

Geometrically Consistent Multi-View Scene Generation from Freehand Sketches

Ahmed Bourouis^{1,2}, Savas Ozkan¹, Andrea Maracani¹, Yi-Zhe Song²,
and Mete Ozay¹

¹Samsung Research, UK ²SketchX, University of Surrey, UK




Fig. 1: Sketch-to-Multi-View generation. Given a single freehand sketch and a text caption, our method generates geometrically consistent multi-view images spanning a full 360° azimuth orbit at four different elevations. Top three rows: in-domain examples from FS-COCO [9]. Bottom two rows: zero-shot generalisation to unseen sketch domains (TU-Berlin [11] and InstantStyle [41]). [Project Page](#).

Abstract. We tackle a new problem: generating geometrically consistent multi-view scenes from a single freehand sketch. Freehand sketches are the most geometrically impoverished input one could offer a multi-view generator. They convey scene intent through abstract strokes while introducing spatial distortions that actively conflict with any consistent 3D interpretation. No prior method attempts this; existing multi-view approaches require photographs or text, while sketch-to-3D methods need multiple views or costly per-scene optimisation.

We address three compounding challenges; absent training data, the need for geometric reasoning from distorted 2D input, and cross-view consistency, through three mutually reinforcing contributions: (i) a curated dataset of $\sim 9k$ sketch-to-multiview samples, constructed via an automated generation and filtering pipeline; (ii) Parallel Camera-Aware Attention Adapters (CA3) that inject geometric inductive biases into the video transformer; and (iii) a Sparse Correspondence Supervision Loss (CSL) derived from Structure-from-Motion reconstructions.

Our framework synthesizes all views in a single denoising process without requiring reference images, iterative refinement, or per-scene optimization. Our approach significantly outperforms state-of-the-art two-stage baselines, improving realism (FID) by over 60% and geometric consistency (Corr-Acc) by 23%, while providing up to a $3.7\times$ inference speedup.

Keywords: Multi-view synthesis · Sketch-based generation · Geometric consistency · Correspondence supervision

1 Introduction

Consider a rough, freehand sketch of a street scene: a few strokes for a car, some wobbly lines representing buildings, a stick figure on the pavement. Despite this inherent abstraction, a human observer effortlessly infers a plausible 3D world, understanding that the car has an occluded side, the buildings extend beyond the visible corners and the figure could be seen from different perspectives. This rises the fundamental question:

Can a generative model, conditioned solely on a single freehand sketch, synthesize a set of geometrically consistent novel views representative of a complete 3D scene?

We define this task as **Sketch-to-Multi-View (S2MV)** generation as illustrated in Fig. 1. This task is a surprisingly hard problem for current generative models, and one that has not been studied before. Freehand scene sketches are arguably the most geometrically impoverished conditioning signal one could provide to multiview generators. Unlike photographs, which provide dense appearance and implicit 3D cues (i.e., shading and perspective), or text prompts, which offer semantic flexibility without spatial constraints, freehand sketches are *simultaneously* spatially informative yet geometrically unreliable. They convey the *intent* of a scene layout while introducing abstraction, non-linear distortions, and stroke-level noise that often conflict with a consistent 3D interpretation.

Existing multi-view diffusion methods do not directly address these inconsistencies. Image-conditioned approaches [12, 28, 30, 38, 39, 42, 55] typically begin with high-fidelity photographs that already resolve geometric ambiguity. Text-conditioned methods [39] operate in a domain where spatial imprecision is expected and tolerated. Conversely, existing Sketch-to-3D methods, such as Sketch2NeRF [7] and SketchDream [27], are largely restricted to single-object reconstruction or require *multiple* sketch views and costly per-scene optimisation. To our knowledge, no prior work generates geometrically consistent multi-view outputs from a single, freehand scene sketch.

Bridging this gap requires solving three compounding challenges. **(i) Data:** no sketch-to-multi-view dataset exists, and freehand sketches cannot be consistently paired with multi-view images through standard rendering pipelines. **(ii) Geometric reasoning:** generative models must infer coherent 3D structures from a 2D input that lacks strict projective alignment, demanding camera-aware generation mechanisms. **(iii) Cross-view consistency:** generating different views simultaneously requires the model to enforce correspondences *between* views, rather than simply conditioning on viewpoint independently.

We address these challenges with the following contributions:

- **A Curation Pipeline for S2MV Data** (Sec. 3.2). We introduce a novel pipeline to bridge the domain gap between freehand sketches and multi-view images. By leveraging FS-COCO [9], we synthesize photorealistic frontal views using a pretrained generative model [21], verify sketch-image alignment through semantic segmentation matching [4, 6], and generate multiple views

- per scene [10]. The resulting dataset includes 9 222 curated samples with 33 views each, providing a rigorous foundation for S2MV training.
- **Camera-Aware Attention Adapters (CA3)** (Sec. 3.3). We propose a lightweight camera attention adapters that inject Projective Rotary Position Encoding (PRoPE) [25] into a pretrained video diffusion transformer [43]. This mechanism adds only 2.7% additional parameters while enabling attention layers to explicitly account for relative projective geometry between views.
 - **Correspondance Supervision Loss (CSL)** (Sec. 3.4). We introduce a sparse InfoNCE-based correspondence loss applied to the adapter’s latent projections to maintain geometric consistency between cross-views. Using positive pairs extracted through Structure-from-Motion [37], this regularization enforces cross-view identity and merges distinct viewpoints into a unified latent representation.

The components of our framework are mutually reinforcing, creating a cohesive system for S2MV synthesis. While our curated dataset provides the foundational sketch-multiview pairs previously missing from the literature, our CA3 provides the necessary geometric inductive bias and the *capacity* to reason about viewpoints. Last but not least, our CSL provides the explicit supervisory signal that teaches the model *which* spatial locations across views should attend to each other, turning geometric capacity into geometric consistency. Removing any one of these components significantly degrades the final performance.

Crucially, our method generates all multi-views jointly in a *single end-to-end denoising process* directly from a sketch, bypassing any intermediate photo-generation stage, requiring no reference photograph, no per-scene optimisation, no iterative refinement. To the best of our knowledge, this is the first method to produce geometrically coherent, photorealistic multi-view scenes from freehand sketch input.

2 Related Work

Multi-View Generation and Camera Conditioning. Diffusion models have been widely applied to multi-view image generation. Early methods fine-tune image diffusion models to synthesise novel views from a single photograph: Zero-1-to-3 [28] conditions on relative camera pose, and follow-up works [24, 30, 31, 38, 39] improve consistency by generating multiple views jointly through synchronised diffusion or multi-view attention. A parallel line of work repurposes video diffusion models, treating the temporal axis as a viewpoint axis, for object-centric orbiting [8, 42] or scene-level [12] generation. More recent DiT-based architectures push the scale further: Hunyuan3D 2.0 [55] trains a 2B-parameter shape-generation DiT, TRELLIS [48] employs rectified flow transformers over structured 3D latents, and See3D [33] learns 3D priors from 16M video clips without pose annotations.

How camera geometry is communicated to the generative model is a critical design axis. Absolute conditioning methods [13, 28, 47] inject camera parameters

as global vectors, while per-pixel Plücker ray representations [13, 49, 56] offer richer spatial encoding but do not inject *relative* inter-view geometry into the attention computation. Geometric approaches address this limitation directly: PRoPE [25] modulates queries, keys, and values with projective matrices so that attention scores reflect relative camera geometry, UCPE [51] extends this to per-token ray-level encoding, and EpiDiff [16] and SPAD [17] restrict cross-view attention to epipolar lines. AC3D [1] shows that camera information concentrates in a subset of DiT layers, enabling a $4\times$ parameter reduction, and the concurrent ReRoPE [23] injects relative pose into underutilised RoPE frequency bands.

We build on the Wan 2.1 video DiT [43] and adopt PRoPE inside lightweight parallel adapters (adding fewer than 3% parameters) following the zero-initialised parallel branch paradigm [52]. We complement the geometric inductive bias with explicit correspondence supervision, a combination not explored in prior work.

Geometric Consistency and Correspondence Supervision. Ensuring cross-view geometric consistency in generated images remains a central challenge. Implicit approaches encode geometric priors through architecture: SyncDreamer [30] constructs a shared 3D feature volume, epipolar attention [16, 24, 40] constrains cross-view attention to geometrically plausible regions, and SpatialCrafter [54] decouples spatial and temporal attention for scene reconstruction. ViewCrafter [50] guides video diffusion with monocular depth-based point clouds, and GeoVideo [2] aligns predicted depth maps across frames during training. In all these cases, geometric reasoning is an emergent property of the architecture rather than an explicitly supervised objective.

In the reconstruction domain, DUS3R [45], MAST3R [22], and VGGT [44] learn to predict dense correspondences from multi-view images, but treat them as the end product rather than as supervision for generation. Correspondent-Dream [18] extracts correspondences from a frozen diffusion UNet to regularise NeRF optimisation, yet does not train the generative model itself with correspondence supervision.

The concurrent CAMEO [20] is closest to our work, supervising attention maps with VGGT-derived correspondences via dense cross-entropy on the main self-attention of a UNet-based model. Beyond the apparent “sparse-vs-dense” contrast, our approach differs along four *causal* axes. (i) Adapter vs. main attention: CAMEO supervises a self-attention layer of CAT3D’s [12] UNet, entangling correspondence with appearance on the *same* parameters that solve denoising; we supervise a *parallel*, zero-initialised adapter whose PRoPE-modulated queries/keys live in the relative-camera projective space, so geometric supervision never competes with appearance synthesis. (ii) Sparse contrastive vs. dense cross-entropy at video scale: CAMEO’s dense map is tractable on $F=4$ views at 32×32 ($\approx 16\text{M}$ entries/layer); at our video-DiT scale ($4N+1=133$ frames, 60×60 latent) it would hold $\approx 2.3\times 10^{11}$ entries (hundreds of GB/layer), whereas our sparse InfoNCE [36] costs $\sim 1\text{ MB/layer}$ and explicitly repels negatives. (iii) Classical vs. learned correspondence: our SfM [37] supervision carries no learned prior, whereas CAMEO inherits VGGT’s pointmap biases. (iv) Video DiT vs. image backbone: CAMEO is validated on image backbones at $F=4$, while we




Fig. 2: Method overview. (a) A freehand sketch is encoded then denoised by a video DiT augmented with our **CA3**. All N views are generated in a single forward denoising process. (b) Each DiT block contains a self-attention with LoRa and a parallel CA3 to inject relative camera geometry. (c) During training, SfM correspondences from pseudo ground-truth views supervise the CA3 query-key projections via a sparse InfoNCE loss. Dashed boxes: frozen modules.

show the supervision transfers to a video DiT generating $4N+1$ frames in a *single* pass, the regime where dense supervision is infeasible.

Sketch-Conditioned Generation. Sketch-based image generation has been widely studied through conditional diffusion models. ControlNet [52] and T2I-Adapter [34] accept edge maps as spatial conditions but target clean line drawings rather than freehand sketches. Recent work addresses the sketch abstraction gap: Koley *et al.* [19] handle amateur sketches for single-image generation, and KnobGen [35] adapts to varying sketch complexity. SketchingReality [3] generates photorealistic images from freehand *scene* sketches using a CLIP-based sketch encoder, but is limited to single-view output. In the 3D domain, Sketch2NeRF [7] requires multiple sketch views and optimises a NeRF via SDS, while SketchDream [27] combines sketch conditioning with depth guidance, both target single objects and require costly per-scene optimisation.

Our method extends the scene-sketch paradigm of [3] to *multi-view* output: from a single freehand scene sketch, we generate N geometrically consistent views through a video diffusion transformer in a single forward pass, with no per-scene optimisation.

3 Method

Given a freehand sketch S and a set of N target camera poses, our method generates N geometrically consistent photorealistic views through a video diffusion

transformer (Fig. 2.a) augmented with lightweight camera-aware adapters **CA3** (Fig. 2.b) and trained with correspondence supervision **CSL** (Fig. 2.c).

S2MV requires solving three compounding challenges: the absence of paired training data, the need for camera-aware generation, and the enforcement of cross-view geometric consistency. In this section, we first review the video diffusion backbone (Sec. 3.1), followed by our dataset construction pipeline (Sec. 3.2), Camera-Aware Attention Adapters (CA3) (Sec. 3.3), and Correspondence Supervision Loss (CSL) (Sec. 3.4).

3.1 Preliminaries: Video Diffusion Transformers

Our framework is built upon a standard Video Diffusion Transformer (Wan 2.1 [43]), which is repurposed to support the multi-view sequence modeling required for S2MV. The model operates in a compressed latent space $\mathcal{Z}$ learned by a VAE encoder $\mathcal{E}$, which first maps a sequence of frames $[\mathbf{x}_i]_{i=0}^N$ into a latent representation $[\mathbf{z}_i]_{i=0}^N$. Later, the DiT architecture processes each $\mathbf{z}_i$ by patchifying it into a sequence of T tokens whose dimension is D . Each of the L transformer blocks computes self-attention as:

$$\text{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}} \right) \mathbf{V}, \quad (1)$$

where $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ are the query, key, and value projections. The training is performed via flow matching [26] by minimizing the velocity field error:

$$\mathcal{L}_{\text{flow}} = \mathbb{E}_{t, \mathbf{z}_0, \epsilon} [\|\mathbf{v}_\theta(\mathbf{z}_t, t, \mathbf{c}) - (\mathbf{z}_0 - \epsilon)\|^2], \quad (2)$$

where $\mathbf{z}_t$ is the noisy latent at timestep $t \in [0, 1]$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and $\mathbf{c}$ denotes the conditioning signals. The VAE applies $4\times$ temporal compression, effective for video where consecutive frames are redundant, but destructive for multi-view generation where neighbouring views differ substantially. We therefore replicate each view four times before encoding (*frame replication*), yielding $4N+1$ input frames and exactly one latent per view.

3.2 Sketch-to-Multiview Dataset Construction

No existing dataset pairs freehand scene sketches with geometrically consistent multi-view images. Unlike photographs, freehand sketches contain non-linear distortions that do not correspond to any standard 3D asset. We therefore propose a pipeline to bridge this domain gap through an automated generation-and-filtering process. The steps in our pipeline are summarized in Fig. 3.

- (a) **Multi-Seed Generation.** For each of the $\sim 10,000$ sketches in FS-COCO [9], we synthesize five photorealistic candidate images using FLUX.2 dev [21], conditioned on both the sketch and its original caption.




Fig. 3: S2WV dataset generation pipeline. (a) Multi-seed image generation from sketch and caption. (b) Segmentation of sketch and generated images. (c) Best seed selection via mIoU. (d) Multi-view synthesis from the selected image. (e) Dataset curation yields 9,222 samples. More details in Sec. 3.2

- (b) **Segmentation.** To ensure spatial fidelity, we perform semantic grounding by segmenting both the input sketch and the candidates. Sketch masks are extracted via a specialized sketch encoder [4]. For the generated images, we use spaCy to extract object nouns from the caption, which serve as prompts for zero-shot segmentation via GroundingDINO [29] and SAM3 [6].
- (c) **Best Seed Selection.** For each sketch, we select the candidate image I^* that best aligns with the sketch’s intended spatial layout. This is determined by maximizing the mean Intersection-over-Union (mIoU) between the sketch segmentation masks $\mathcal{M}_s$ and the image segmentation masks $\mathcal{M}_g$:

$$I^* = \arg \max_{i \in \{1..5\}} \text{mIoU}(\mathcal{M}_s, \mathcal{M}_{g,i}). \quad (3)$$

- (d) **Multi-View Generation.** The selected frontal view is processed by Qwen Image Edit Angles [10] to synthesize N multi-view images at varying azimuth θ and elevation ϕ .
- (e) **Dataset Curation.** We apply a minimum threshold on the mIoU to remove samples where the subjects are poorly aligned.

The result of our pipeline is a final corpus of **9 222 high-quality samples**.

3.3 Camera-Aware Attention Adapters (CA3)

A video diffusion model has no intrinsic notion of camera viewpoint as its self-attention treats all tokens identically regardless of their 3D spatial relationship. For photograph-conditioned generation, dense appearance cues in the input can partially compensate; for freehand sketches, which provide only sparse,

geometrically distorted strokes, the model receives no implicit 3D signal whatsoever. We therefore equip the model with an explicit geometric inductive bias through lightweight Camera-Aware Attention Adapters (CA3). Below we describe the construction of virtual camera matrices, the projective position encoding that consumes them, the adapter architecture, and a complementary domain-adaptation strategy.

Virtual Camera Matrix Construction. Since our multi-view images are generated via image editing rather than 3D rendering, ground-truth camera parameters are unavailable. We construct virtual cameras by placing each viewpoint on a sphere of fixed radius d around the scene origin. For a view at azimuth θ and elevation ϕ , we compose a rotation $\mathbf{R}^{(n)} = \mathbf{R}_y(\theta) \cdot \mathbf{R}_x(\phi)$ and set the camera position to $\mathbf{t}^{(n)} = \mathbf{R}^{(n)}[0, 0, d]^\top$. The camera-to-world extrinsic is then:

$$\mathbf{T}^{(n)} = \begin{bmatrix} \mathbf{R}^{(n)} & \mathbf{t}^{(n)} \\ \mathbf{0}^\top & 1 \end{bmatrix} \in SE(3). \quad (4)$$

The intrinsic matrix $\mathbf{K}$ is shared across all views, constructed from a 60° horizontal field-of-view with the principal point at the image centre. These virtual matrices provide CA3 with a consistent geometric coordinate system, enabling viewpoint reasoning in the absence of ground-truth calibration

Viewpoint Awareness via P_{Ro}PE. We integrate Projective Rotary Position Encoding (P_{Ro}PE) [25], which modulates attention using the scene’s epipolar geometry rather than pixel coordinates alone. For each target view n , we construct a projective matrix from $\mathbf{T}^{(n)}$ and $\mathbf{K}$:

$$\mathbf{P}^{(n)} = \begin{bmatrix} \mathbf{K} & \mathbf{0} \\ \mathbf{0}^\top & 1 \end{bmatrix} \cdot (\mathbf{T}^{(n)})^{-1}. \quad (5)$$

By partitioning the attention head dimensions into projective and spatial subspaces, P_{Ro}PE ensures that the affinity between any two tokens is conditioned on their relative projective transform $\mathbf{P}^{(n_q)}\mathbf{P}^{(n_k)^{-1}}$, embedding epipolar constraints directly into the attention computation.

CA3 Architecture. To preserve the powerful generative priors of the pretrained backbone, we implement CA3 as a parallel branch alongside the attention layers as illustrated in Fig. 2-(b). For each transformer block ℓ , the hidden state $\mathbf{h}_\ell$ is updated as:

$$\mathbf{h}'_\ell = \mathbf{h}_\ell + \mathcal{A}_\ell^{\text{self}}(\mathbf{h}_\ell) + \mathcal{A}_\ell^{\text{cam}}(\mathbf{h}_\ell, \{\mathbf{P}^{(n)}\}_{n=1}^N), \quad (6)$$

where $\mathcal{A}_\ell^{\text{self}}(\cdot)$ is LoRA-adapted backbone self-attention and $\mathcal{A}_\ell^{\text{cam}}(\cdot, \cdot)$ is our trainable camera adapter that uses the projective matrices of the views. To maintain computational efficiency, each adapter employs an $8\times$ dimensionality bottleneck ($1536 \rightarrow 192$) and utilizes only $H'=2$ attention heads, and its output projection is zero-initialised [52] to preserve the pretrained prior at the start of training. The adapter architecture above addresses *where* to attend across views; however, the backbone features themselves must also be adapted to the sketch domain.

Sketch-Domain Adaptation via LoRA. The video backbone is pretrained exclusively on natural videos; freehand sketches represent a significant domain




Fig. 4: SfM correspondences across views. Each correspondence image shows two views with colored lines connecting matched keypoints identified by Structure-from-Motion (COLMAP). The view angles (azimuth, elevation) for each pair are indicated above the correspondence visualizations. These correspondences serve as supervision targets for our camera-aware attention adapters (Sec. 3.4).

shift that the camera adapters alone cannot bridge. We therefore apply LoRA [15] with rank 16 to the query, key, value, and output projections of the backbone’s self-attention layers. While CA3 provides explicit geometric viewpoint control via PRoPE, LoRA handles the complementary task of adapting the backbone’s feature representations for sketch-conditioned generation. Training follows a curriculum that lets domain adaptation stabilise before geometric supervision is introduced (Sec. 3.4). We validate this design choice in Sec. 4.4.

3.4 Correspondence Supervision Loss (CSL)

The camera adapters give the model the *capacity* to distinguish viewpoints, but this geometric inductive bias alone does not teach the model *which* spatial locations across views correspond to the same 3D point. Without explicit supervision, the adapter’s attention remains diffuse and spatially unstructured, as we verify empirically in Sec. 4.5. We therefore introduce CSL derived from SfM that directly teaches cross-view geometric reasoning (Fig. 4).

Correspondence Extraction. Given the pseudo ground-truth multi-view images for each training sample, we run COLMAP [37] to obtain sparse 3D reconstructions. For each reconstructed 3D point $\mathbf{X}_j$, COLMAP provides its projection coordinates $(u_j^{(n)}, v_j^{(n)})$ in each observing view n , along with a confidence score $c_j^{(n)}$ derived from reprojection error. We construct pairwise supervision targets:

$$\mathcal{C} = \{(q_i, k_i, w_i)\}_{i=1}^M,$$

where M is the total number of correspondence pairs, $q_i = (n_q, s_q)$ and $k_i = (n_k, s_k)$ specify the (view index, spatial patch index) for each pair and $w_i = \min(c_j^{(n_q)}, c_j^{(n_k)})$ is the confidence weight.

Sparse InfoNCE loss. Directly supervising the full attention matrix $\mathbf{A}^{(\ell)} \in \mathbb{R}^{T \times T}$ is memory-prohibitive ($T > 29,000$). We instead formulate correspondence

supervision as a contrastive objective on the adapter’s query and key projections. For each correspondence pair (q_i, k_i) , we extract query and key embeddings $\mathbf{q}_i$ and $\mathbf{k}_i$ from the CA3 module at a subset of layers $\mathcal{U} = \{10, 15, 20, 25\}$, corresponding to middle-to-deep layers where geometric reasoning is most pronounced. The sparse InfoNCE loss is:

$$\mathcal{L}_{\text{corr}} = -\frac{1}{M} \sum_{i=1}^M w_i \cdot \log \frac{\exp(\mathbf{q}_i^\top \mathbf{k}_i^+ / \tau)}{\exp(\mathbf{q}_i^\top \mathbf{k}_i^+ / \tau) + \sum_{j=1}^{N_{\text{neg}}} \exp(\mathbf{q}_i^\top \mathbf{k}_j^- / \tau)},$$

where $\mathbf{k}_i^+$ is the positive (corresponding) key, $\{\mathbf{k}_j^-\}_{j=1}^{N_{\text{neg}}}$ are negative keys sampled from non-corresponding token positions, and τ is a temperature parameter. This formulation has complexity $O(M \times N_{\text{neg}})$ rather than $O(T^2)$, reducing memory from 14 GB to ~ 1 MB.

Total training objective. The final training objective combines the flow matching loss with correspondence supervision:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{flow}} + \lambda_{\text{corr}} \cdot \mathcal{L}_{\text{corr}},$$

where λ_{corr} follows a curriculum schedule: it is held at zero during an initial warmup phase, then linearly ramped to its target value of 0.01. This allows the flow matching loss and LoRA domain adaptation to stabilise before correspondence supervision is introduced.

4 Experiments

We evaluate our method on sketch-to-multiview generation, assessing both per-view image quality and cross-view geometric consistency. We describe the experimental setup (Sec. 4.1), compare against two strong baselines (Sec. 4.2), present ablations (Sec. 4.4), visualise learned attention (Sec. 4.5), and test generalisation to unseen sketch domains (Sec. 4.6).

4.1 Experimental Setup

Dataset and splits. We train on our S2MV dataset described in Sec. 3.2 with a total of 9,222 sketch-multiview samples ($N = 33$ views each). We split these into 8,304 training, 441 validation, and 477 test samples, following [5]. All reported metrics are computed on the test set across 3 seeds.

Evaluation metrics. We report standard per-view image quality metrics against ground-truth views: peak signal-to-noise ratio (**PSNR** $\uparrow$), structural similarity index (**SSIM** $\uparrow$) [46], and learned perceptual image patch similarity (**LPIPS** $\downarrow$) [53] using an AlexNet backbone. To measure distributional realism, we report **FID** $\downarrow$ [14] computed over all generated views versus all ground-truth views. For semantic fidelity, we compute **CLIP-I** $\uparrow$ (cosine similarity between CLIP ViT-B/32 embeddings of generated and GT views). We additionally report **Corr-Acc** $\uparrow$, a geometric consistency metric that measures the percentage of COLMAP correspondences whose reprojection error on generated views falls below a threshold

Table 1: Quantitative comparison on S2MV test set (477 samples, $N=33$ views each, 3 seeds). Our single-stage method is compared against two-stage baselines that first translate the sketch to a photograph via FLUX.2-dev.

Method	#Stages	Time/sample	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FID $\downarrow$	CLIP-I $\uparrow$	Corr-Acc $\uparrow$
SEVA [56]	2-stage	~ 3.1 min	11.310	0.265	0.705	46.34	0.756	0.161
ViewCrafter [50]	2-stage	~ 35 min	11.148	0.338	0.737	48.22	0.773	0.136
Ours	1-stage	~ 50 s	12.169	0.302	0.632	18.49	0.828	0.199

of 5 pixels; directly evaluating whether the model preserves the 3D structure visible in ground truth. All per-view metrics are averaged across views, seeds and test samples.

Baselines. No prior method generates geometrically consistent multi-view output from a single freehand scene sketch. Therefore, we evaluate two state-of-the-art novel-view synthesis methods adapted to our setting via a shared two-stage pipeline: FLUX.2-dev [21] first translates each sketch to a photorealistic front view, then either **SEVA** [56] (Plücker-conditioned latent diffusion, 576×576) or **ViewCrafter** [50] (point-cloud-guided video diffusion via DUS3R [45], 576×1024) generates the remaining views. Full baseline configurations are provided in the supplementary material.

Implementation details. We set $N=33$ views at 480×480 resolution, building on the Wan 2.1 video DiT (1.3B parameters) [43] with frame replication ($N \rightarrow 4N+1$). We train on $8 \times$ NVIDIA A100 GPUs with AdamW [32], LoRA [15] rank 16 (~ 5.9 M parameters) alongside ~ 35.4 M CA3 parameters. The correspondence loss follows a curriculum schedule, ramping λ_{corr} from 0 to 0.01. Extended training details are provided in the supplementary material.

4.2 Comparison with Baselines

Quantitative results. Tab. 1 presents the main comparison across all metrics. Our single-stage method outperforms both two-stage baselines on five of six metrics despite operating directly from sketches without an intermediate photorealistic reference. The improvement is most pronounced on realism (FID 18.49 vs. 46–48) and geometric consistency (Corr-Acc 0.199 vs. 0.161/0.136), confirming that CA3 and CSL produce views that are individually plausible and mutually coherent. Perceptual quality likewise favours our method (LPIPS 0.632 vs. 0.705/0.737), as does semantic alignment with ground truth (CLIP-I 0.828 vs. 0.756/0.773). ViewCrafter achieves the highest SSIM (0.338), which we attribute to its tendency to produce smoother, lower-frequency outputs that score well on pixel-level structural similarity despite worse perceptual fidelity.

Our single-stage design also provides a $42\times$ inference speedup over ViewCrafter and $3.7\times$ over SEVA. More details in the supplementary material.

Qualitative comparison. Fig. 5 presents visual comparisons for representative test sketches. We visualize more examples in the supplementary material.




Fig. 5: Qualitative comparison of multi-view generation from sketches. For each example we show the input sketch with its text prompt (left) and eight azimuth views (0°–315° at 45° intervals) generated by each method. **GT**: ground-truth views; **Ours**: our single-pass sketch-to-multiview method; **VC**: ViewCrafter; **SEVA**: Stable Video Diffusion. Additional comparisons are provided in the supplementary material.

Table 2: Two-stage FLUX+Qwen vs. ours on S2MV. “Frontal” evaluates the single front view (mIoU = sketch fidelity); “Multi-view” evaluates all N views. Ours wins on both stages despite seeing only a sketch at inference.

Method	Frontal				Multi-view ($N=33$)		
	FID↓	CLIP-I↑	LPIPS↓	mIoU↑	FID↓	CLIP-I↑	Corr-Acc↑
FLUX+Qwen (3 seeds)	47.64	0.776	0.731	0.338	51.77	0.804	0.174
Ours	36.33	0.781	0.722	0.353	19.12	0.822	0.191

Table 3: Ablation study. We remove each contribution independently. Row (a) is the full model. Removing any single component degrades both per-view quality and geometric consistency, confirming that the contributions are mutually reinforcing.

Configuration	PSNR↑	SSIM↑	LPIPS↓	FID↓	CLIP-I↑	Corr-Acc↑	T. params
(a) Full model (CA3 + CSL + LoRA)	12.169	0.302	0.632	18.49	0.828	0.199	~41.3M
(b) w/o CSL ($\lambda_{\text{corr}}=0$)	12.073	0.287	0.664	20.37	0.817	0.175	~41.3M
(c) w/o CA3 (LoRA only)	5.026	0.266	0.819	266.06	0.632	0.183	~5.9M
(d) w/o LoRA (CA3 only)	12.211	0.304	0.644	19.27	0.823	0.188	~35.4M
(e) w/o frame replication	12.198	0.304	0.652	42.54	0.786	0.174	~41.3M

4.3 Comparison to a Two-Stage FLUX+Qwen Pipeline

Our data pipeline (Sec. 3.2) decomposes sketch-to-multiview into two off-the-shelf stages: FLUX.2-dev [21] (sketch→frontal) followed by Qwen Image Edit Angles [10] (frontal→ N views). To test whether our trained single-stage model is merely a faster substitute for this pipeline, we evaluate a strong two-stage **FLUX+Qwen** variant on the S2MV test set (Tab. 2) averaged across 3 different seeds.

Despite seeing strictly less information at inference, a sketch rather than a curated photograph, our model wins on *every* metric, even surpassing the pipeline on *frontal* sketch fidelity (mIoU 0.353 vs. 0.338) and realism (FID 36.33 vs. 47.64); the gap widens for multi-view. The bottleneck is thus the sketch→frontal step, which our five-seed mIoU curation (Sec. 3.2) closes during training. The LoRA branch is therefore *not* a speed-only convenience: it absorbs the curated S2MV statistics, so inference needs no FLUX seed search. The same argument applies to any sketch→image→multi-view decomposition (*e.g.* a sketch-ControlNet front-end), of which FLUX+Qwen is a strong public instance.

4.4 Ablation Study

We ablate each component under identical training settings and report results on the S2MV test set averaged across 3 seeds (Tab. 3).

- **CSL** (Tab. 3b). Removing the correspondence loss ($\lambda_{\text{corr}}=0$) retains geometric capacity but removes explicit supervision. Row (b) shows the largest




Fig. 6: Attention correspondence visualisation. Query pixel (red dot) in the front view; heatmaps show attention at layer 20 over three target viewpoints. Left: with correspondence supervision. Right: without (ablation Tab. 3 row b).

Corr-Acc drop among non-catastrophic ablations (0.175 vs. 0.199) and degraded LPIPS (0.664 vs. 0.632), confirming that **CSL** is essential for translating adapter capacity into cross-view consistency.

- **CA3** (Tab. 3c). Without our CA3 camera adapter, the model has no mechanism to distinguish viewpoints. Row (c) is nearly catastrophic: PSNR drops to 5.026, FID degrades to 266.06, and CLIP-I drops to 0.632, confirming that LoRA alone is insufficient for the multi-view setting.
- **LoRA** (Tab. 3d). Row (d) is competitive on per-view quality (PSNR 12.211, FID 19.27), suggesting our CA3 carry most of the representational burden. However, Corr-Acc drops to 0.188, indicating that LoRA’s domain adaptation complements the adapters for fine-grained geometric consistency.
- **Frame replication** (Tab. 3e). Without frame replication, the VAE’s $4\times$ temporal compression merges neighbouring views into shared latent frames. Per-view PSNR remains competitive (12.198), but FID more than doubles (42.54 vs. 18.49) and Corr-Acc degrades to 0.174, confirming that preserving view identity through the VAE bottleneck is critical.

Taken together, all components are mutually reinforcing: removing any one degrades geometric consistency.

4.5 Geometric Consistency Evaluation

To understand *how* correspondence supervision shapes attention, we visualise camera attention maps at layer 20 (Fig. 6). Given a query pixel in the front view, we extract attention weights over three target views at varying viewpoints. With correspondence supervision, attention concentrates sharply on the geometrically correct region, even for the challenging 180° back view. Without it (ablation row b), attention is diffuse and spatially unstructured, confirming that CSL directly teaches the adapters where to attend.




Fig. 7: Generalisation to unseen sketch domains. 12 views per sketch at four elevations. (Top) TU-Berlin object sketches [11]: consistent multi-view output despite a distinctly different sketch style. (Bottom) InkScenes [41]: robust to dense, textured scene compositions unseen during training.

4.6 Generalisation to Unseen Sketch Domains

Our model is trained exclusively on FS-COCO freehand sketches. To test whether the learned geometric reasoning transfers to unseen sketch distributions or not, we evaluate on sketches from two out-of-distribution datasets (see the supplementary material) and the qualitative results are illustrated in Fig. 7. From these results, it is clear that per-view image quality and consistency between multi-views are plausible, even if the data distributions are different than FS-COCO.

4.7 Limitations

Our method is limited by (i) training scale (8 304 samples vs. $> 1\text{M}$ for baselines), (ii) generated rather than rendered ground truth, and (iii) 480×480 resolution; we discuss these in the supplementary material.

5 Conclusion

We presented the first method to generate geometrically consistent multi-view scenes from a single freehand sketch in one forward pass, unifying a curated S2MV dataset, lightweight parallel CA3, and a sparse CSL into three mutually reinforcing components. Beyond multi-view synthesis, we believe this opens a broader direction: sketching as a lightweight, accessible interface for 3D content creation. More generally, our work demonstrates that combining camera-aware attention adapters with explicit correspondence supervision is an effective paradigm for lifting 2D generative models into geometrically grounded multi-view generators.

References

1. Bahmani, S., Skorokhodov, I., Qian, G., Siarohin, A., Menapace, W., Tagliasacchi, A., Lindell, D.B., Tulyakov, S.: Ac3d: Analyzing and improving 3d camera control in video diffusion transformers. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 22875–22889 (2025)
2. Bai, Y., Fang, S., Yu, C., Wang, F., Huang, Q.: Geovideo: Introducing geometric regularization into video generation model. arXiv preprint arXiv:2512.03453 (2025)
3. Bourouis, A., Bessmeltsev, M., Gryaditskaya, Y.: Sketchingreality: From freehand scene sketches to photorealistic images. arXiv preprint arXiv:2602.14648 (2026)
4. Bourouis, A., Fan, J.E., Gryaditskaya, Y.: Open vocabulary semantic scene sketch understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4176–4186 (2024)
5. Bourouis, A., Fan, J.E., Gryaditskaya, Y.: Data used in the paper" open vocabulary scene sketch semantic understanding"(fscoco-seg) (2024)
6. 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)
7. Chen, M., Yuan, W., Wang, Y., Sheng, Z., He, Y., Dong, Z., Bo, L., Guo, Y.: Sketch2nerf: multi-view sketch-guided text-to-3d generation. arXiv preprint arXiv:2401.14257 (2024)
8. Chen, Z., Wang, Y., Wang, F., Wang, Z., Liu, H.: V3d: Video diffusion models are effective 3d generators. arXiv preprint arXiv:2403.06738 (2024)
9. Chowdhury, P.N., Sain, A., Bhunia, A.K., Xiang, T., Gryaditskaya, Y., Song, Y.Z.: Fs-coco: Towards understanding of freehand sketches of common objects in context. In: European conference on computer vision. pp. 253–270. Springer (2022)
10. dx8152: Qwen-Edit-2509-Multiple-angles. Hugging Face model card (2025), <https://huggingface.co/dx8152/Qwen-Edit-2509-Multiple-angles>, accessed: 2026-01-28
11. Eitz, M., Hays, J., Alexa, M.: How do humans sketch objects? ACM Transactions on graphics (TOG) **31**(4), 1–10 (2012)
12. Gao, R., Holynski, A., Henzler, P., Brussee, A., Martin-Brualla, R., Srinivasan, P., Barron, J.T., Poole, B.: Cat3d: Create anything in 3d with multi-view diffusion models. arXiv preprint arXiv:2405.10314 (2024)
13. He, H., Xu, Y., Guo, Y., Wetzstein, G., Dai, B., Li, H., Yang, C.: Cameractrl: Enabling camera control for text-to-video generation. arXiv preprint arXiv:2404.02101 (2024)
14. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
15. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. ICLR **1**(2), 3 (2022)
16. Huang, Z., Wen, H., Dong, J., Wang, Y., Li, Y., Chen, X., Cao, Y.P., Liang, D., Qiao, Y., Dai, B., et al.: Epidiff: Enhancing multi-view synthesis via localized epipolar-constrained diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9784–9794 (2024)
17. Kant, Y., Siarohin, A., Wu, Z., Vasilkovsky, M., Qian, G., Ren, J., Guler, R.A., Ghanem, B., Tulyakov, S., Gilitschenski, I.: Spad: Spatially aware multi-view diffusers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10026–10038 (2024)

18. Kim, S., Li, K., Deng, X., Shi, Y., Cho, M., Wang, P.: Enhancing 3d fidelity of text-to-3d using cross-view correspondences. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10649–10658 (2024)
19. Koley, S., Bhunia, A.K., Sekhri, D., Sain, A., Chowdhury, P.N., Xiang, T., Song, Y.Z.: It’s all about your sketch: Democratising sketch control in diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7204–7214 (2024)
20. Kwon, M., Choi, J., Park, J., Jeon, S., Jang, J., Seo, J., Kwak, M., Kim, J.H., Kim, S.: Cameo: Correspondence-attention alignment for multi-view diffusion models. *arXiv preprint arXiv:2512.03045* (2025)
21. Labs, B.F.: FLUX.2: Frontier Visual Intelligence. <https://bfl.ai/blog/flux-2> (2025)
22. Leroy, V., Cabon, Y., Revaud, J.: Grounding image matching in 3d with mast3r. In: *European Conference on Computer Vision*. pp. 71–91. Springer (2024)
23. Li, C., Yang, Y., Shao, J., Zhou, H., Schwarz, K., Liao, Y.: Rerope: Repurposing rope for relative camera control. *arXiv preprint arXiv:2602.08068* (2026)
24. Li, P., Liu, Y., Long, X., Zhang, F., Lin, C., Li, M., Qi, X., Zhang, S., Xue, W., Luo, W., et al.: Era3d: High-resolution multiview diffusion using efficient row-wise attention. *Advances in Neural Information Processing Systems* **37**, 55975–56000 (2024)
25. Li, R., Yi, B., Liu, J., Gao, H., Ma, Y., Kanazawa, A.: Cameras as relative positional encoding. *arXiv preprint arXiv:2507.10496* (2025)
26. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747* (2022)
27. Liu, F.L., Fu, H., Lai, Y.K., Gao, L.: Sketchdream: Sketch-based text-to-3d generation and editing. *ACM Transactions on Graphics (TOG)* **43**(4), 1–13 (2024)
28. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9298–9309 (2023)
29. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Jiang, Q., Li, C., Yang, J., Su, H., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In: *European conference on computer vision*. pp. 38–55. Springer (2024)
30. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453* (2023)
31. Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using cross-domain diffusion. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9970–9980 (2024)
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017)
33. Ma, B., Gao, H., Deng, H., Luo, Z., Huang, T., Tang, L., Wang, X.: You see it, you got it: Learning 3d creation on pose-free videos at scale. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 2016–2029 (2025)
34. Mou, C., Wang, X., Xie, L., Wu, Y., Zhang, J., Qi, Z., Shan, Y.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 38, pp. 4296–4304 (2024)

35. Navard, P., Monsefi, A.K., Zhou, M., Chao, W.L., Yilmaz, A., Ramnath, R.: Knobgen: controlling the sophistication of artwork in sketch-based diffusion models. arXiv preprint arXiv:2410.01595 (2024)
36. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
37. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4104–4113 (2016)
38. Shi, R., Chen, H., Zhang, Z., Liu, M., Xu, C., Wei, X., Chen, L., Zeng, C., Su, H.: Zero123++: a single image to consistent multi-view diffusion base model. arXiv preprint arXiv:2310.15110 (2023)
39. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512 (2023)
40. Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. *Advances in Neural Information Processing Systems* **36**, 1363–1389 (2023)
41. Tang, M., Vinker, Y., Yan, C., Zhang, L., Agrawala, M.: Instance segmentation of scene sketches using natural image priors. In: Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers. pp. 1–10 (2025)
42. Voleti, V., Yao, C.H., Boss, M., Letts, A., Pankratz, D., Tochilkin, D., Laforte, C., Rombach, R., Jampani, V.: Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. In: European Conference on Computer Vision. pp. 439–457. Springer (2024)
43. 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)
44. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: Vggt: Visual geometry grounded transformer. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5294–5306 (2025)
45. Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: Dust3r: Geometric 3d vision made easy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20697–20709 (2024)
46. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
47. Wang, Z., Yuan, Z., Wang, X., Li, Y., Chen, T., Xia, M., Luo, P., Shan, Y.: Motionctrl: A unified and flexible motion controller for video generation. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11 (2024)
48. Xiang, J., Chen, X., Xu, S., Wang, R., Lv, Z., Deng, Y., Zhu, H., Dong, Y., Zhao, H., Yuan, N.J., et al.: Native and compact structured latents for 3d generation. arXiv preprint arXiv:2512.14692 (2025)
49. Xu, D., Nie, W., Liu, C., Liu, S., Kautz, J., Wang, Z., Vahdat, A.: Camco: Camera-controllable 3d-consistent image-to-video generation. arXiv preprint arXiv:2406.02509 (2024)
50. Yu, W., Xing, J., Yuan, L., Hu, W., Li, X., Huang, Z., Gao, X., Wong, T.T., Shan, Y., Tian, Y.: Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis. arXiv preprint arXiv:2409.02048 (2024)
51. Zhang, C., Li, B., Wei, M., Cao, Y.P., Gambardella, C.C., Phung, D., Cai, J.: Unified camera positional encoding for controlled video generation. arXiv preprint arXiv:2512.07237 (2025)

52. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3836–3847 (2023)
53. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
54. Zhang, S., Xu, H., Guo, S., Xie, Z., Bao, H., Xu, W., Zou, C.: Spatialcrafter: Unleashing the imagination of video diffusion models for scene reconstruction from limited observations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 27794–27805 (2025)
55. Zhao, Z., Lai, Z., Lin, Q., Zhao, Y., Liu, H., Yang, S., Feng, Y., Yang, M., Zhang, S., Yang, X., et al.: Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation. arXiv preprint arXiv:2501.12202 (2025)
56. Zhou, J., Gao, H., Voleti, V., Vasishta, A., Yao, C.H., Boss, M., Torr, P., Rupperecht, C., Jampani, V.: Stable virtual camera: Generative view synthesis with diffusion models. arXiv preprint arXiv:2503.14489 (2025)