FlashVSR applied to that prior, and FrescoDiffusion. The bottom row zooms on the lower-left quadrant of each frame. FlashVSR sharpens the texture of the prior but reproduces its layout almost exactly, including the same static figures and the empty foreground. FrescoDiffusion instead populates the same region with new content (e.g., the horses and additional figures), confirming that VSR can refine an already-determined trajectory but cannot invent content beyond it.




Fig. 12: VSR baseline qualitative comparison. Top row: full frame for the low-resolution prior, FlashVSR, and FrescoDiffusion. Bottom row: zoom on the lower-left quadrant of each frame above. FlashVSR stays close to the prior’s layout, while FrescoDiffusion introduces new content (e.g., the horses) absent from both the prior and the VSR output.

E Motion Quantity at Full Resolution

We measure optical-flow magnitude using RAFT-Large [35] on frames resized to 2560×1440 . For each method, we report the median flow magnitude after removing outliers above $Q_3 + 1.5 \text{ IQR}$, quantifying how much net motion each method introduces at full resolution.

Table 6: Optical flow magnitude (RAFT-Large, median over pixels, $Q_3 + 1.5 \text{ IQR}$ outliers removed) at 2560×1440 on FrescoArchive (FA) and VBench-I2V.

Method	FA	VBench-I2V
Wan (native resolution)	0.288	5.837
FlashVSR	0.247	6.421
DemoFusion	1.251	4.638
DynamicScaler	24.085	17.358
MultiDiffusion	26.217	18.469
FrescoDiffusion	1.318	4.424
R-FrescoDiffusion	0.885	4.374

Table 6 reveals three regimes. FlashVSR and native Wan produce near-zero flow on FA, below 0.3 px: VSR refines texture but inherits the motion of its low-resolution prior. MultiDiffusion and DynamicScaler produce an order of magnitude more flow, reaching 26 and 24 px on FA, yet score lowest on motion preference in the user study (Tab. 1), because the large flow reflects tiles drifting independently rather than coherent animation. FrescoDiffusion and R-FrescoDiffusion occupy the intermediate regime, with 1.3 and 0.9 px respectively, producing motion that users consistently prefer. Motion coherence, not motion quantity, drives preference. For the coherent-motion methods, namely native Wan, FlashVSR, DemoFusion, FrescoDiffusion, and R-FrescoDiffusion, flow magnitudes are lower on FrescoArchive than on VBench-I2V. This gap stems from the different content of the two image sets rather than from less motion. FrescoArchive frescoes spread many small subjects across a large composition, so each animated element is small relative to the frame and its displacement spans only a few pixels. VBench-I2V images instead center on one or a few large subjects that occupy much of the frame, so an equivalent proportional motion produces a larger pixel displacement. The two drift baselines, MultiDiffusion and DynamicScaler, are the exception, with higher flow on FrescoArchive: their motion is dominated by tiles drifting in uncorrelated directions, an effect amplified by the many independent scenes of a fresco rather than by genuine subject motion.

F FrescoDiffusion Additional Examples



Fig. 13: Additional examples obtained on our FrescoArchive dataset. First 2 rows obtained with FrescoDiffusion and last 2 rows obtained with R-FrescoDiffusion; columns show the same frame indices.




Fig. 14: Additional examples obtained on VBench I2V dataset. First 3 rows obtained with FrescoDiffusion and 2 last rows obtained with R-FrescoDiffusion; columns show the same frame indices.

G FrescoDiffusion closed-form solution in noise prediction setting

Our proposed approach can be used with ϵ -prediction diffusion models. We modify the FrescoDiffusion loss in Eq. (5) to include the one-step approximation of the ϵ -diffusion formulation:

$$\ell_{\text{FD}}(y^*; t) = \left\| \sqrt{\lambda} \odot \left[\frac{1}{\sqrt{\alpha_t}} (x_t^{4K} - \sqrt{1 - \alpha_t} y^*) - x_{\text{prior}} \right] \right\|_2^2 + \ell_{\text{MD}}(y^*; t), \quad (9)$$

where $\alpha_t = \prod_{i=1}^t (1 - \beta_i)$ and β_t are the schedule variances, and the one-step approximation is given by $\frac{1}{\sqrt{\alpha_t}} (x_t^{4K} - \sqrt{1 - \alpha_t} y^*)$. Next, we set the derivative of Eq. (9) to 0 to solve for the optimal noise output. Hence, the optimal ϵ noise is:

$$y_{\text{FD}}(x_t^{4K}) = \frac{\sqrt{\frac{1 - \alpha_t}{\alpha_t}} \lambda \odot \left(\frac{1}{\sqrt{\alpha_t}} x_t^{4K} - x_{\text{prior}} \right) + \sum_{i=1}^n w_i \odot y_i}{\frac{1 - \alpha_t}{\alpha_t} \lambda + \sum_{i=1}^n w_i}. \quad (10)$$

We adopt the same variable definitions, as in the main text, for the current noisy state x_t^{4K} , the prior, x_{prior} , the weighting tensors w_i , and the prior regularization strength λ . As in the flow-matching formulation, when $\lambda = 0$, y_{FD} reduces to y_{MD} .