

Motivation

Goal: Learn generative 3D priors from scalable multi-view training

Input: Any number of 2D views with poses

Output: Any number of 3D consistent novel views

✓ **Generative:** conditional generative model (more given, more constrained)

✓ **Consistency:** reference-to-target and target-to-target view consistencies

✓ **Scalability:** not scene-specific (NeRF); large-scale 2D posed images data

✓ **Generalisation:** trained with fixed number views (3-to-3), infer any-to-any

Points

Voxels

SDF

NeRF

Zero-1-to-3

EscherNet

Input

Any Number Novel View Synthesis

Multiview Diffusion Model

➤ Inheriting 2D diffusion model (SD 1.5) priors

Projection

Reshape

Lightweight Vision Encoder

ConvNeXt-V2 Tiny

Feed-Forward Block

Cross Attention Block with CaPE

Self-Attention Block with CaPE

Residual Block × 2

Linear

Reshape

Dot-Product Attention with CaPE

Linear

Linear

Linear

MatMul

SoftMax

Scale

MatMul

CaPE for Ref. Poses

CaPE for Tar. Poses

Value

Key

Query

Value

Key

Query

CaPE is applied on Key/Query

Camera Positional Encoding (CaPE)

➤ 3D Novel View Synthesis as Seq Image-to-Image Generation

Language

Token:

Position:

3D Vision

Token:

Position:

	Type	Length	Origin
Language	Discrete	Linear	Unbounded
3DV	Rotation	Cyclic	Bounded
	Translation	Linear	Unbounded

➤ Each view's absolute position rotates its token features

➤ Only relative position is obtained in dot-product self attention

Query / Key

Absolute Position

Rotate Feature

Position Encoded Query / Key

Results

Novel View Synthesis (NeRF)

InstantNGP (Scene Specific Training)

3D Gaussian Splatting (Scene Specific Training)

EscherNet (Zero Shot Inference)

Novel View Synthesis (Text-to-3D)

3D Reconstruction

Novel View Synthesis (GSO/RTMV)

Novel View Synthesis (Real-world 6DoF)