SESSION 5: VISUAL REPRESENTATION GENERATION AND EXPLORATION

Introduction

- Input:
  - Images associated with the room
    - Spatially referenced
  - 3D room model or pixel-wise information
    - Single scene
    - Walls, ceilings, floor
    - Multi-modal information for specific tasks
- Output:
  - Editable representations
  - VR exploration, Extended Reality, Editing appearance

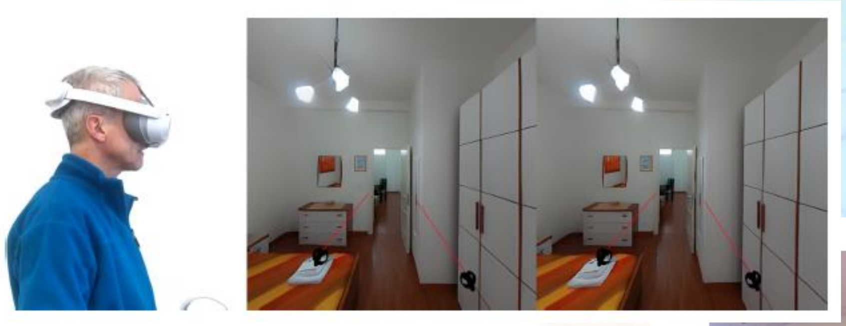

PanoStereo, Elsevier CAG 2024

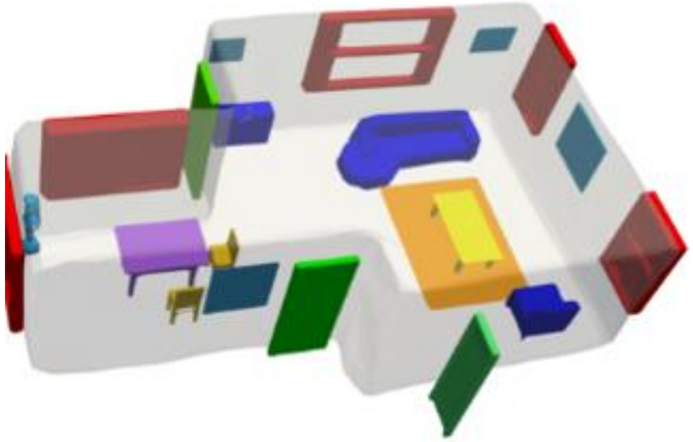

PanoContextFormer, IEEE CVPR 2024

Pano2Room, ACM TOG 2025

Application context

- **Omnidirectional imagery**
  - Fundamental component for creating immersive content from real-world scenes
- **Virtual tour popular in the real-estate domain**
  - Presentation to virtual visitors
  - Popularized during Covid pandemic
- **Other application domains:**
  - Tourism, architecture, construction

https://matterport.com/industries/real-estate

The Met 360° Project | The Metropolitan Museum of Art Visit

Images may be subject to copyright. Learn More

Outline

- Single image solutions for supporting seamless exploration and editing in multiple devices
  - 3-DOF motion
    - Viewpoint rotation without translation on enriched image-based representations
  - Small displacement 6-DOF motion
    - Limited viewpoint translation
    - Stereo generation for immersive exploration
  - Large displacement 6-DOF motion
    - Full 3D reconstructions for full explorations
- Overview of SOTA methods and our recent contributions

Background on 3-DOF solutions

Static solutions: Pano, Omni, Sphere viewers

- Limitation: they support only viewpoint rotation
  - No parallax
- Available online and using various representations
- Integration with WebVR and WebXR for direct usage with VR devices
- Cubemaps (krpano)
- Stereo panoramic couples (sphere stereo viewer)
  - A-frame

Pannellum.org

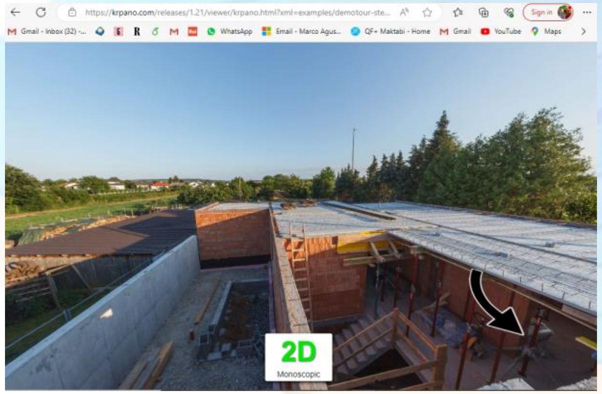

krpano

360° | CSE-LAB-ICT (renderstuff.com)

Static solutions: mesh-based rendering

- Spherical dome tessellation
  - Iterative subdivision from icosahedron
  - Subdivision level 8 leads to ~1.3M verts and ~1.3 M triangles
- Basic rendering mode for original images and head rotation movements (viewer in the camera position)

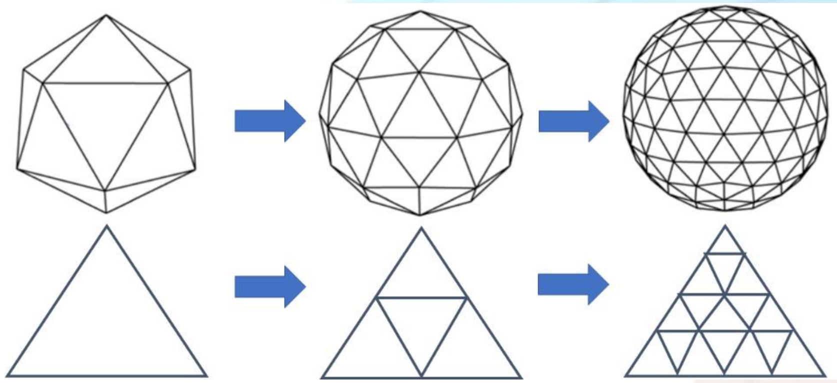

Static solutions: cubemaps

- From equirectangular image to six textures to be mapped to the faces of a cube
- Graphics hardware accelerates texture fetching in shaders (GL_TEXTURE_CUBEMAP)
  - Popular for environment maps in games
- Used in popular panoramic image viewer like krPano

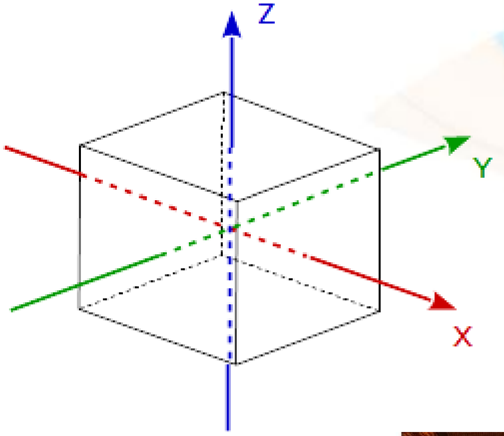

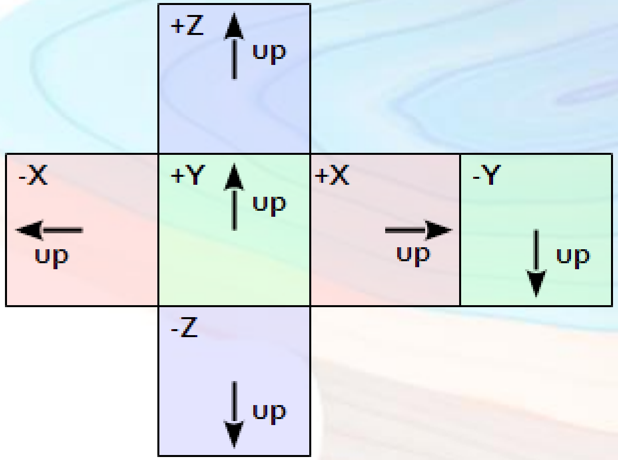

Courtesy: paulbourke.net

Viewport rendering: single-pass ray casting (e.g., on equirectangular image)

- Draw a quad in screen coordinates
- Fragment shader:
  - Pass view and perspective parameters: fov, distance of view plane
  - For each fragment:
    - cast a ray from eye position to intersect the spherical scene
    - fetch the corresponding texels from equirectangular images through inverse spherical mapping

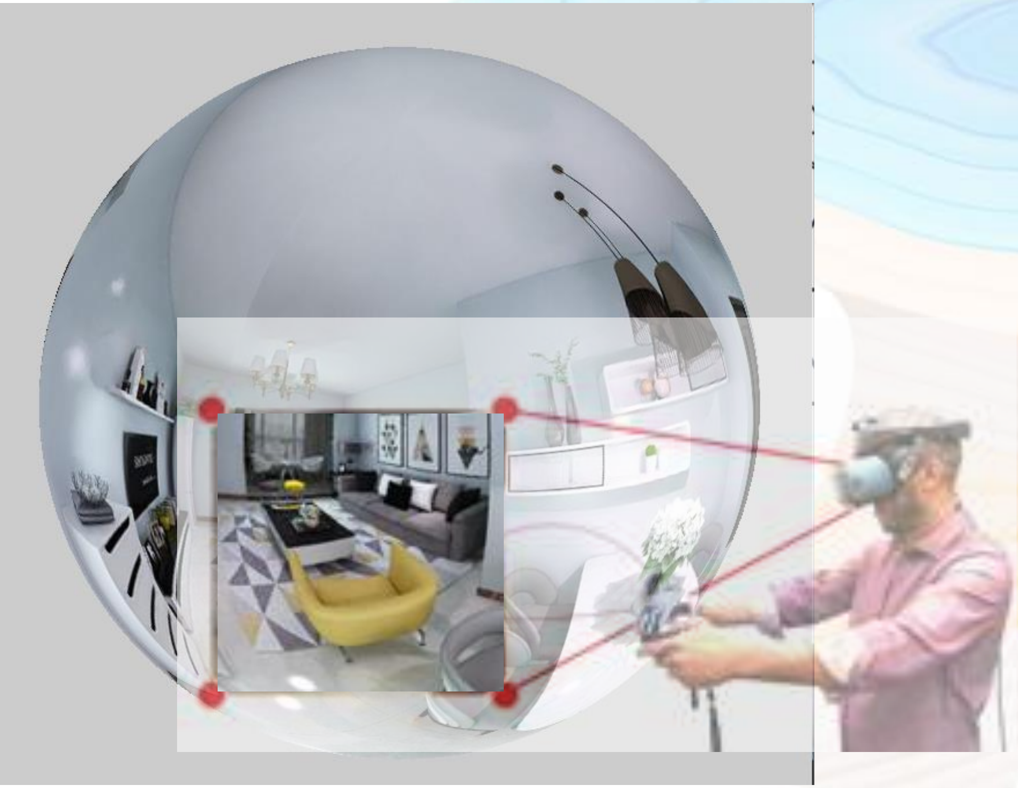

Image-based editing using 3-DOF solutions

Applications: Diminished Reality

- Instant photorealistic view and depth of a panoramic indoor scene emptied of furniture and clutter
- Enables compelling and immersive XR applications, such as re-furnishing or planning of interior spaces

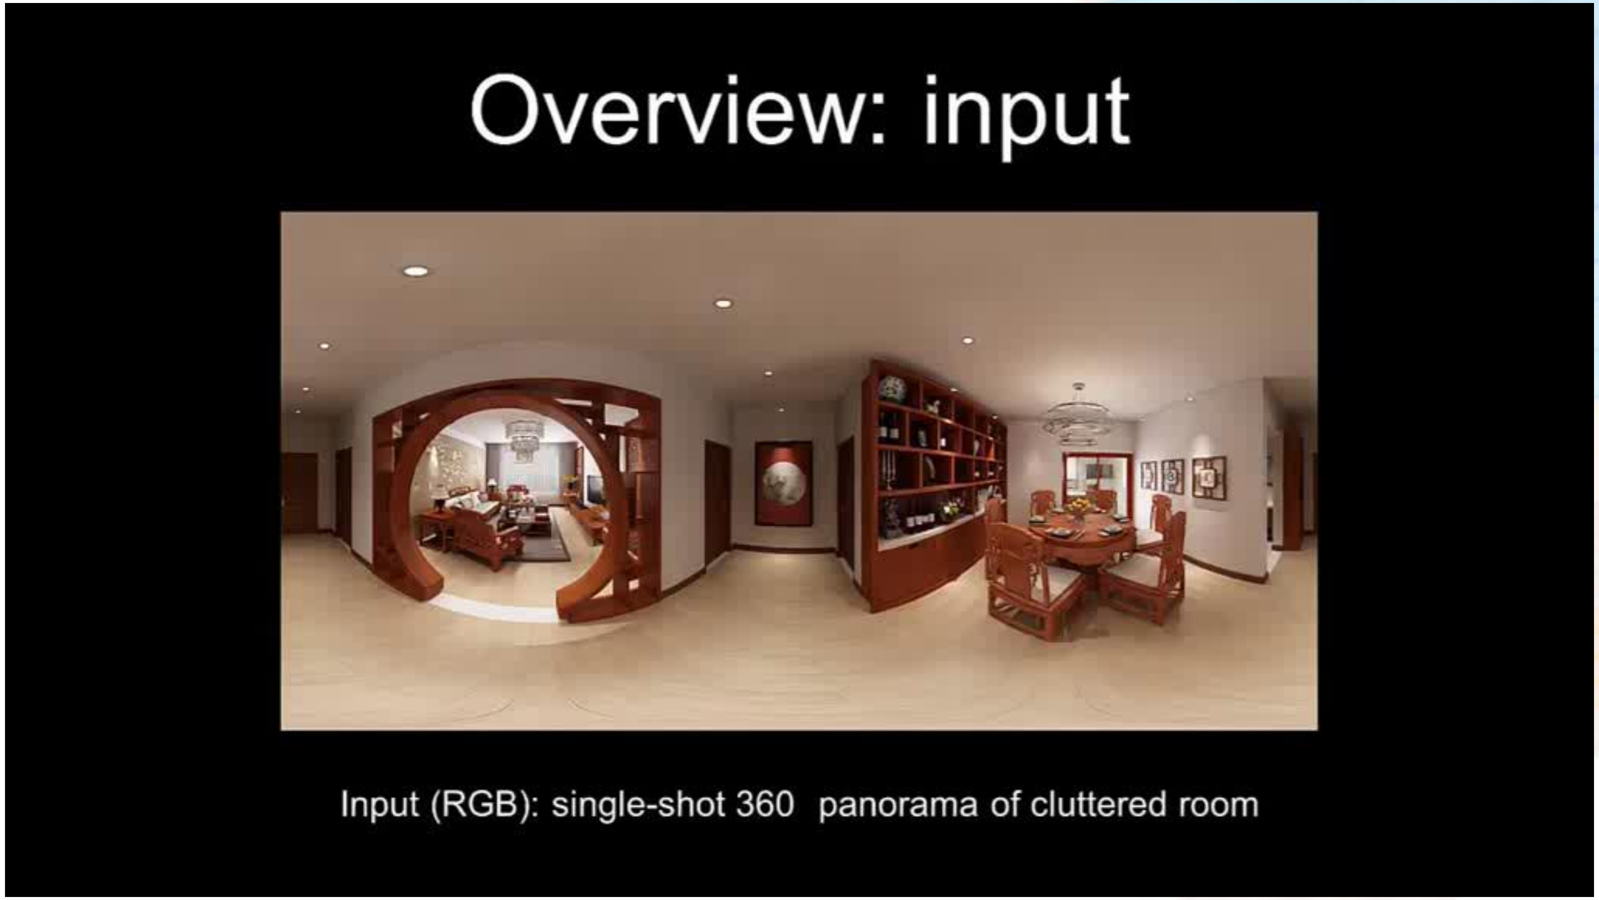

Giovanni Pintore, Marco Agus, Eva Almansa, and Enrico Gobbetti.
Instant Automatic Emptying of Panoramic Indoor Scenes. IEEE TVCG, 28(11): 3629-3639, 2022.
DOI: 10.1109/TVCG.2022.3202999. Proc. ISMAR.

Applications: diminished reality

- **Light-weight end-to-end deep network**
  - Input: 360 image of a furnished indoor space
  - Output: 360 photorealistic view and architecturally plausible depth of the same scene emptied
    - Very low latency
  - NB. Learning on synthetic dataset transferred to real-world cases

Pintore, Almansa, Agus, Gobbetti. IEEE TVCG 2022

Key contributions

- **End-to-end network providing, at interactive rate, a panoramic indoor scene emptied automatically without user intervention**
  - Linear fashion and depth-separable gating
  - Visual and geometric constraints are applied only at training time
- **Geometric representation of the scene as additional output**
  - Basis for further processing in XR application
  - Enables robust and effective pixel-wise geometric priors
- **Loss function that combines photorealistic and geometric terms**
  - Virtual normals to recover the salient characteristics of indoor structures
  - Flatness and smoothness, less restrictive than Manhattan World, etc.

Model architecture

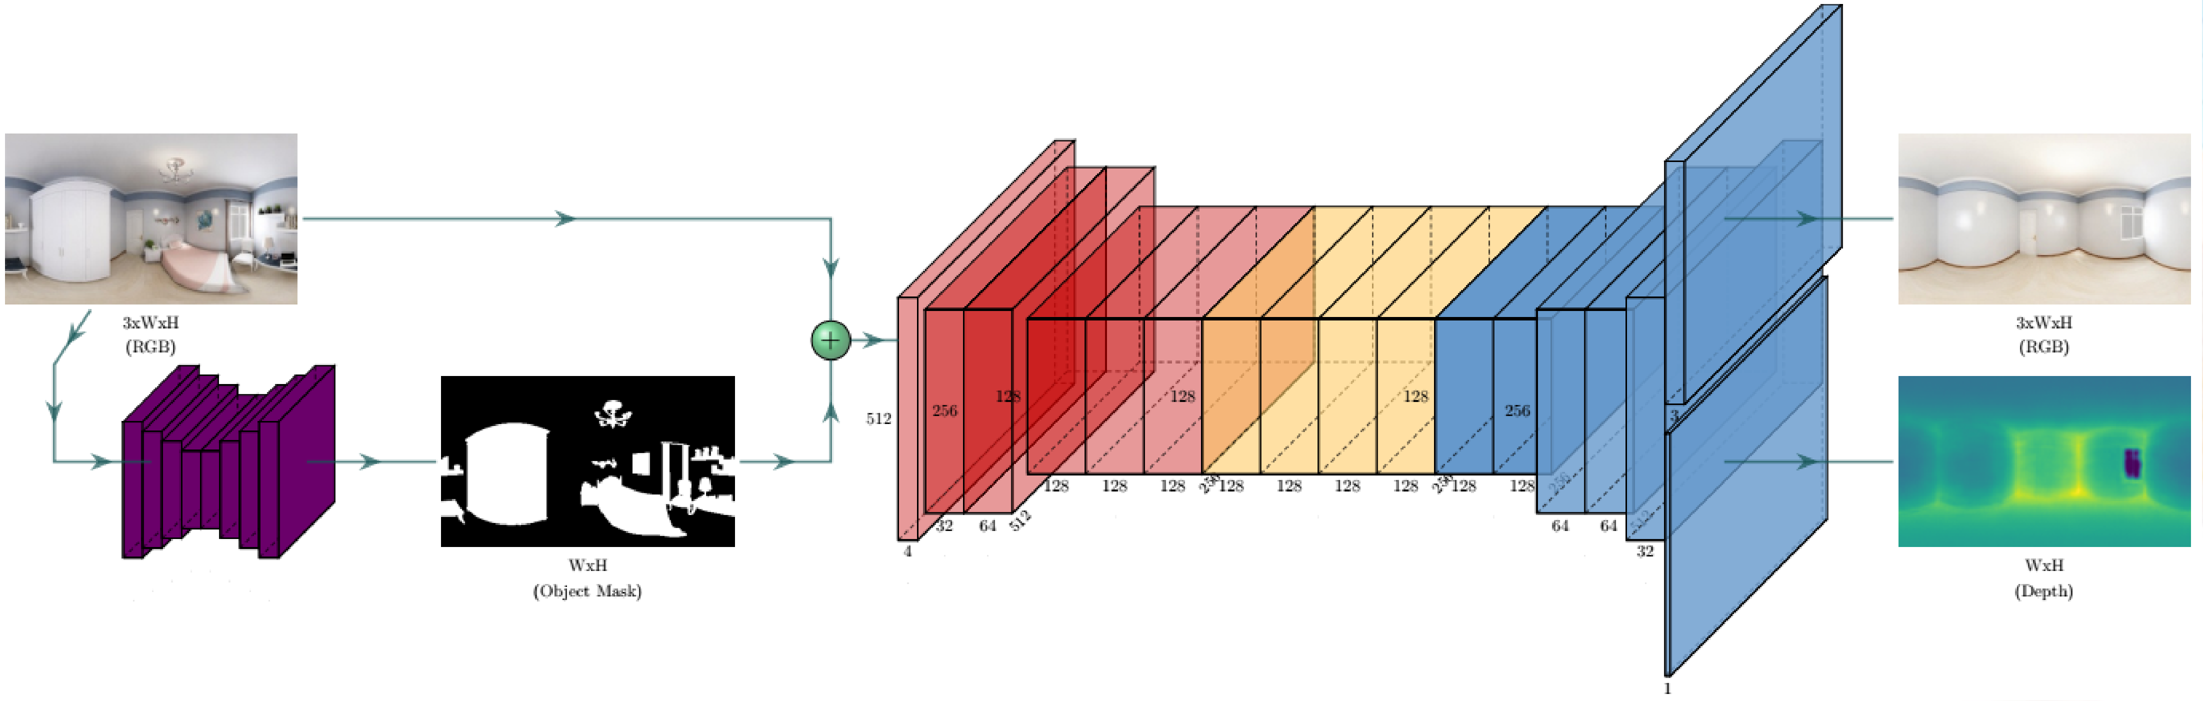

Pintore, Almansa, Agus, Gobbetti. IEEE TVCG 2022

Methods

- **Clutter identification**
  - Automatic binary mask
  - Geometric mask obtained by comparing the ground-truth depths
  - Very lightweight encoder-decoder network
  - Binary cross-entropy loss

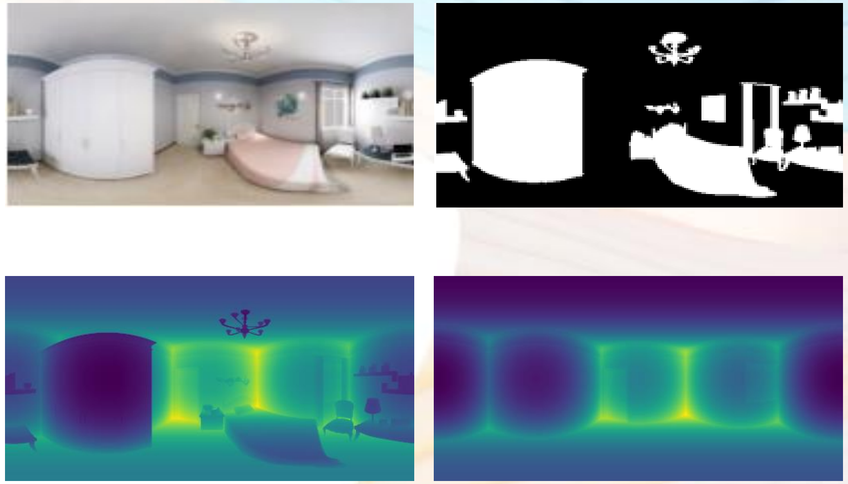

$$-\frac{1}{n}\sum_{p\in {D_m}^c}\left(\hat{p}\log p+(1-\hat{p})\log(1-p)\right)$$

Methods

- **Empty scene synthesis**
  - Image inpainting baseline
    - Learnable gating
  - Light Weight Gated Convolutions (LWGC)
    - simplify training
    - low latency at inference time
  - Repeated dilations used for the bottleneck
    - Aggregates multi-scale contextual information without losing resolution
    - Avoid increasing number of weights

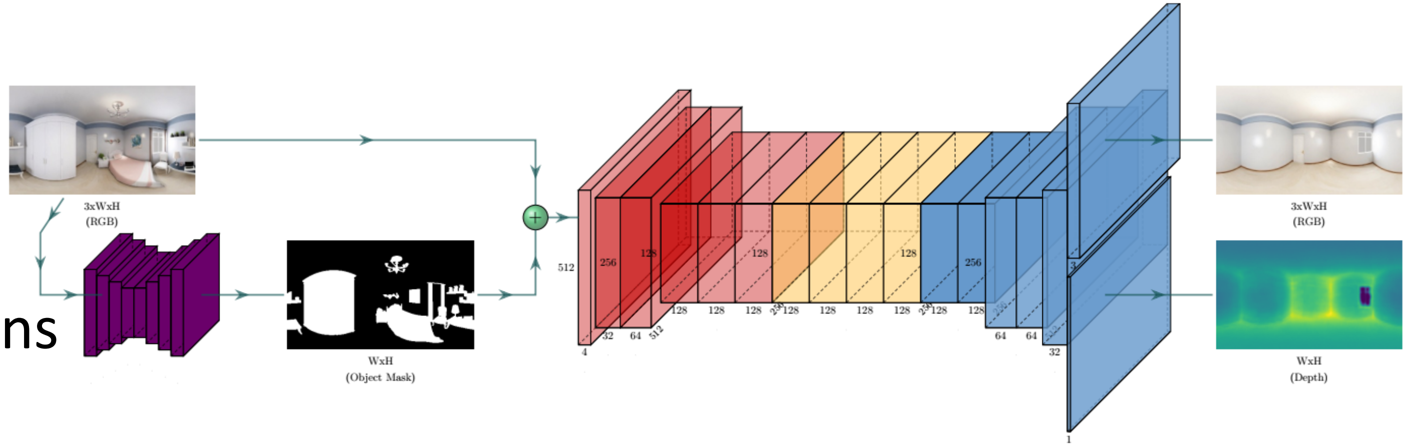

$$
\begin{aligned}
G &= conv(W_g, I)\\
F &= conv(W_f, I)\\
O &= \sigma(G) \odot \psi(F)
\end{aligned}
$$

$$D_{y,x} = \sigma(b + \sum_{i=-k'_h}^{k'_h} \sum_{j=-k'_w}^{k'_w} W_{k'_h+i, k'_w+j} \cdot I_{y+\eta i, x+\eta j})$$

Methods

- **Training and losses**
  - Combination of a visual term and a geometric term
  - Visual term
    - L1 with data-driven perceptual and style losses
  - Geometric term
    - combination of low- and high-order 3D constraints
    - High-order based on virtual normal consistency

$$\mathscr{L}_{vis} = \lambda_{px}\mathscr{L}_{px} + \lambda_{perc}\mathscr{L}_{perc} + \lambda_{style}\mathscr{L}_{style}$$

$$\mathscr{L}_{geom} = \lambda_d\mathscr{L}_d + \lambda_n\mathscr{L}_n$$

$$n_i = \frac{\overrightarrow{P_aP_b} \times \overrightarrow{P_aP_c}}{\left\|\overrightarrow{P_aP_b} \times \overrightarrow{P_aP_c}\right\|} \qquad \mathscr{L}_n = \frac{1}{N}\sum_{i=1}^{N}\left\|{n_i}^{pred} - {n_i}^{gt}\right\|$$

$$C = \{\alpha \geq \angle(\overrightarrow{P_aP_b}, \overrightarrow{P_aP_c} \leq \beta, \alpha \geq \angle(\overrightarrow{P_bP_c}, \overrightarrow{P_bP_a} \leq \beta\}$$

$$\mathscr{L}_{perc} = \sum_{n}^{N-1}\left\|\psi_n(I_{out}) - \psi_n(I_{gt})\right\|_1$$

$$\mathscr{L}_{style} = \sum_{n}^{N-1}\left\|K_n(\psi_n(I_{out})^T\psi_n(I_{out})) - \psi_n(I_{gt})^T\psi_n(I_{gt})\right\|_1$$

Some results

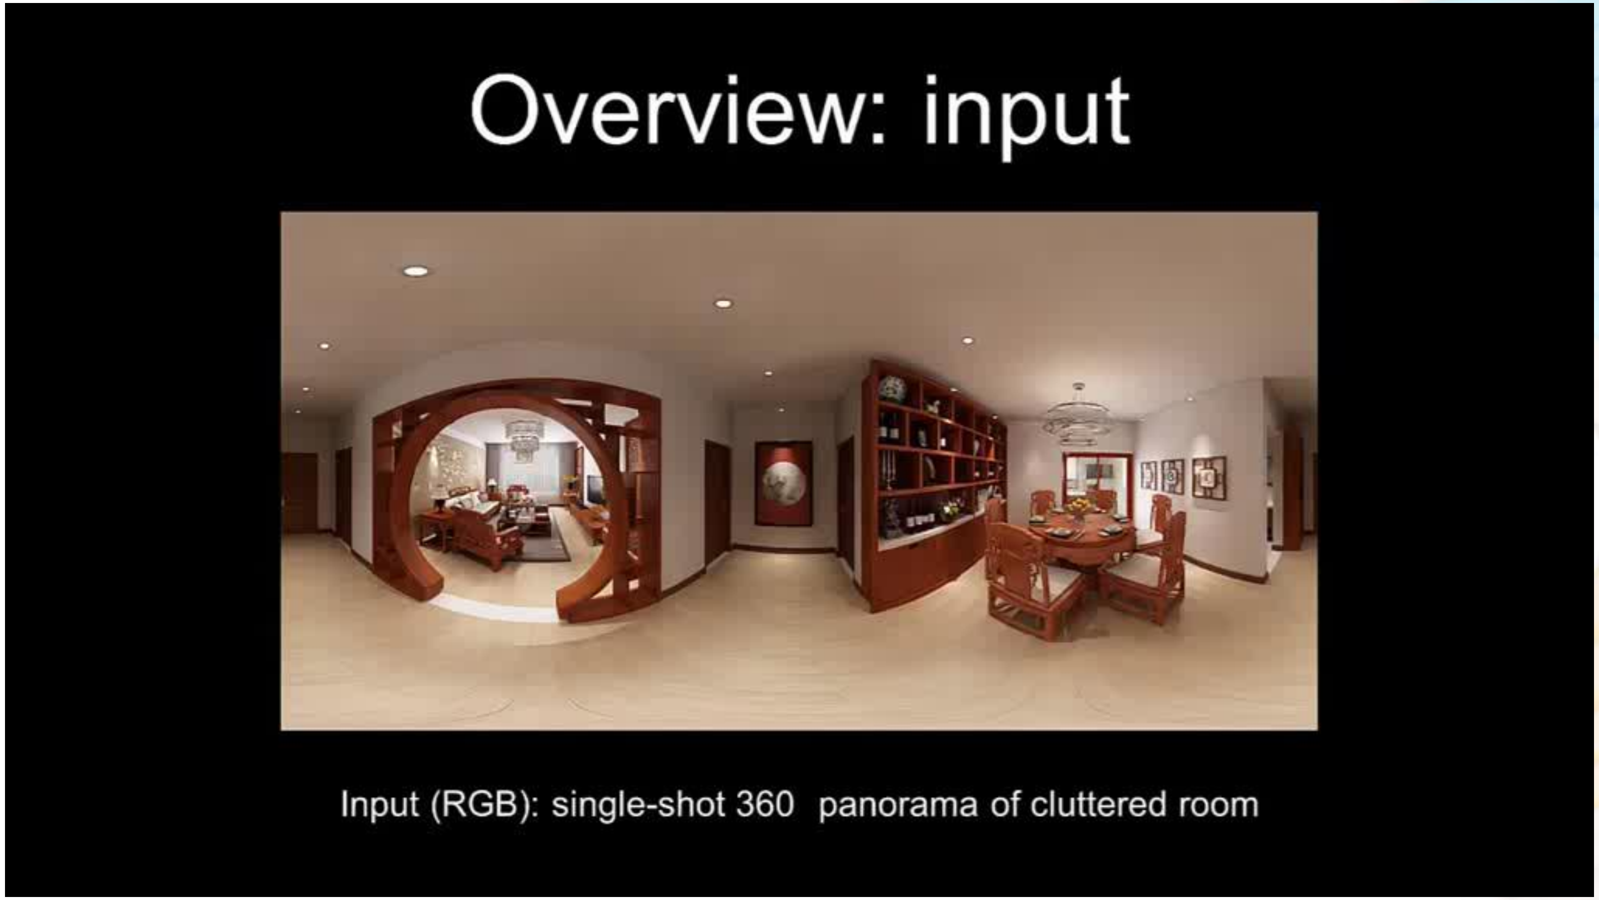

Giovanni Pintore, Marco Agus, Eva Almansa, and Enrico Gobbetti.
Instant Automatic Emptying of Panoramic Indoor Scenes. IEEE TVCG, 28(11): 3629-3639, 2022.
DOI: 10.1109/TVCG.2022.3202999. Proc. ISMAR.

Applications: scene modification

- Support editing and modifications
  - Adding/removing clutter/objects
  - Place POI/annotations
  - 3D multimedia hyperlinks
- Appearance modification
  - Lighting (Zhi et al, ACM TOG 2022)
  - Style Transfer (Tukur et al, IEEE ICCVW 2023)
- Virtual staging as emerging field

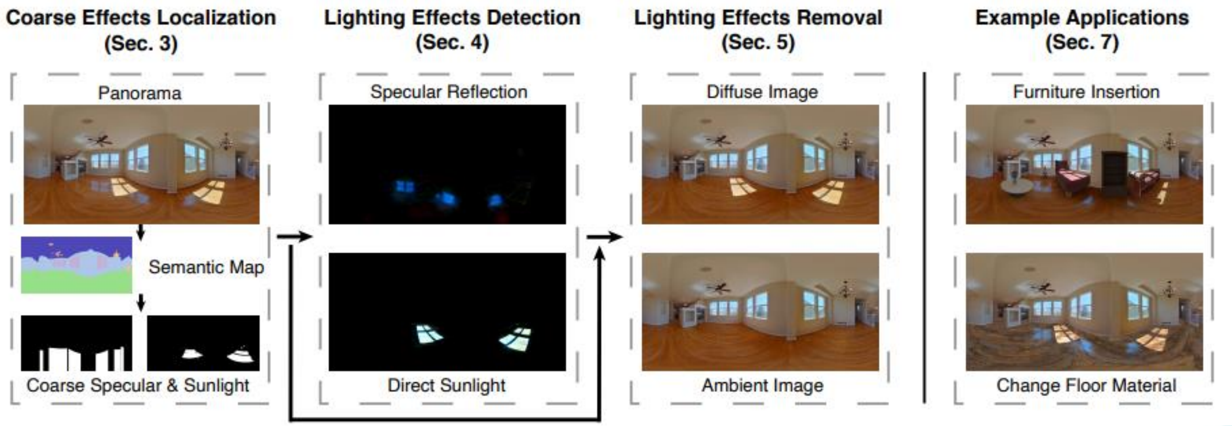

Zhi et al. , ACM TOG 2022

Applications: virtual staging

- An interactive editing and rendering system for indoor DR/XR applications from a single panoramic image

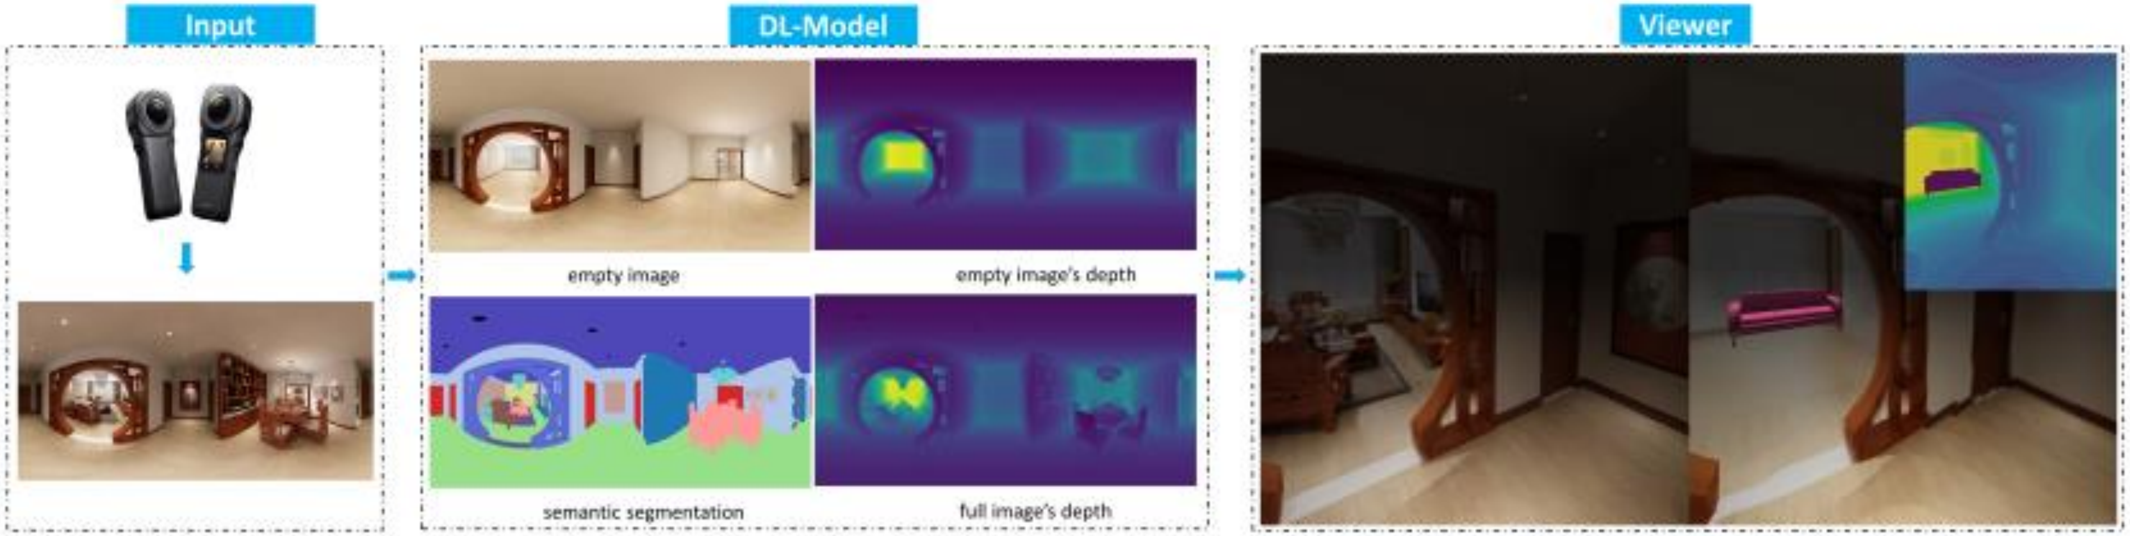

Muhammad Tukur, Giovanni Pintore, Enrico Gobbetti, Jens Schneider, and Marco Agus.
SPIDER: A framework for processing, editing and presenting immersive high-resolution spherical indoor scenes.
Graphical Models, 128: 101182:1-101182:11, July 2023. DOI: 10.1016/j.gmod.2023.101182.

Virtual object placement

- Basic operations for Virtual Staging
  - Placement of synthetic objects
  - Transfer of semantic content from cluttered scene to empty scene

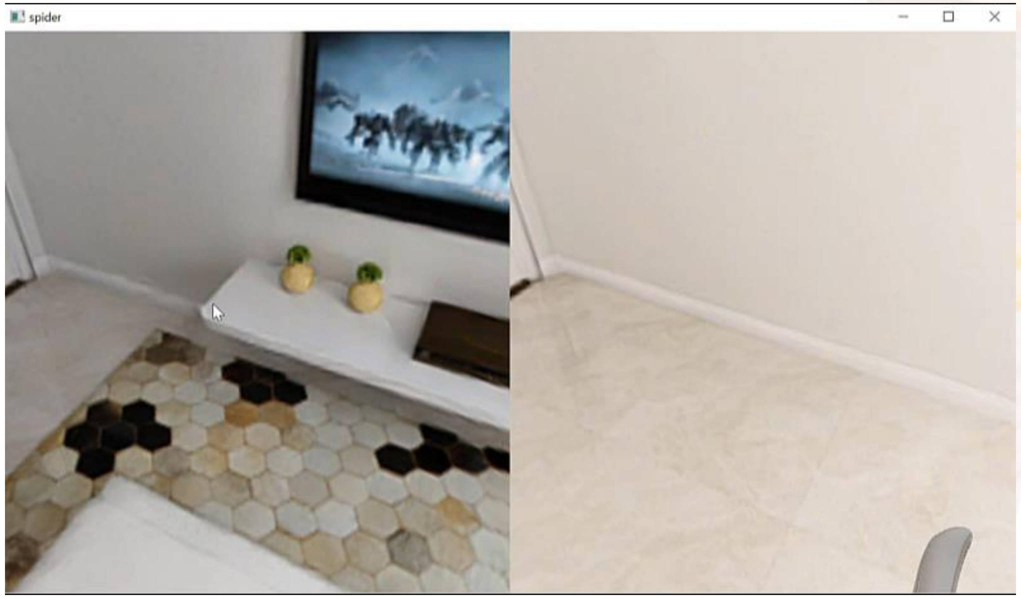

Applications: editing indoor panoramas

- GAN-based photorealistic style transfer

Tukur, Ur Rehman, Pintore, Gobbetti, Schneider, and Agus.PanoStyle, ICCVW 2023

GAN-based photorealistic style transfer

- Two main additions on top of a classical GAN-based style transfer architecture:
  - Shading decomposition
  - Geometry constraints

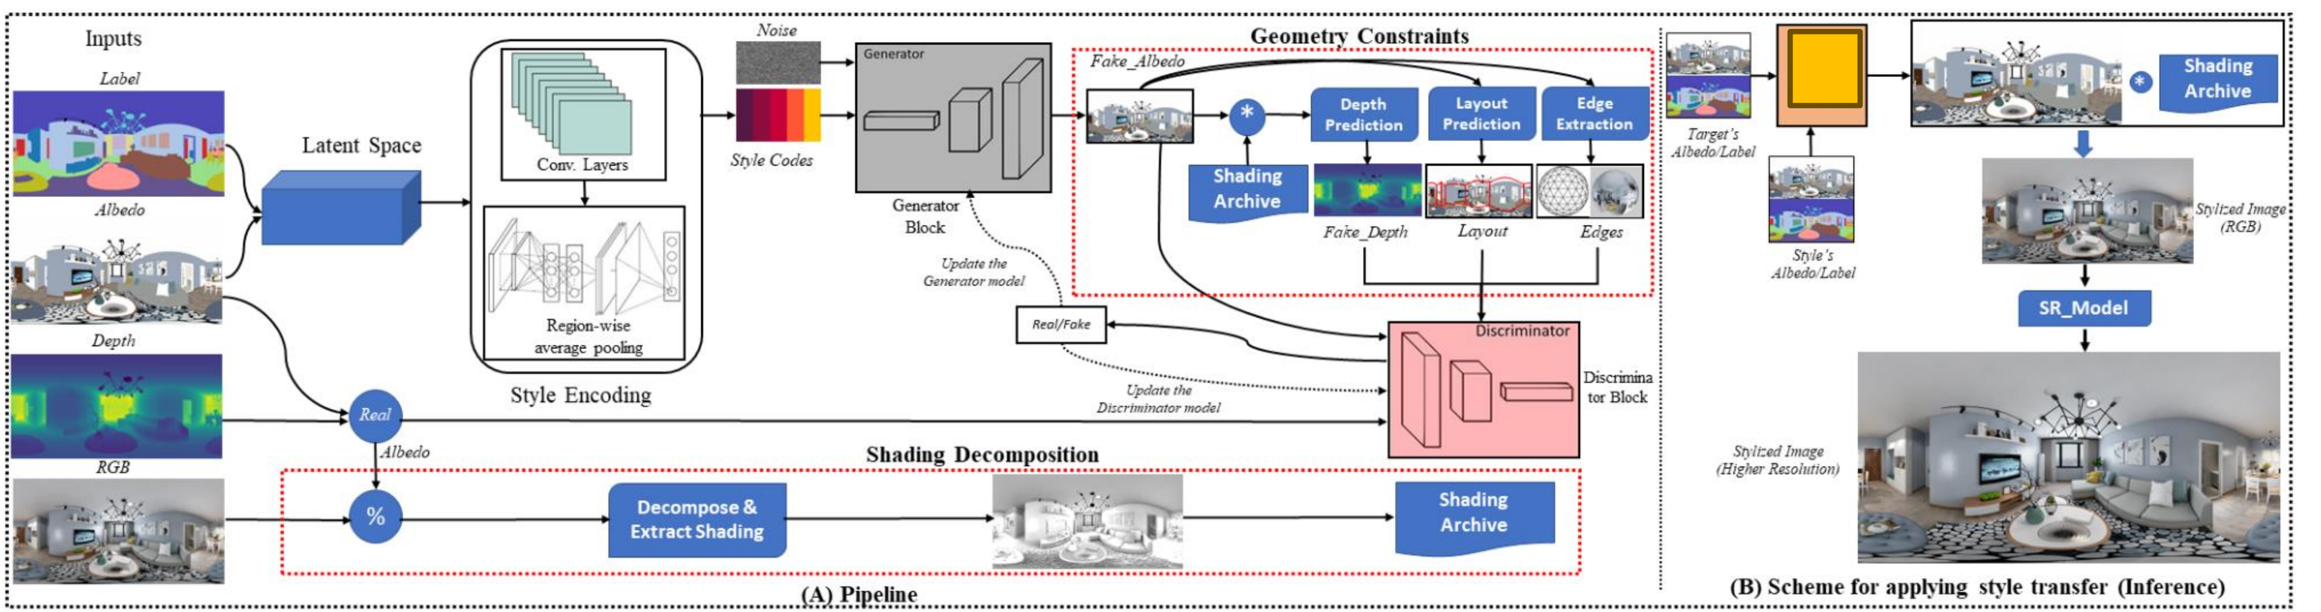

Tukur, Ur Rehman, Pintore, Gobbetti, Schneider, and Agus.PanoStyle, ICCVW 2023

Intrinsic shading decomposition

- Normalized shading signal for removing secondary effects

$$I_{\text{shad}} := \max\left(\left\|I_{\text{rgb}} \oslash I_{\text{alb}}\right\|_2, 1\right) \Rightarrow \hat{I}_{\text{rgb}} = I_{\text{shad}} \cdot I_{\text{alb}}$$

- Style codes computed on albedo and shading a-posteriori

Geometry constraints

- Enforce scene depth, layout and edge consistency with additional geometry losses
  - For depth prediction, SliceNet [Pintore et al, 2021]
  - For layout prediction, HorizonNet [Sun et al., 2019]

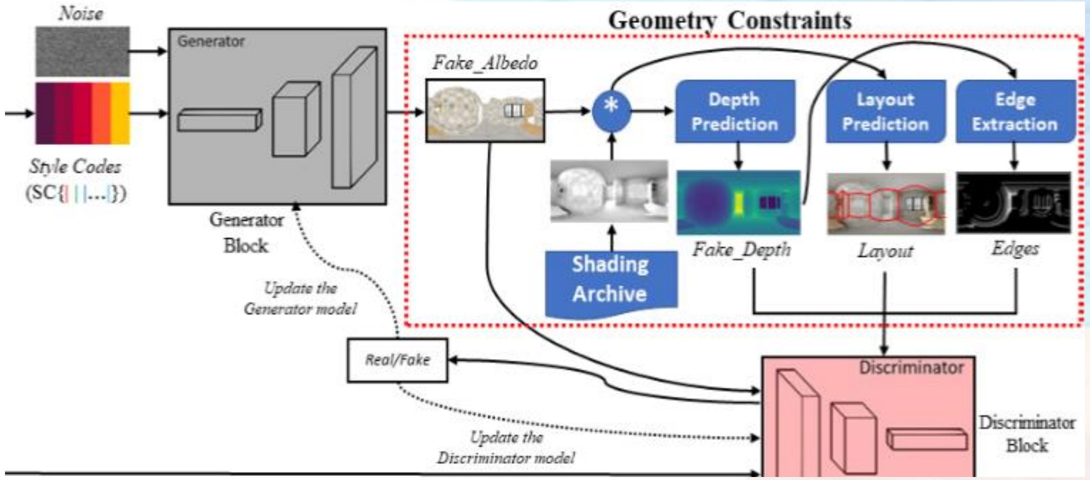

$$\mathcal{L}^{\text{geo}}_{\text{depth}} = \sum_{ij} w_{ij} \left\| D^{\text{G}}_{ij} - D^{\text{R}}_{ij} \right\|_1$$

$$\mathcal{L}^{\text{glob}}_{\text{depth}} = \sum \left\| F_n(D^{\text{G}}) - F_n(D^{\text{R}}) \right\|_1$$

$$\mathcal{L}^{\text{loc}}_{\text{depth}} = \sum_n \left\| K_n \left(F_n \left(D^{\text{G}}\right)^T F_n \left(D^{\text{G}}\right) - F_n \left(D^{\text{R}}\right)^T F_n \left(D^{\text{R}}\right) \right) \right\|_1$$

$$\mathcal{L}^{\text{geo}}_{\text{layout}} = \left\| L^{\text{G}} - L^{\text{R}} \right\|_1,$$

$$\mathcal{L}^{\text{glob}}_{\text{layout}} = \sum_n \left\| H_n \left(L^{\text{G}}\right) - H_n \left(L^{\text{R}}\right) \right\|_1,$$

$$\mathcal{L}^{\text{loc}}_{\text{layout}} = \sum_n \left\| K_n \left(H_n \left(L^{\text{G}}\right)^T H_n \left(L^{\text{G}}\right) - H_n \left(L^{\text{R}}\right)^T H_n \left(L^{\text{R}}\right) \right) \right\|_1$$

Some results

Limitations of static solutions

- Indoor usage scenario
    - Users want to perceive layout, scene composition, metric distances
    - Users want to explore behind corners and check clutter details
- Lack of stereo and parallax cues
    - Limited immersive capabilities on VR devices
- Spherical mapping
    - Flat appearance is a limitation for indoor applications
    - Clutter close to viewer as wells as architecture details not correctly represented

6-DOF solutions

6-DOF solutions

- They try to provide consistent viewpoint translation or accurate stereo cues for immersive exploration
- Two main categories according to input data:
  - Integration of multiple images and videos and multistereo setups (Multi-input)
  - Data-driven solutions working on single panoramic images (Single-input)
- Two main categories according to the output representation:
  - Image-based solutions (Novel View Synthesis)
  - Representation-based solutions (Renderable data structures)
- Mix and match of these categories

Input: multiple spherical images

- Multi-spherical images
  - Extension of Multi-planar images for spherical shells (Attal et al, 2020)
  - Conversion to layered mesh representation (Broxton et al., 2020)
  - 360 sweep video
    (Bertel et al., 2020)

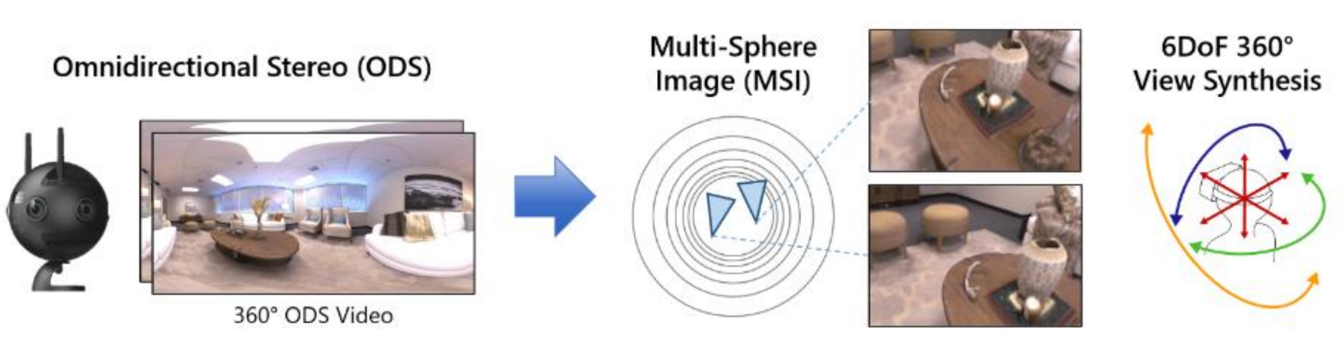

Attal et al. , Matryodska, ECCV 2020

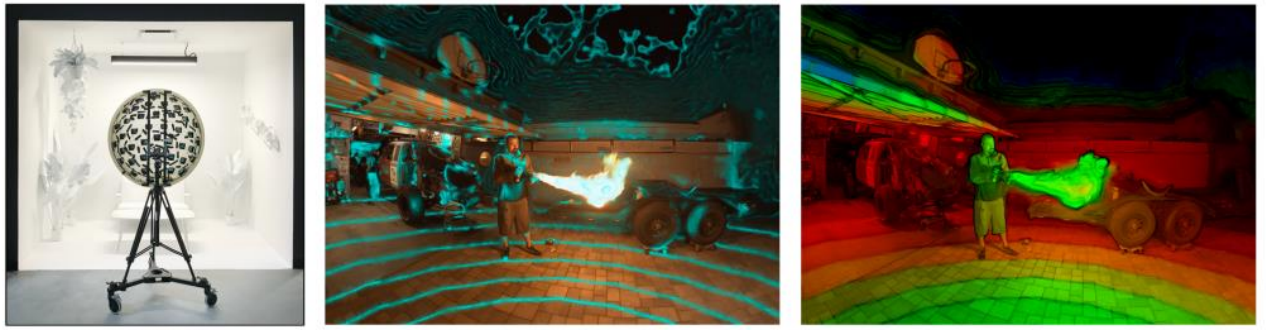

(a) Capture Rig (b) Multi-Sphere Image (c) Layered Mesh Representation

Broxton et al. ACM TOG 2020

Our focus today:
6-DOF from a single image

Input: single spherical image

- Importance: single or sparse panoramic capture is fast and convenient
- Usage scenario: quick capture of a single room, or acquisition of large multi-room environment with one image/room
- Goals: make single panorama more immersive and increase exploration freedom
- Two classes of approaches according to the allowed motion:
  - Small displacement:
    - support stereo cues: eyes stay in a circle around the original capture position
    - Support small head movements: eyes stay close to the original capture position (range of 50 cm or less – to support head motion for static viewer)
  - Large displacement:
    - From limited viewer motion to full exploration (also of originally unseen scene portions!)

6-DOF with small displacement

Restricted small displacement motion

- Specialized solutions for moderately small change of perspective or disocclusion of limited portion of the scene
- Creation of compact representation optimized for rendering and valid only for the known viewing environment

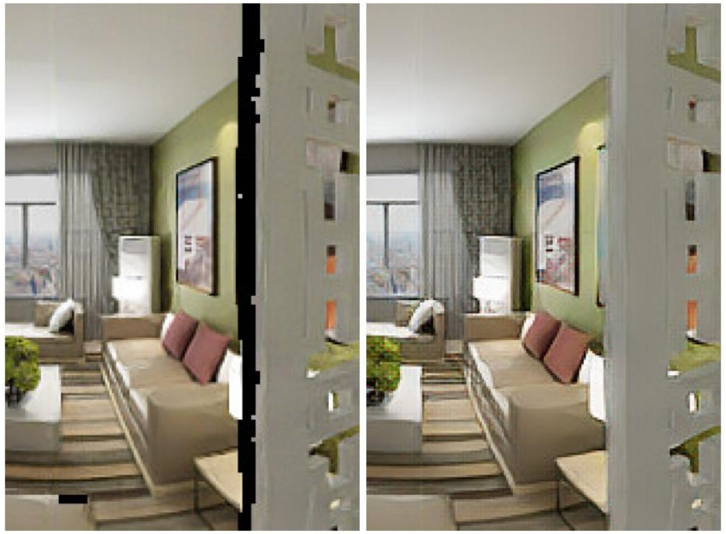

Geometry/Proxy solutions

- Panoramic images + depth map
  - Point cloud rendering [Huang, 2017]
  - View-dependent meshes from point clouds [Tukur et al, GMOD 2023]
- Integrate depth maps with multiple image signals to create a geometry proxy
  - OmniPhotos[Bertel et al., 2020]
  - It requires single sweep with a consumer 360° video camera as input

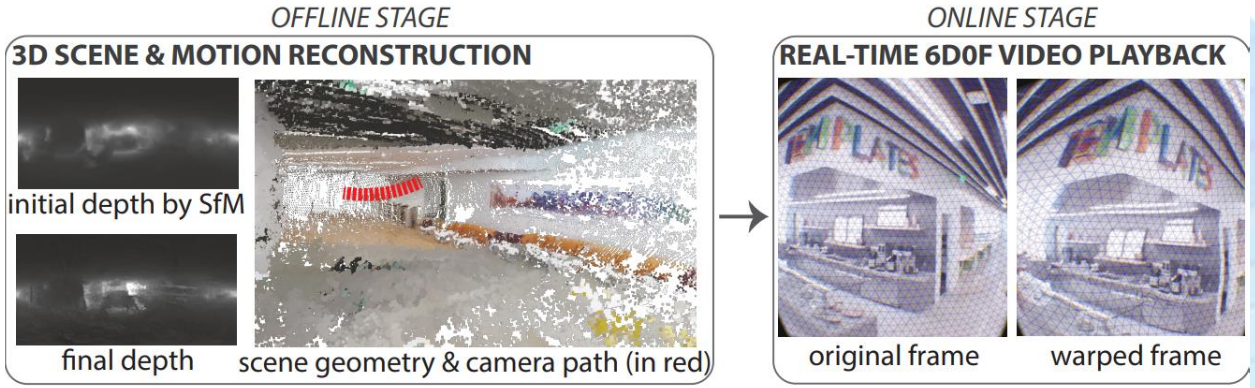

Huang et al., IEEE VR 2017

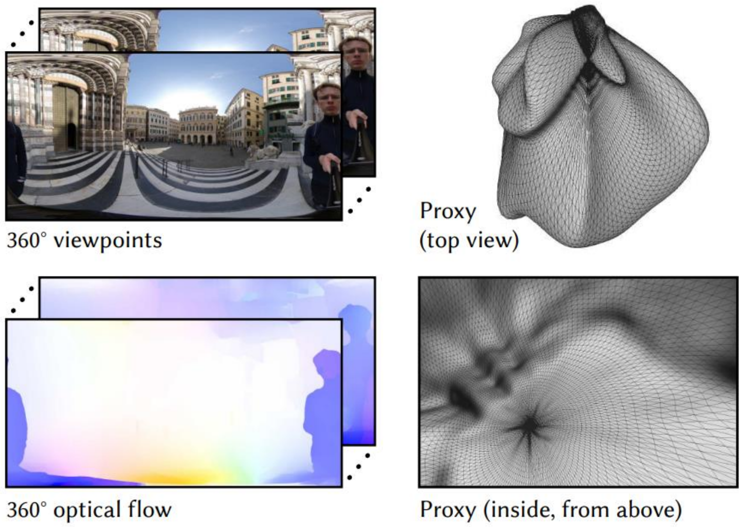

Bertel et al., ACM TOG 2020

Image synthesis solutions

- End-to-end synthesis networks
  - [Pintore, Bettio, Agus, Gobbetti, IEEE TVCG 2023]
  - Panoramic Novel View Synthesis [Xu et al.,IEEE CVPR 2021]
- Novel views through interpolation or blending of nearby viewpoints
  - FILM [Reda et al., 2022]
  - Gram matrix loss that measures the correlation difference between features

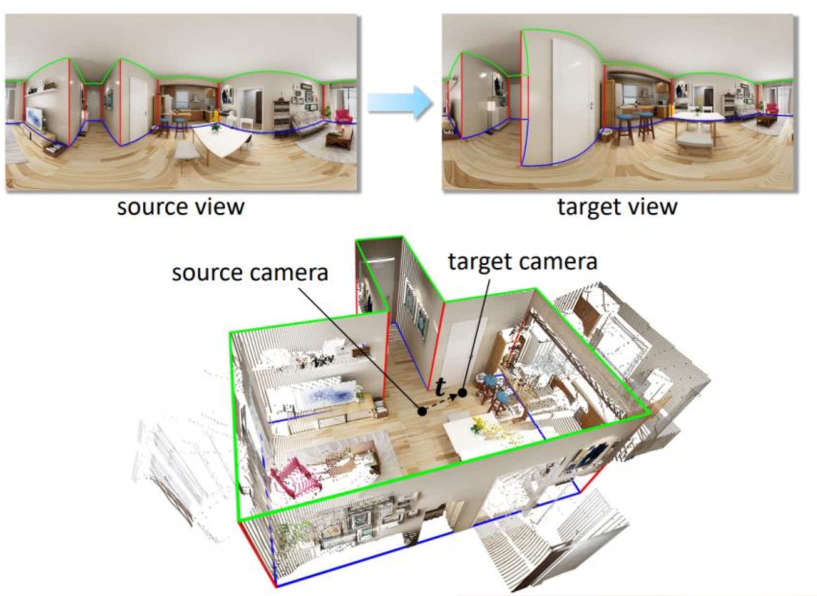

Xu et al., IEEE CVPR 2021

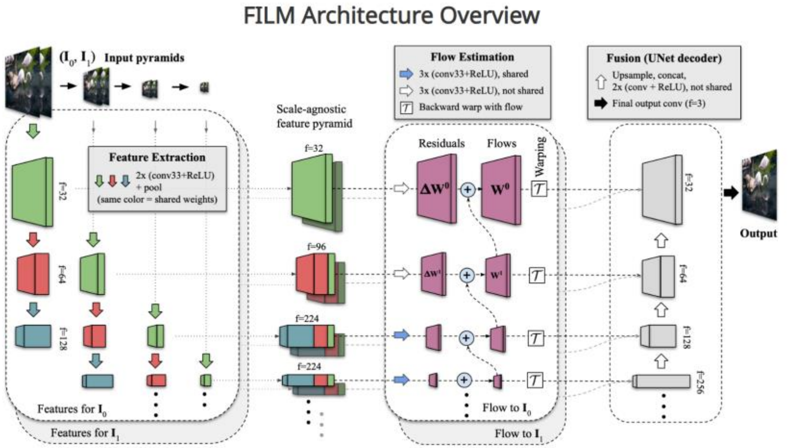

Reda et al., ECCV 2022

Image-based GAN-based view synthesis

- **Main Idea:** Explicit or implicit geometry estimation, to perform occlusion-aware reprojection and synthesize the disoccluded content
  - Complex training and inference
  - Room layout prior, to guide the generation of target views [PNVS, Xu et sl., IEEE CVPR 2021]
- **Application:** Low-latency extraction of novel poses to extract perspective images in real-time responding to both translation and rotation for small movements
  - [Pintore, Bettio, Agus, Gobbetti, IEEE TVCG 2023]

Features

- Client-server architecture
  - Thin WebGL client manages head motion
  - Server computes images for head translation
  - 70 Hz refresh, 10 fps panorama updates, workspace ~30 cm
- Novel views synthesis respecting Atlanta World model constraints
  - Model exploits Gravity Aligned Features and LSTM for managing spatial relationships
  - Depth and layout prediction for constraining view synthesis

Forward pipeline

- **Signal extraction module:** concurrent estimation of scene depth, scene latent representation, 3D room shape and floor occupancy map
- **View Synthesis module:** lightweight approach to generate novel panoramic views
  - Limited number of layers, combining gated and dilated convolutions

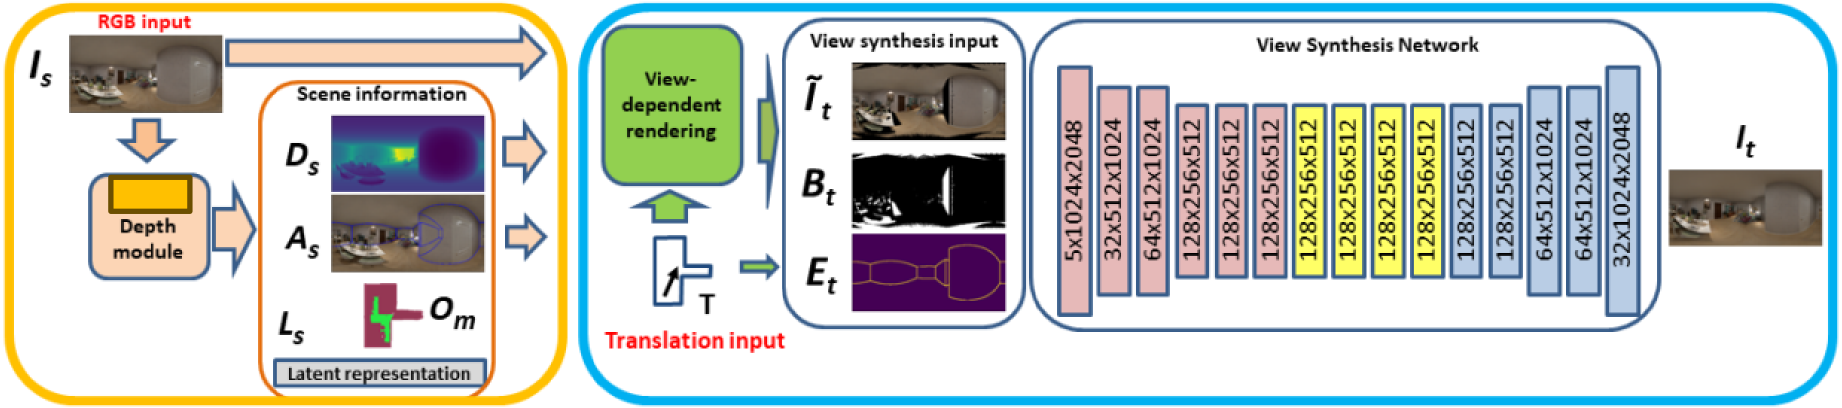

Training stage

- Objective functions for indoor structural consistency
  - Design of losses based on direct estimation and latent-space features
  - Geometric perceptual and geometric style loss

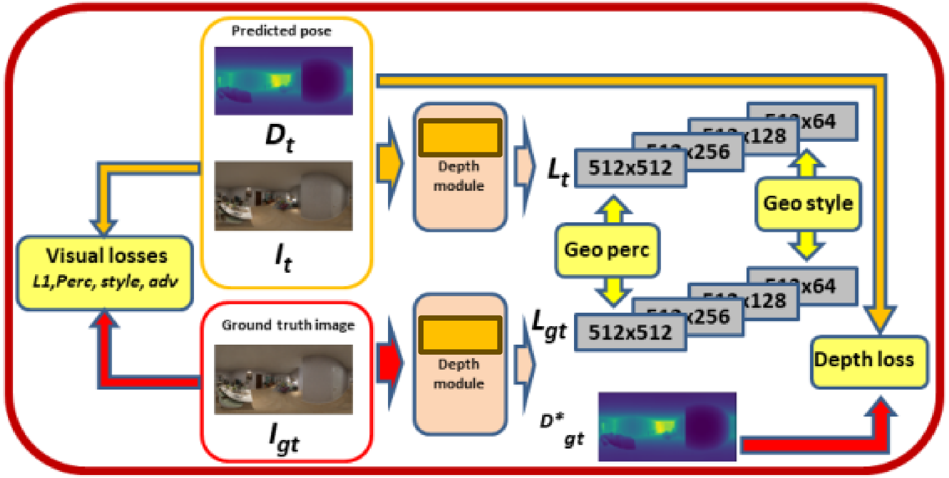

$$\mathscr{L}_{adm} = \lambda_d \mathscr{L}_d - \lambda_{ss} \mathscr{L}_{ss} + \lambda_l \mathscr{L}_l + \lambda_h \mathscr{L}_h$$

$$\mathscr{L}_{geocont} = \sum_{n}^{4} \| L_n(I_t) - L_n(I_{gt}) \|_1$$

$$\mathscr{L}_{geostyle} = \sum_{n}^{4} \left\| K_n(L_n(I_t)^T L_n(I_t)) - L_n(I_{gt})^T L_n(I_{gt}) \right\|_1$$

Results

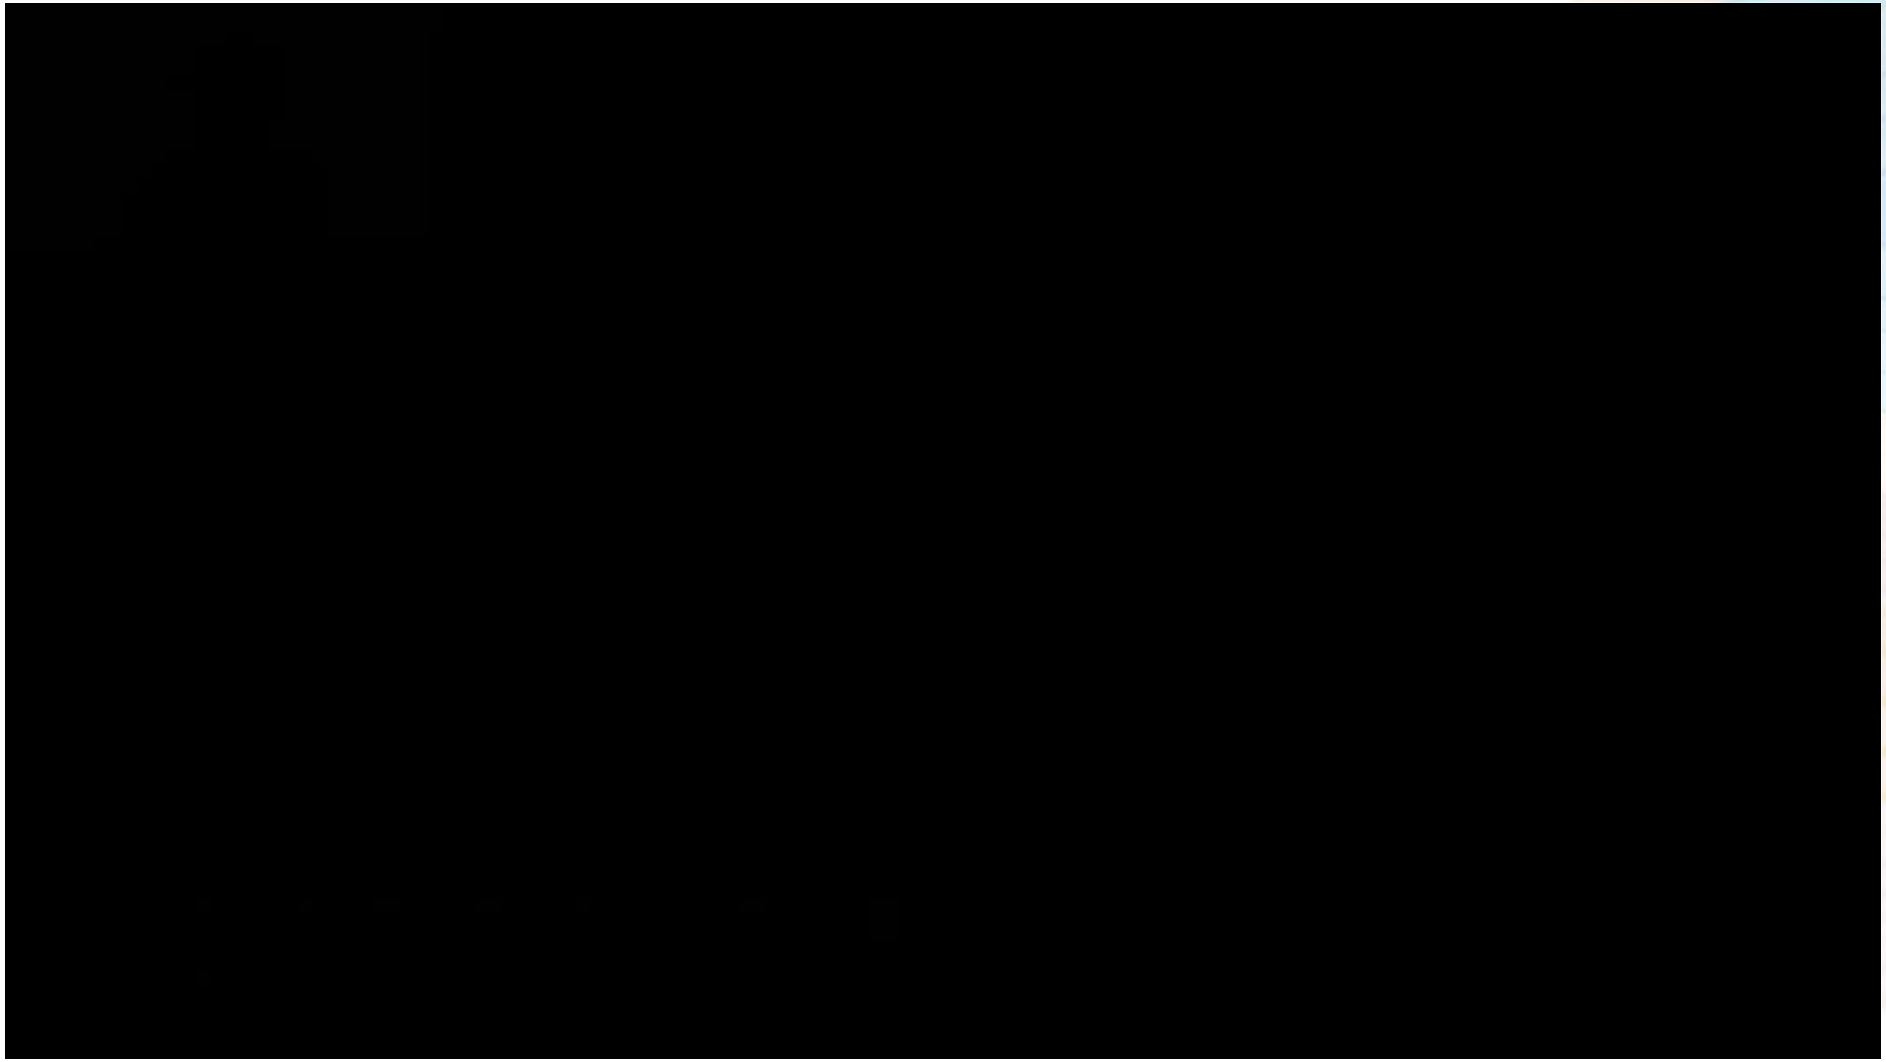

Limitations

- Latency and bandwidth issues in generating interactive frames in response to fast head motions
- Solution: use novel view synthesis as building blocks for computing intermediate representations fast to render

Layered depth representations

- **Key concept**: Each pixel is associated with multiple depth values
- Single panoramic images [Serrano et al, 2019]
- Create light field videos and layered meshes [Broxton et al., 2020]

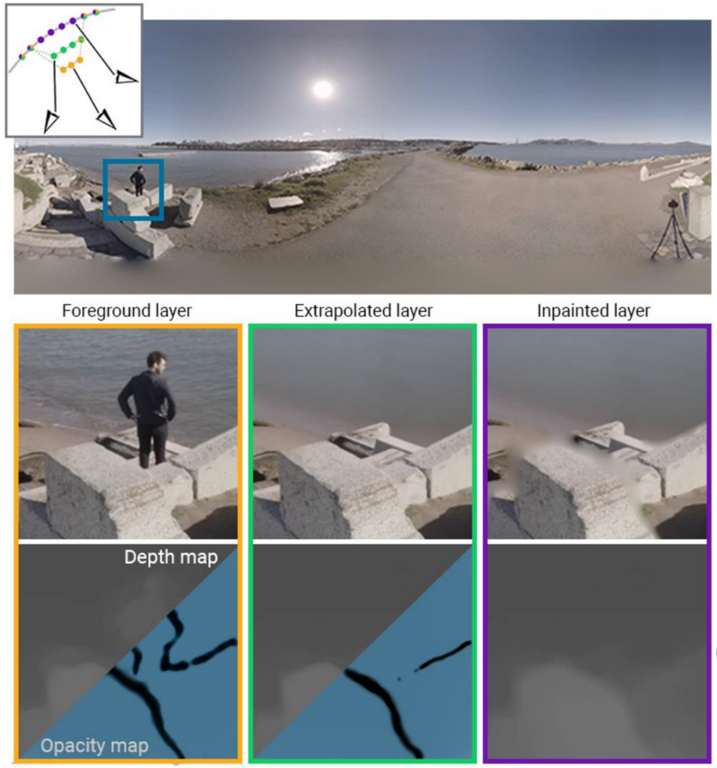

Serrano et al., IEEE TVCG 2019

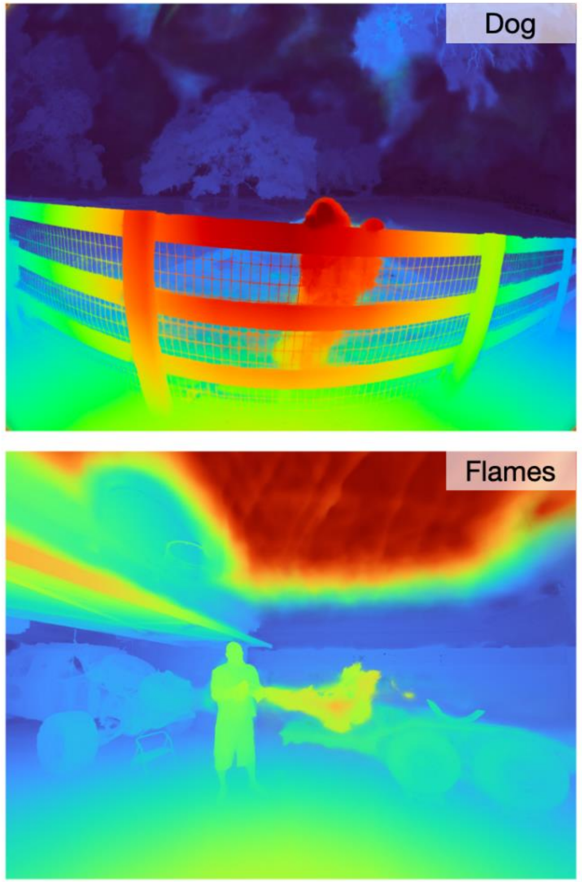

Broxton et al., ACM TOG 2020

Multi-planar Images

- **Key concept**: capturing the scene at multiple fixed depths and integrate during rendering
  - Multi-planar Image (MPI) [Tucker and Snavely, 2020]
- Extension to panoramic images by changing the proxy representation
  - Multi-Spherical Images (MSI) [Attal et al., 2020]
  - Multi-Cylinder Images (MCI) [Waidhofer et al, 2022]
- View-dependent rendering of depth images
  - [Marrinan and Papka, 2021]

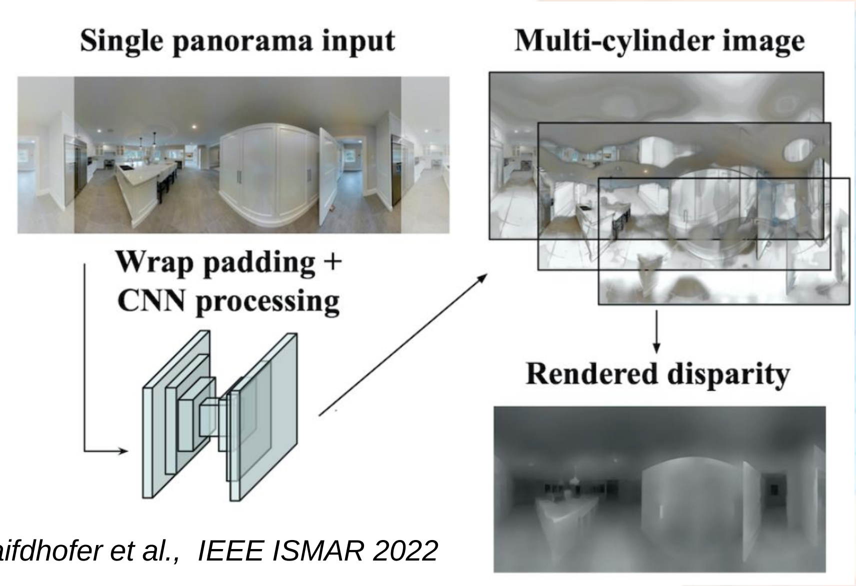

Waifdhofer et al., IEEE ISMAR 2022

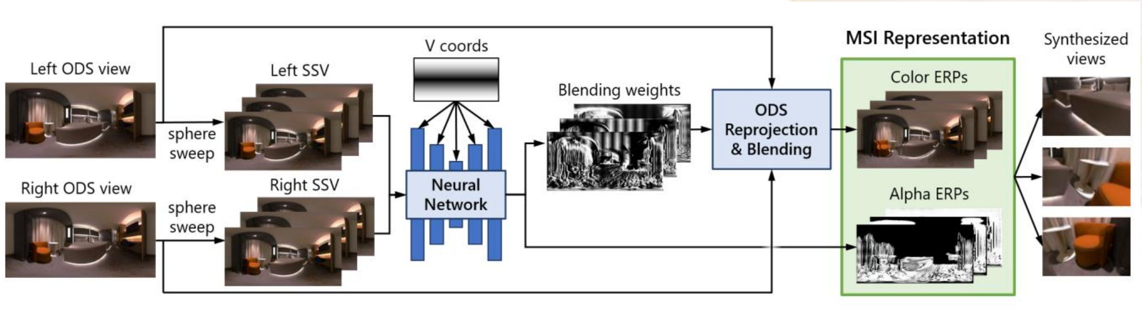

Attal et al., ECCV 2020

MCOP: Multiple-center-of-projection images for omnidirectional stereo

- **Key concept**: generate images by rotating around the head center and create a multi-perspective image by blending the viewing rays
  - [Original setup, also with physical cameras, Ishiguro et al 1992]
- Later: many different geometries and blending modes
- Pintore et al., 2024: use deep view synthesis networks to generate virtual views, blend them into omnidirectional images

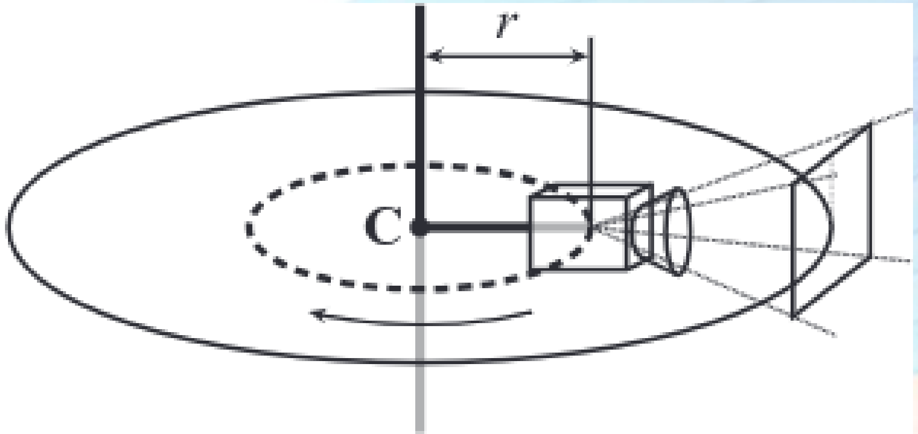

Ishiguro et al., Omnidirectional Stereo, TPAMI 1992

Pintore et al. Deep Pano Stereo, CAG 2024

MCOP generation: discrete set of panoramic slices

- Deep learning architecture for generating shifted views of indoor panoramic images
  - Estimation of depth, reprojection, and inpainting
- Similar PNVS network, based on a lightweight gated and dilated architecture
  - Novel photometric loss function + GAN for view synthesis
- Stereo exploration through precomputed MCOP image
  - Seamless stereo couples responding to the head motion
  - Web-compliant viewer using WebXR

PanoStereo Pipeline

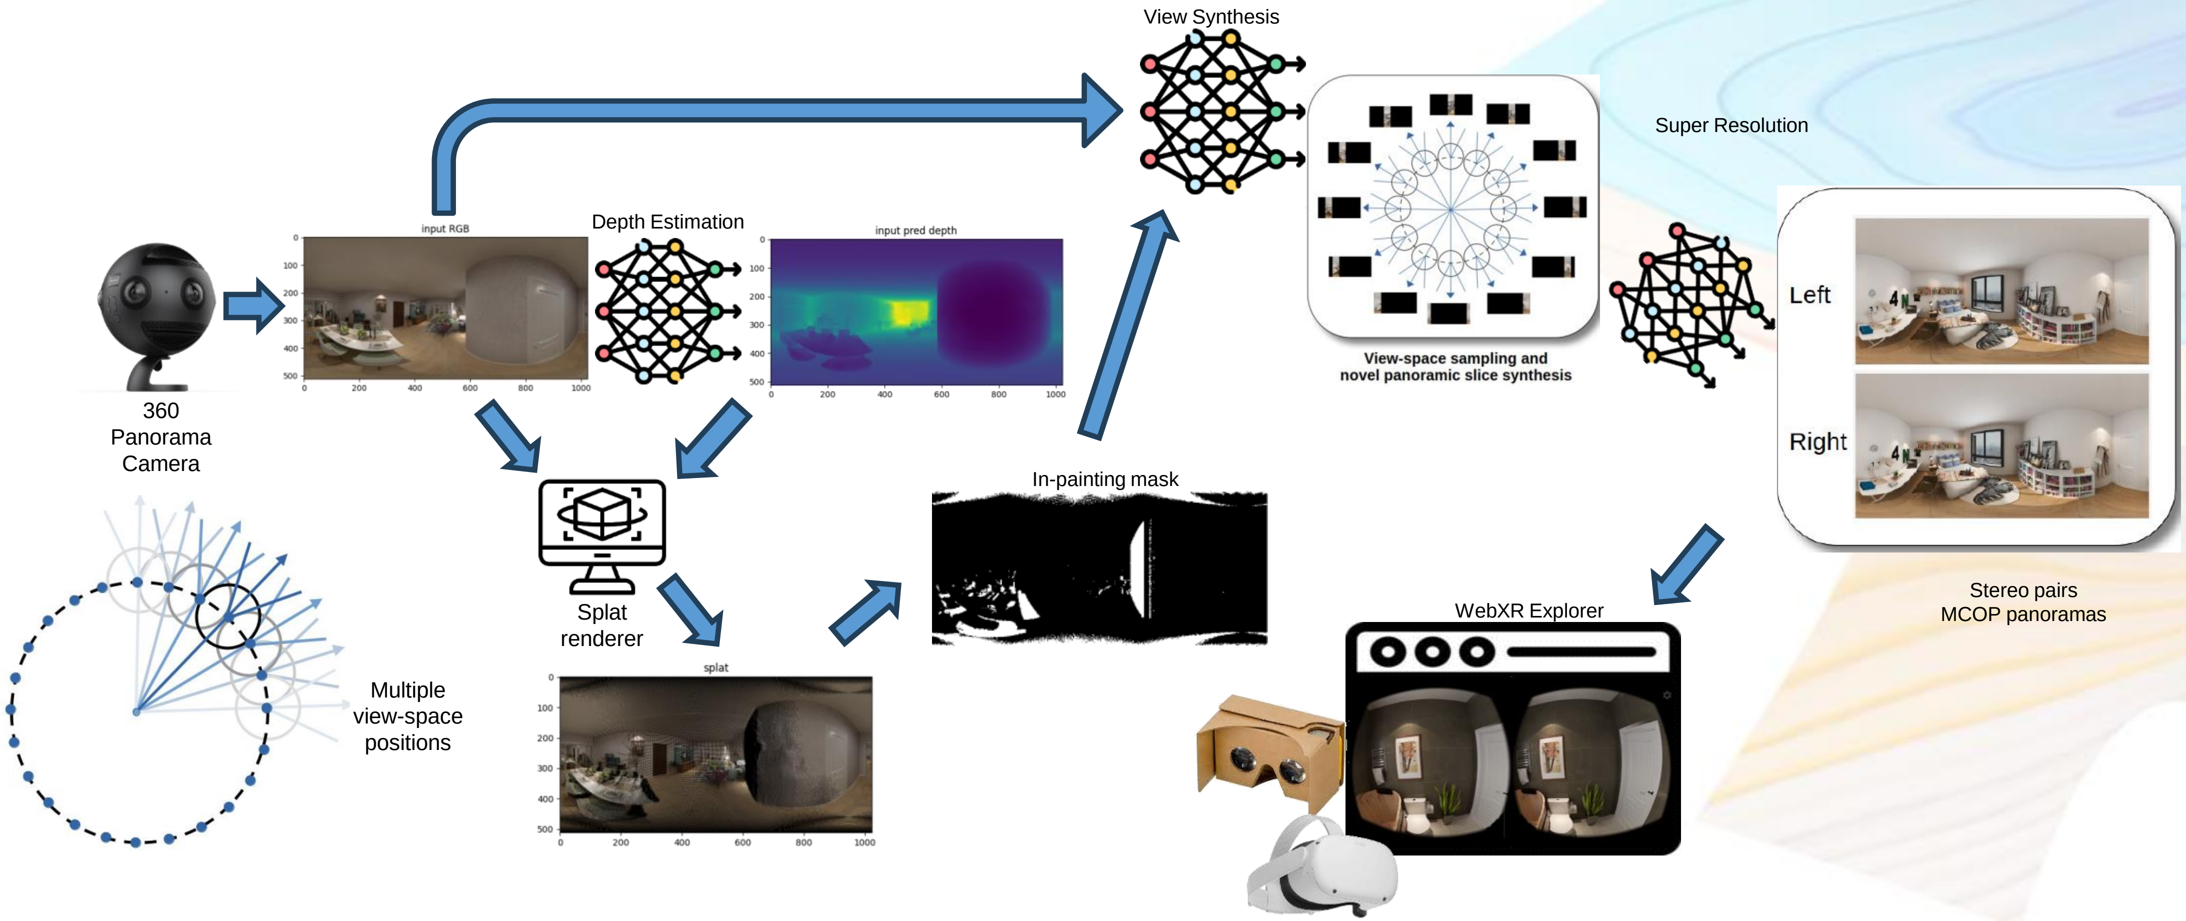

Pintore, Jaspe-Villanueva, Hadwiger, Schneider, Agus, Marton, Bettio, and Gobbetti. Deep synthesis and exploration of omnidirectional stereoscopic environments from a single surround-view panoramic image. Computers & Graphics, 119: 103907, March 2024. DOI: 10.1016/j.cag.2024.103907.

PanoStereo Pipeline

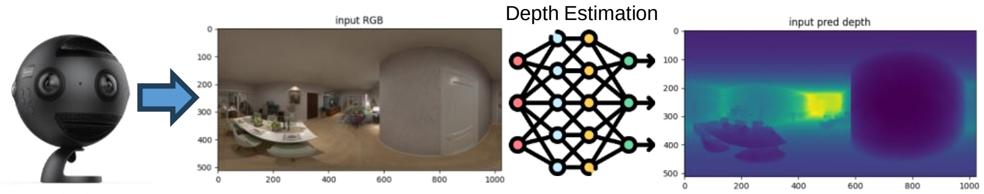

360
Panorama
Camera

Depth Estimation

- Many solutions available – see section on depth estimation
    - Pintore el al. 2024 uses a gated architecture, with an encoder-decoder scheme follows the same design adopted for view-synthesis, with adaptations to the specific task of spherical depth estimation
- Computed once for the central view, then used to support reprojection before view synthesis at each of the translated views

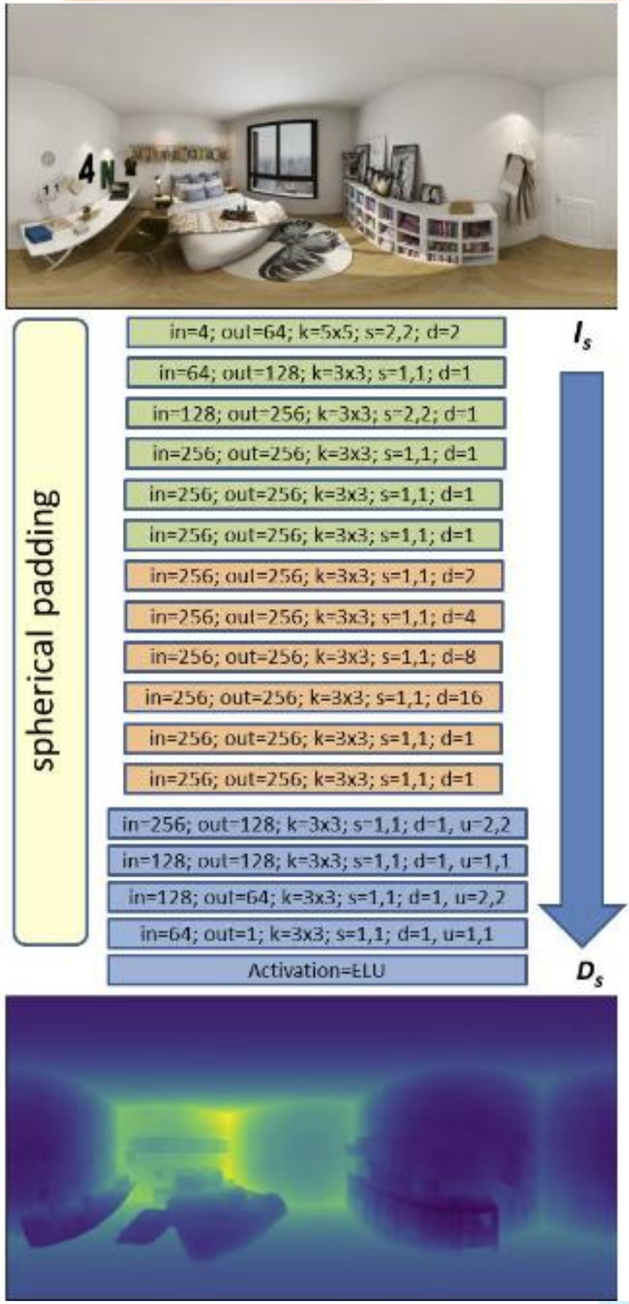

PanoStereo Pipeline

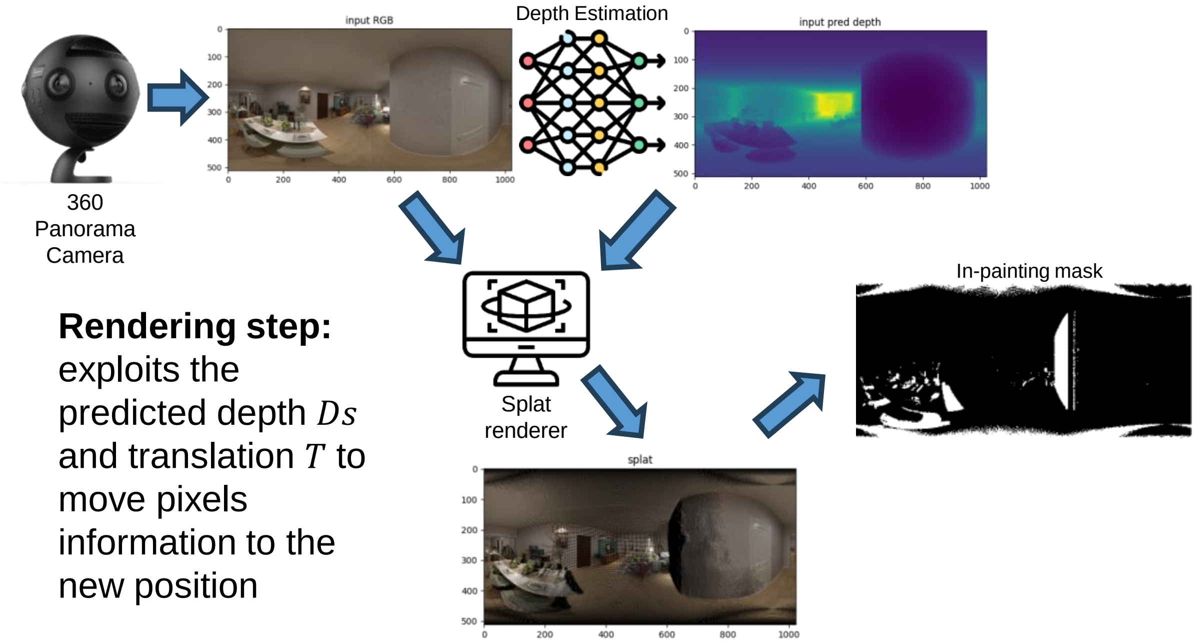

Rendering step:
exploits the predicted depth Ds and translation T to move pixels information to the new position

Generation of stereo view

- Stereo generation is “easy” for planar images (displacement in the eyes’ axis)
- Not applicable for panoramic images
- **Solution**
  - Simulation of head rotation
  - Multi-Center-Of-Projection
  - Blending of multiple thin slices

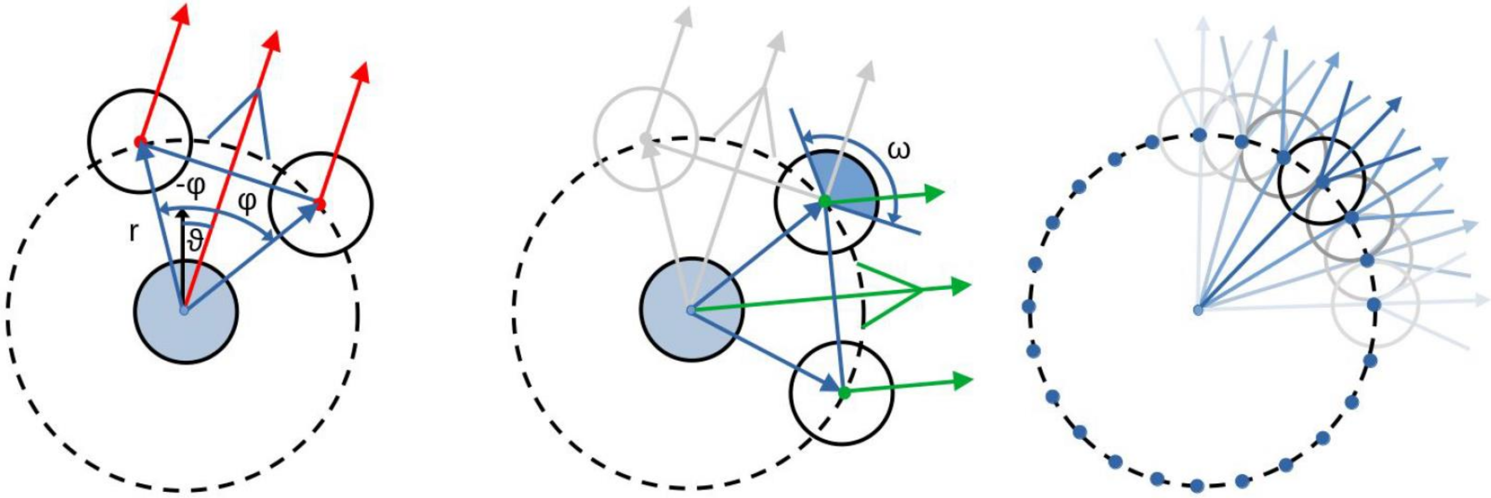

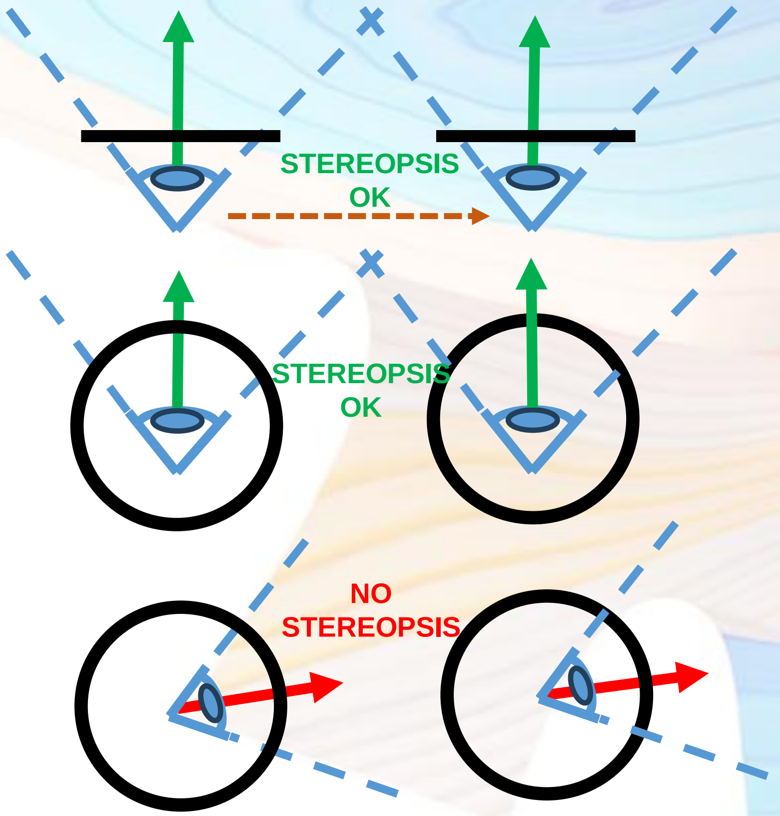

PanoStereo Pipeline

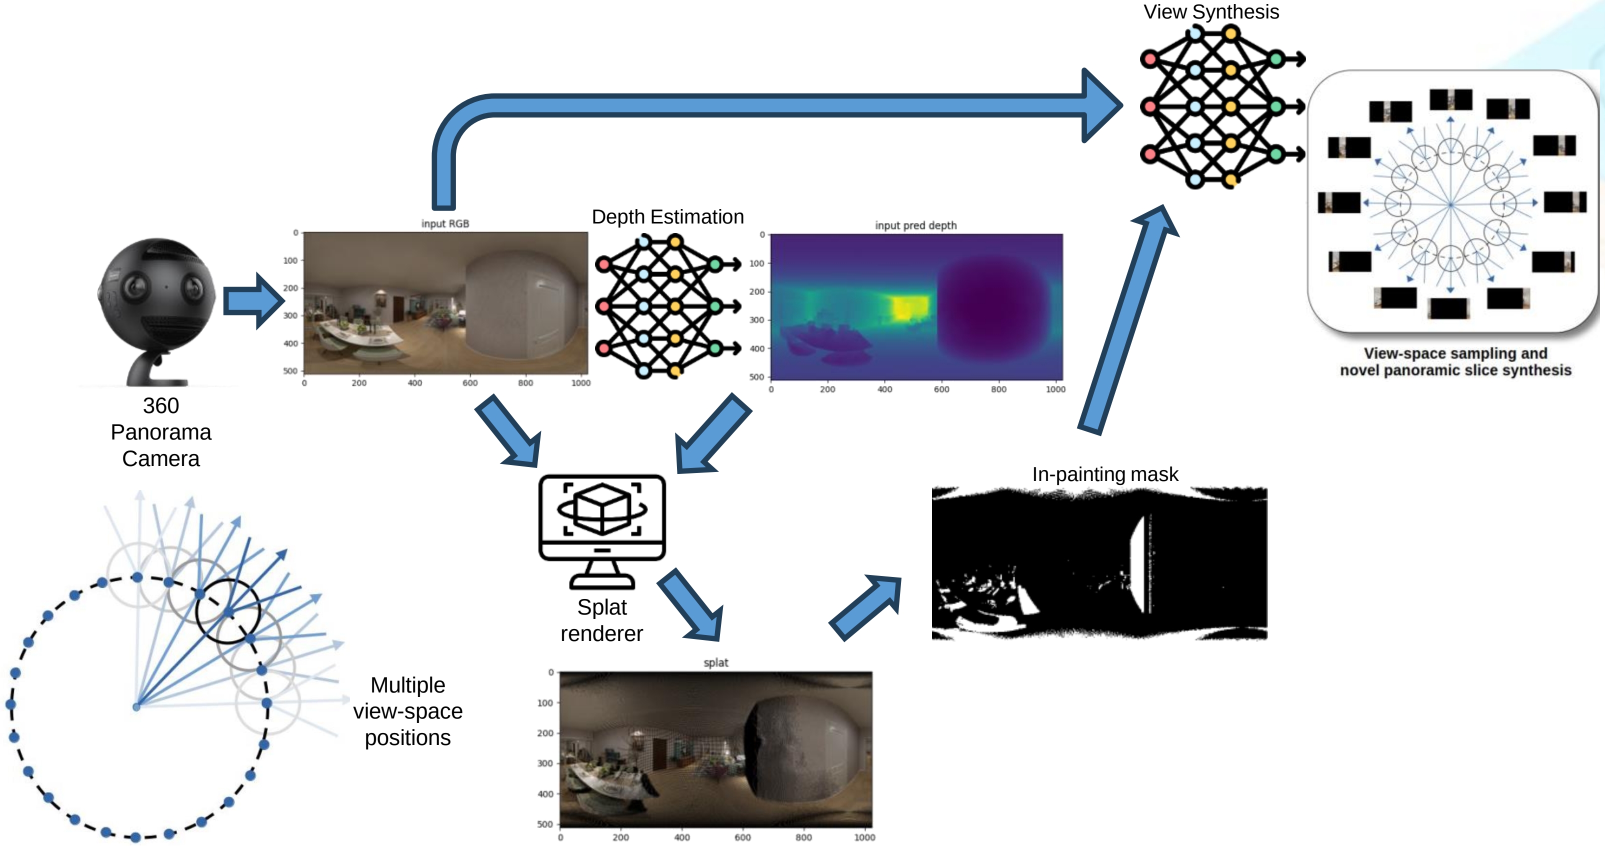

View synthesis

- Compute a novel, photorealistic, spherical image from a translated position
- Inpainting of the masked disoccluded holes
- Similar encoder-decoder architecture with perceptual and geometry losses
- Trained with synth dataset + **novel photometric loss function**
  - The synthetic image is back-projected to the initial position and compared with the original
  - Combination of L1 and SSIM (structure similarity)

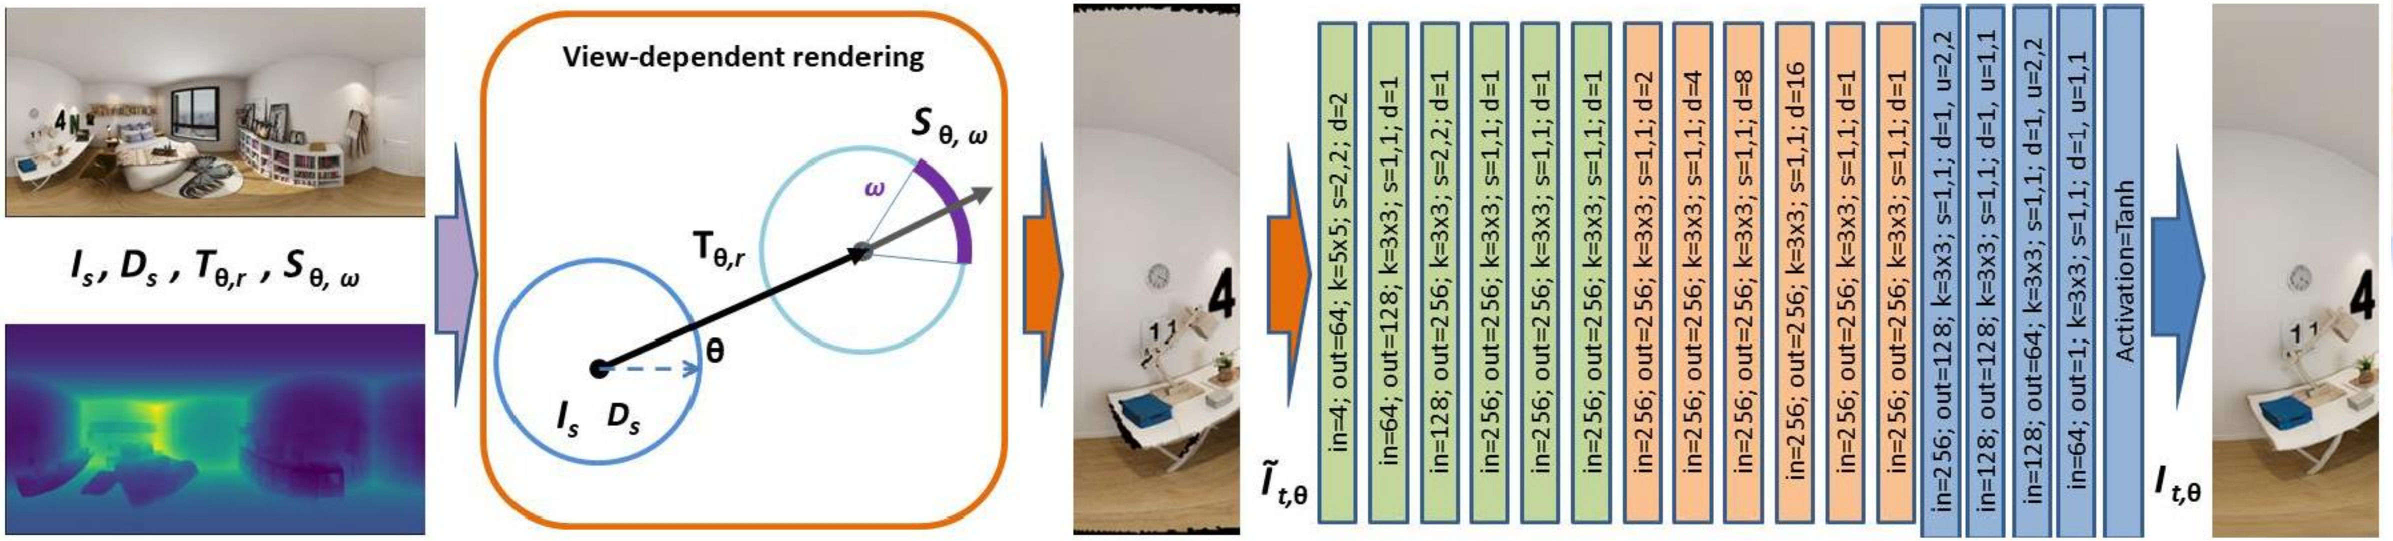

Scene composition (MCOP image)

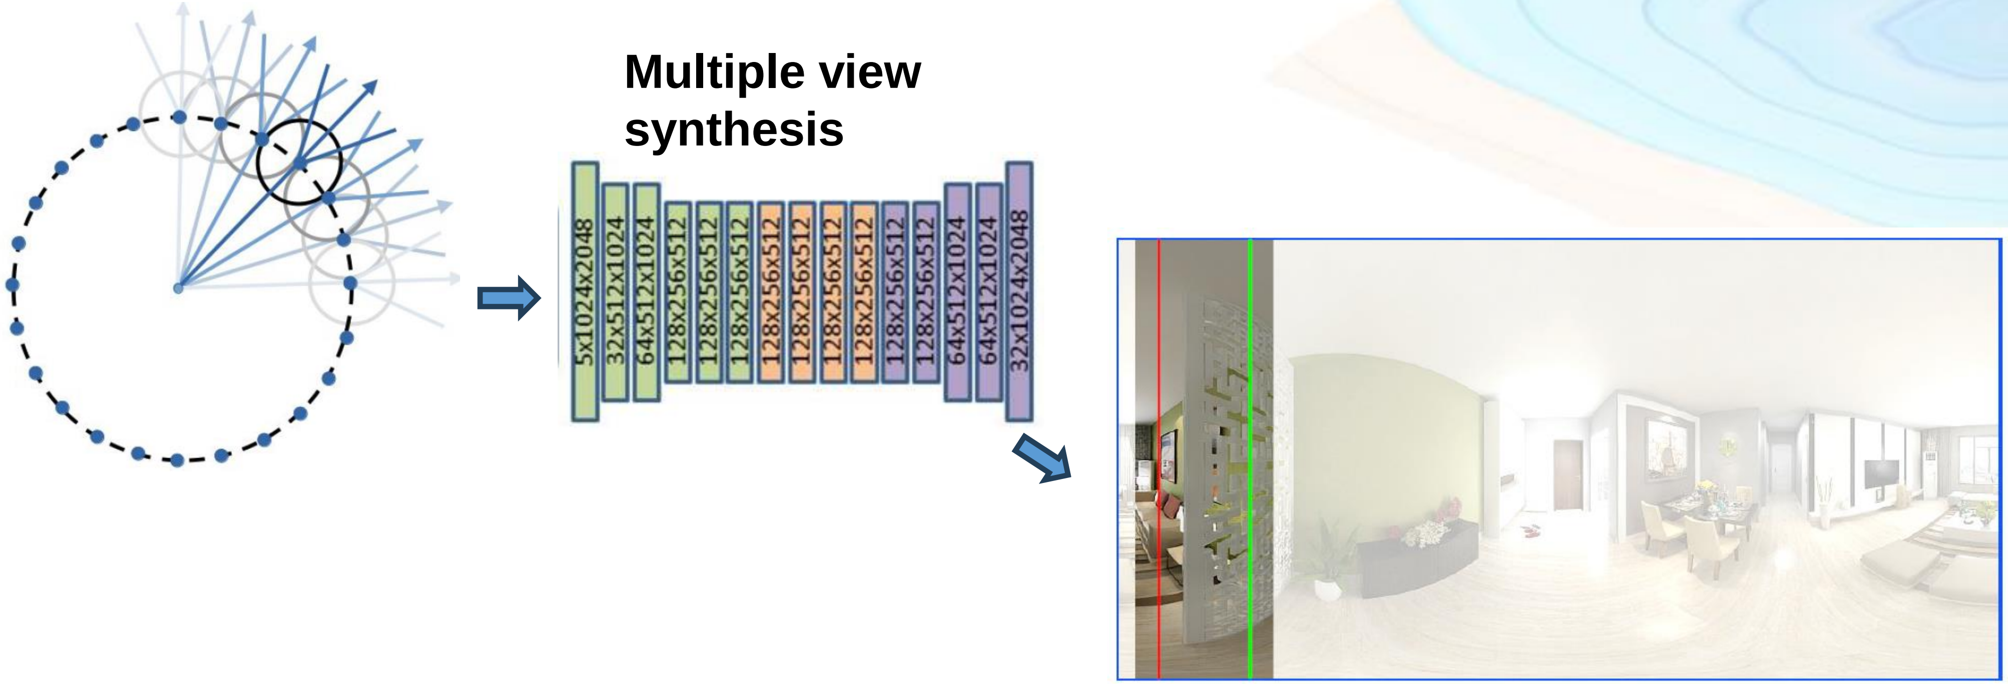

PanoStereo Pipeline

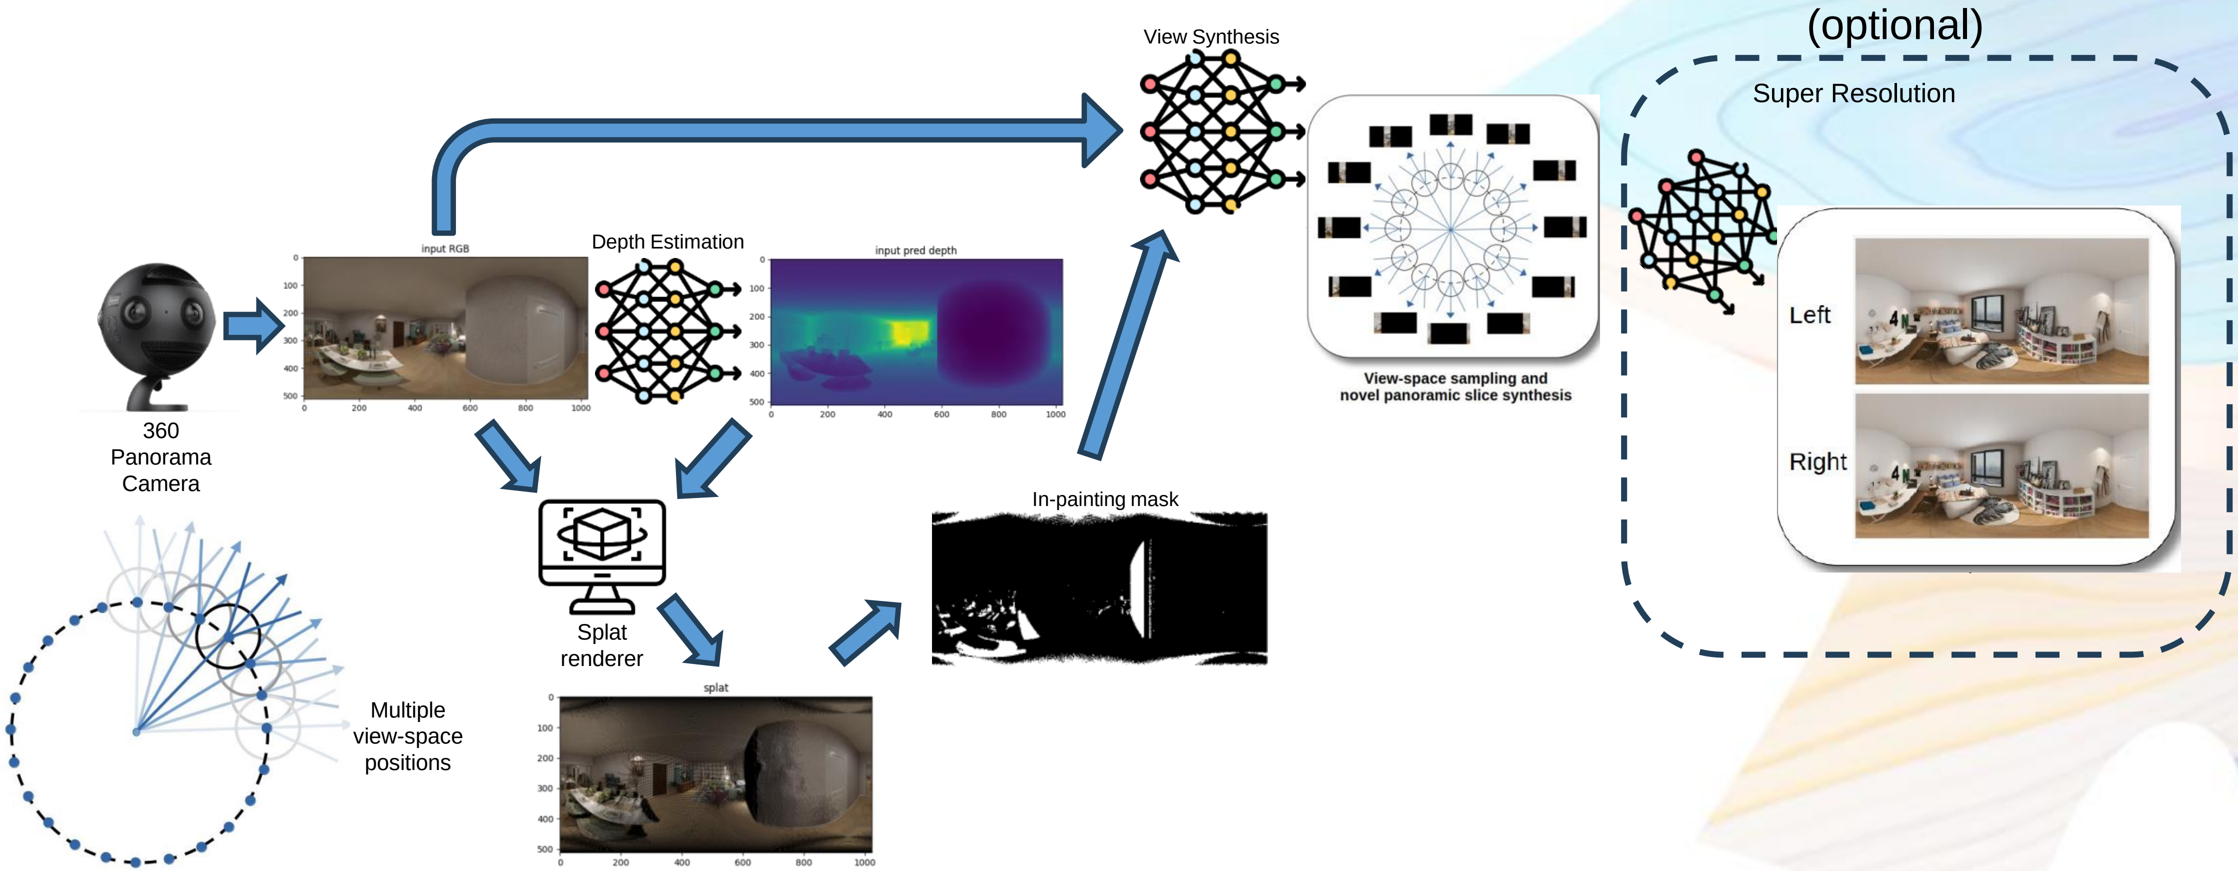

Results

Deep synthesis and exploration of omnidirectional stereoscopic environments from a single surround-view panoramic image

Giovanni Pintore (1,2), Alberto Jaspe-Villanueva (3),
Markus Hadwiger (3), Jens Schneider (4), Marco Agus (4),
Fabio Marton (1,2), Fabio Bettio (1),
Enrico Gobbetti (1,2)

(1) CRS4, Italy
(2) National Research Center in HPC, Big Data and Quantum Computing, Italy
(3) KAUST, Saudi Arabia
(4) HBKU, Qatar

Pintore, Jaspe-Villanueva, Hadwiger, Schneider, Agus, Marton, Bettio, and Gobbetti. Deep synthesis and exploration of omnidirectional stereoscopic environments from a single surround-view panoramic image. Computers & Graphics, 119: 103907, March 2024. DOI: 10.1016/j.cag.2024.103907.

Wrap-up of small displacement solutions

- Valid only on small area around the head
  - … lack of freedom
- Real-time synthesis may provide the best possible quality
  - … but is too slow to support hires rendering on current HMDs
- Precomputed representations with fast rendering allow for much higher resolution of individual images
  - … but work well only for limited image counts (e.g., stereo, or small set of multi-planer images)
  - So, more constraints (e.g., stereo only) or artefacts when moving out of sweet spot

Attal et al., ECCV 2020

6-DOF with large displacement

Large displacement methods

- Enable large displacements from the original capture position
  - Input: single panoramic image
  - Output: complete 3D renderable model
- Solutions categorized according to the need of prior knowledge
  - Known (inferred) scene semantic and structure
    - … which must be inferred from the image
    - Semantic and structure then guides 3D scene generation
  - Semantic-free approaches
    - Requires some way to synthesize plausible and consistent views even far from the original panorama
    - May exploit Neural Radiance Fields, 3D Gaussian Splats, or similar representations as intermediate or final renderable representation

Solutions exploiting additional knowledge

- They exploit the knowledge of layout of the room, and the type and pose of known objects
  - DeepPanoContext [Zhang et al., 2021]
  - PanoContextFormer [Dong et al., 2024]
- Limitations: they fail in real-world scenarios when they need to fill the scene with undefined/unrecognized semantics
  - Need for open vocabulary solutions
  - Diffusion models for text-to-3D generation (hot topic)

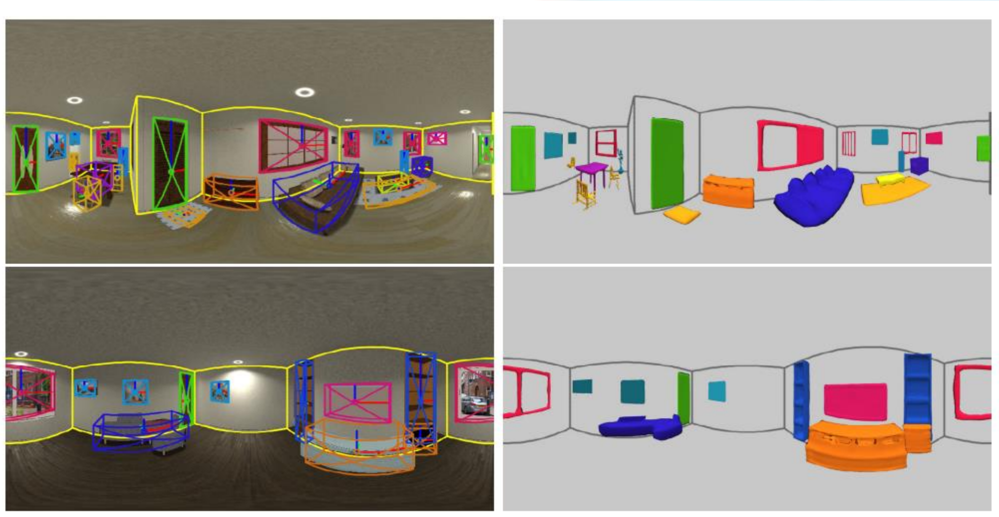

Zhang et al., IEEE ICCV 2021

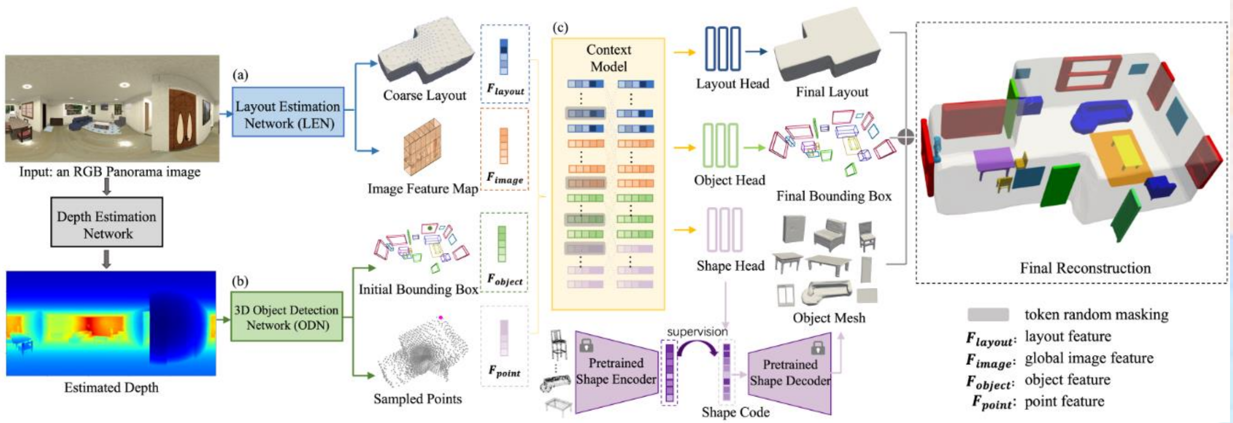

Dong et al., IEEE CVPR 2024

Semantic-free approaches

- They use novel view synthesis blocks described before to generate data for training Neural Radiance Fields of 3D Gaussian Splats
- NERF and 3DGS are optimized rendering representations able to perform blending and filtering in real-time
  - PERF [Wang et al., 2024]
  - PixelSplat [Charatan et al., 2024]
  - Pano2Room [Pu et al., 2024]

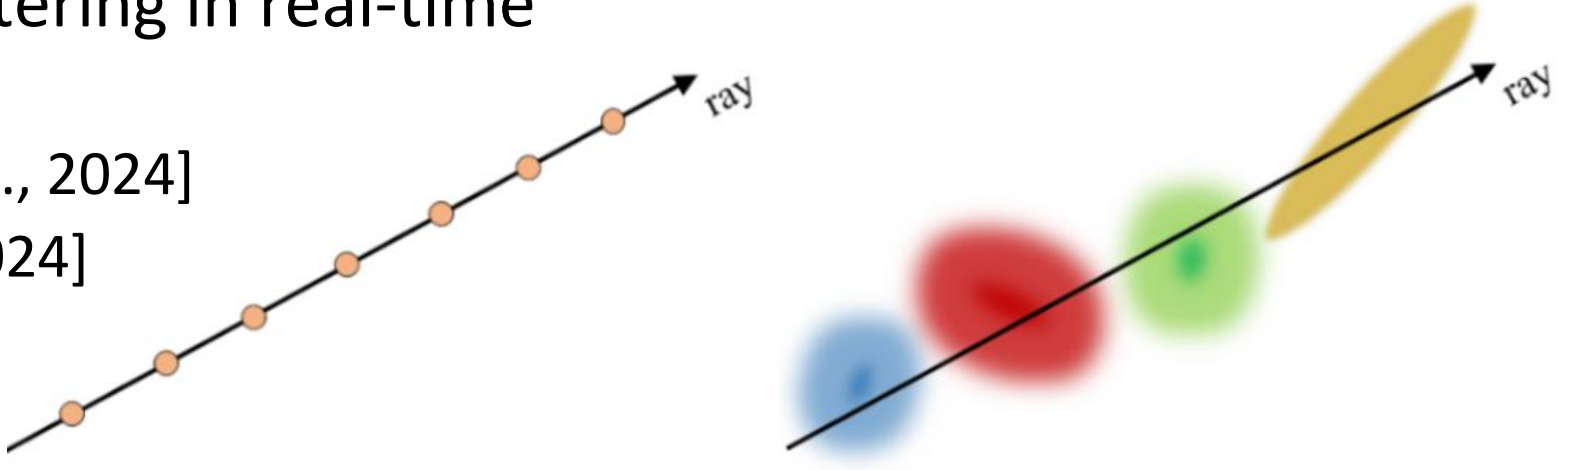

PERF

- Generate images with depth corresponding to novel views
  - [Wang et al., IEEE TPAMI, 2024]
- Perform collaborative RGB-D inpainting
- Limitations
  - Even if allowing for large motion, the camera needs to stay in predefined trajectories
  - Ghost geometries in undersampled areas
  - Heavy blurring in occluded areas

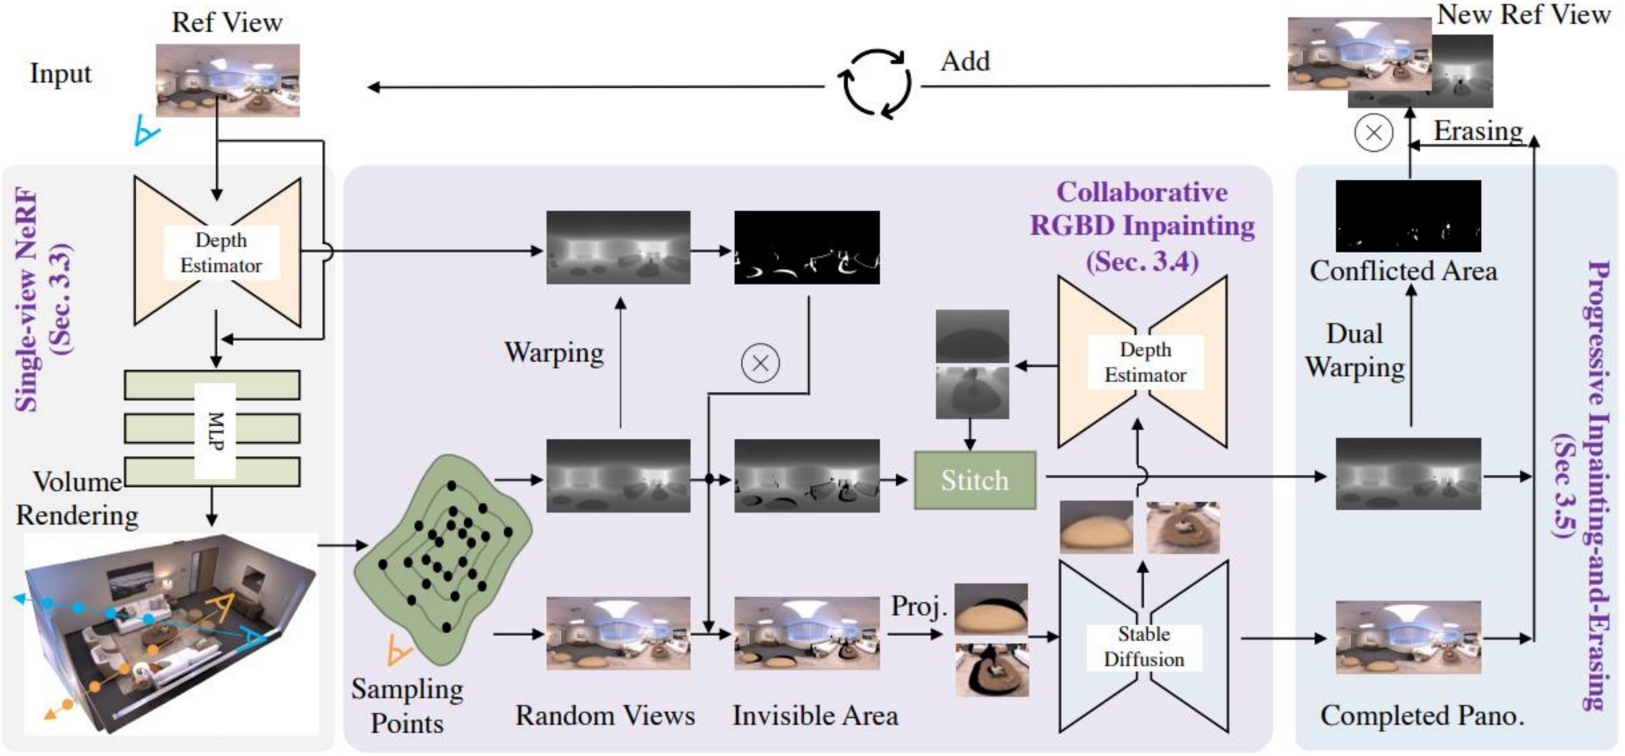

Wang et al., IEEE TPAMI 2024

Pano2Room

- Convert input panorama in a 3D mesh through depth estimation
  - [Pu et al., ACM TOG, 2025]
- Iteratively refine the mesh through a RGB-D inpainter to generate occluded color and geometry
- General incorporation and check for visibility conflicts at each step
- Inpainted mesh converted to 3D Gaussian Splats, trained on 3D consistent pseudo-novel views

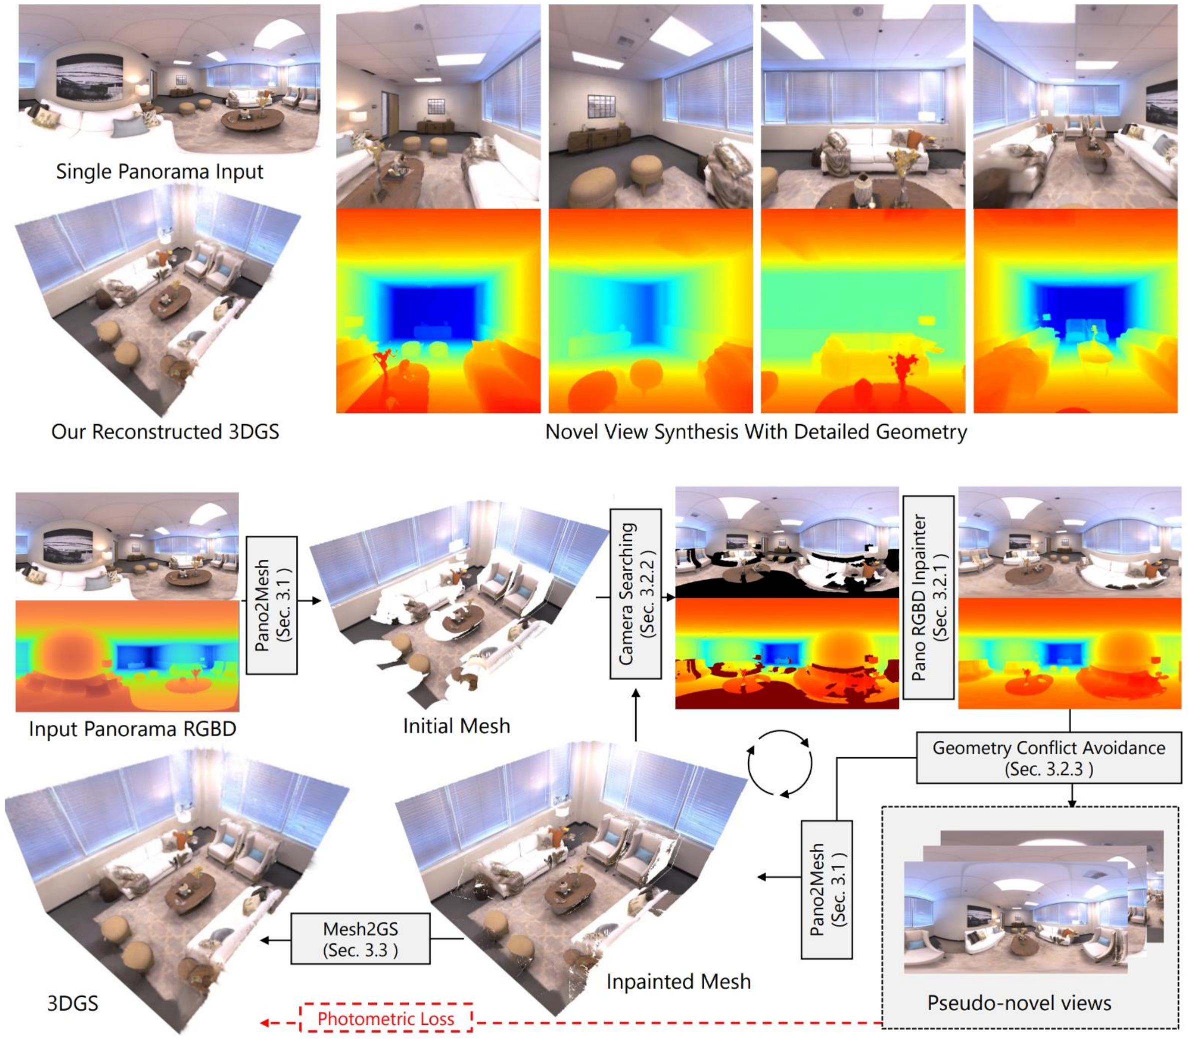

Pu et al., Siggraph Asia 2024

Recap

- AI-based technologies for performing immersive exploration and editing of scenes obtained through spherical imaging
  - Methods using different priors and assumptions
- Limitations:
  - Data hungry methods (rely on high-quality time-consuming data acquisition campaigns and processing)
    - We still mostly rely on synthetic datasets, like Structured3D
  - Resolution (most methods still work on 1024x512 input images!)
    - Partial workaround (usage of superresolution methods, like ESRGan or LAUNet)

Take-home messages

- The field is developing very fast
  - Thanks also to academic efforts
- Many challenges to address
  - Generalization to real-world scenarios
  - Integration with diffusion models
  - Open-vocabulary solutions
  - Increasing resolution
- Tech companies are investing huge resources
  - New solutions for XR
  - Automatic solutions for virtual staging

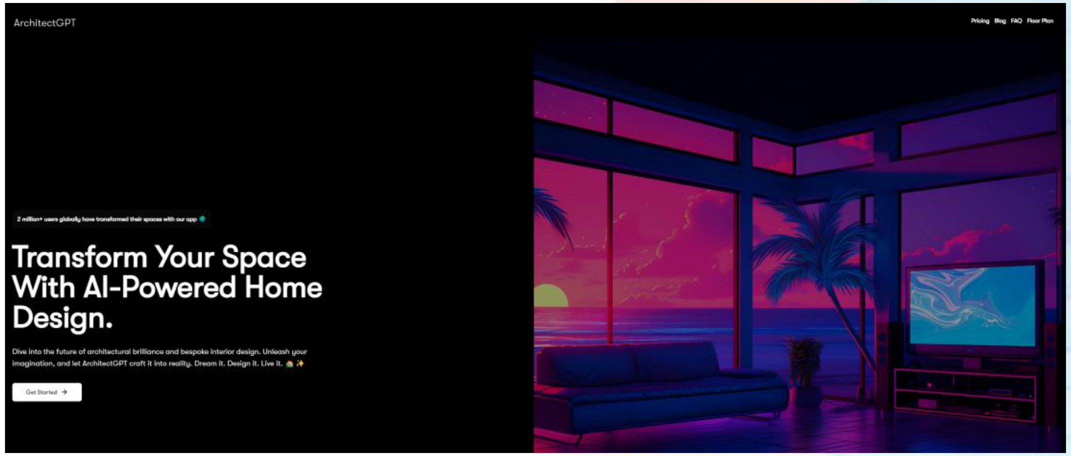

From ArchitectGPT.io, 2024

Apple VisionPro, 2024

Meta Quest Pro, 2024

Meta, Project ARIA, 2024

NEXT SESSION:
CLOSING