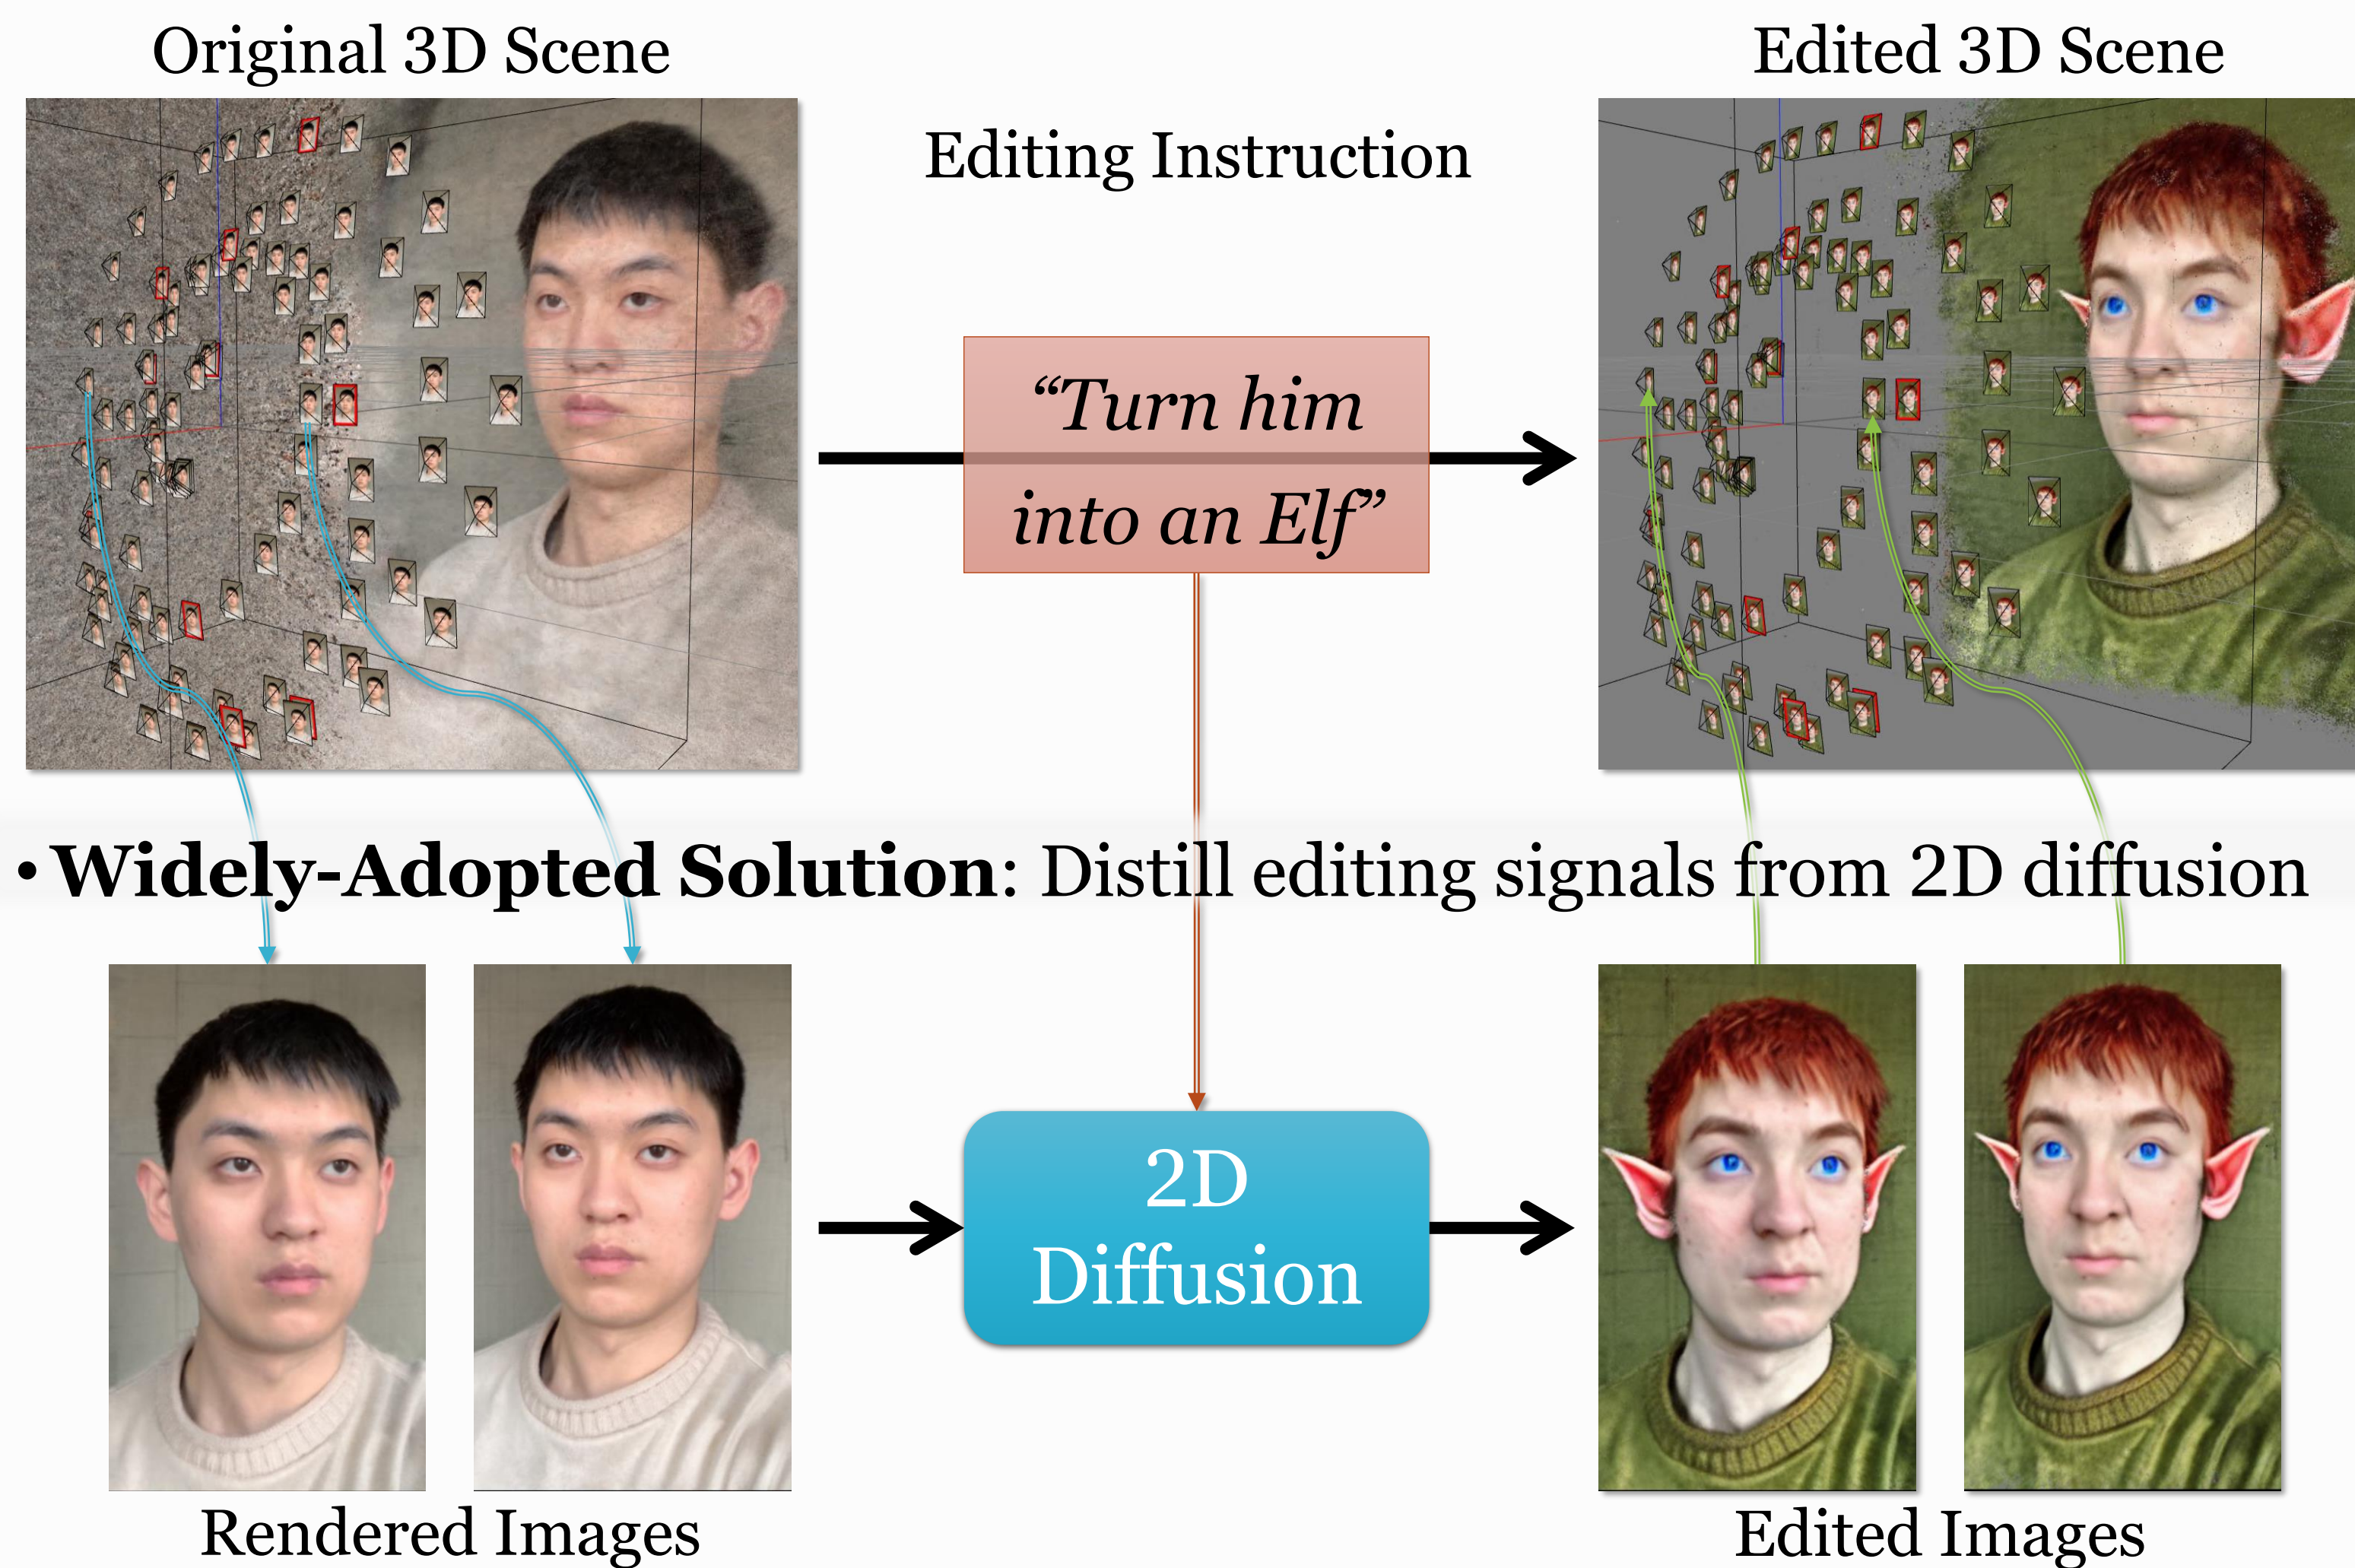



Task and Motivation

- **Task:** Instruction-guided 3D scene editing



- **Fundamental Limitation:** Lack of 3D consistency

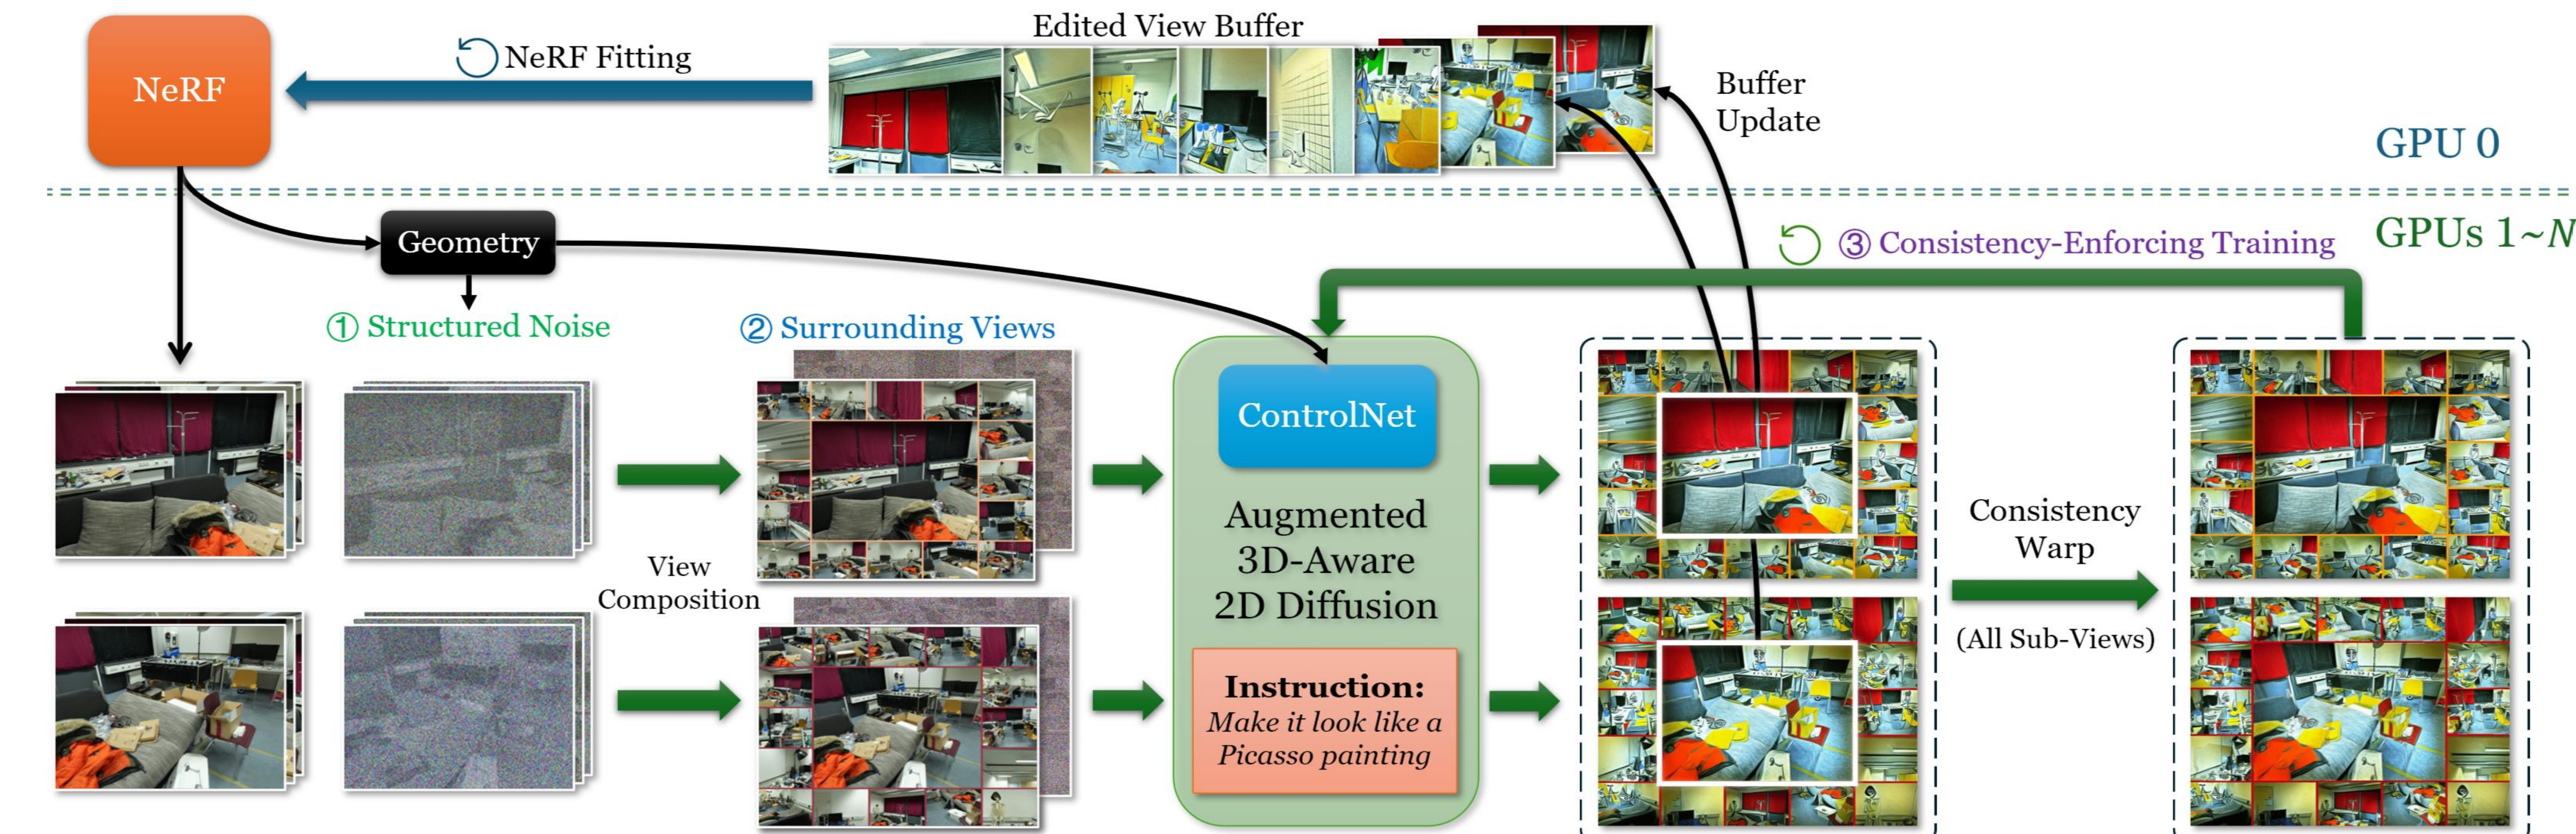
- 2D model acting **independently** on different views produces **inconsistent** results
- Converge towards an “**average**” of inconsistent results
- Common failure case:
Plaid pattern, *first solved* by our ConsistDreamer →



Key Insights: Three Synergistic Strategies



- ① Generate **3D-consistent structured noise** as foundation of consistency
- ② Compose **surrounding views** as 3D-context-rich input
- ③ Enforce **consistent denoising procedure** through concurrent **training**



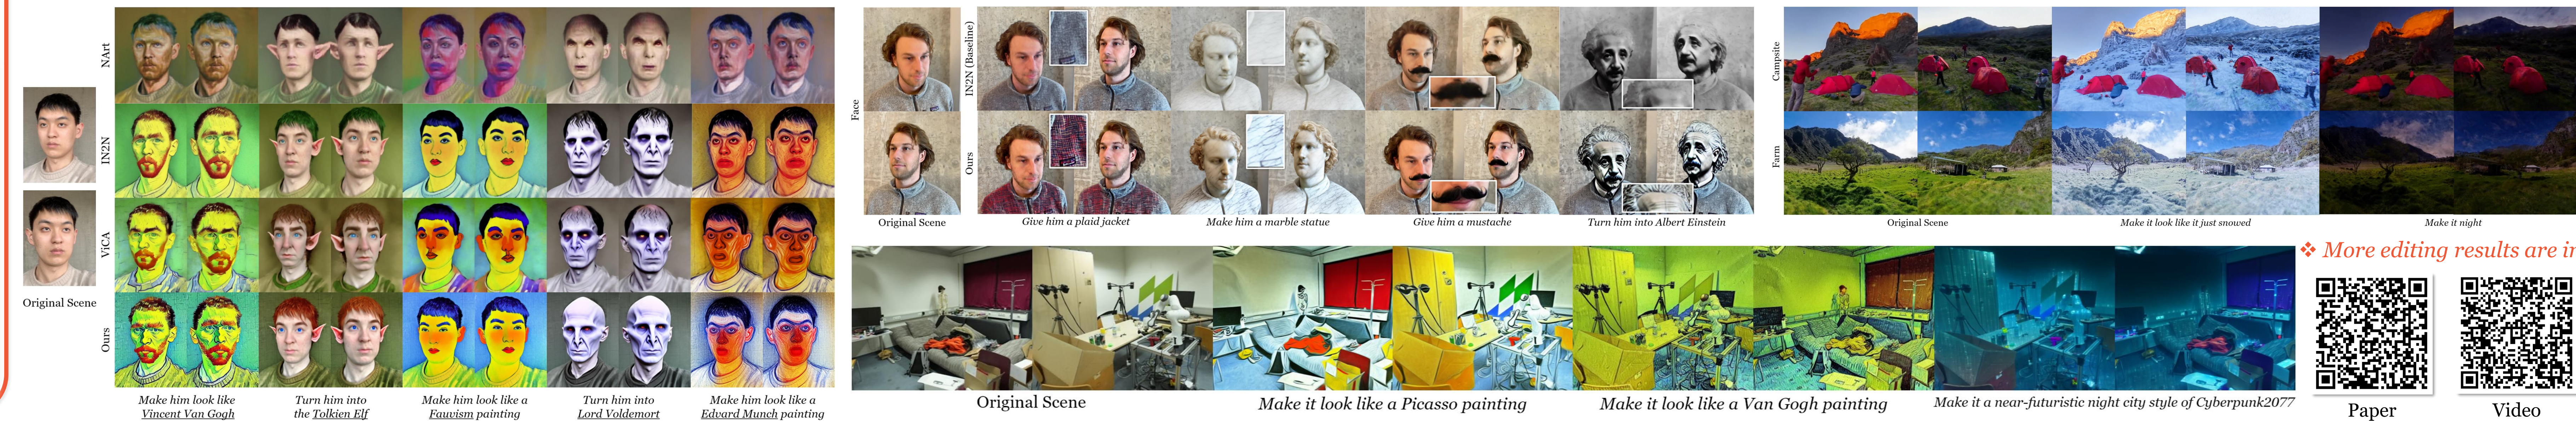
① 3D-Consistent Structured Noise

- Render 3D-consistent noise maps for each view
- Enable a 3D-consistent denoising procedure

② Surrounding Views

- Composed by surrounding a central view with 12 small reference views
- Provide rich context to the diffusion model and facilitates ③ training

Our ConsistDreamer generates high-quality edited scenes with **clear textures**, **bright colors**, and **high instruction fidelity**, achieving state-of-the-art scene editing results




❖ *More editing results are in*

Quantitative Evaluation & Ablation Study

Our Novel Metric:

Distillation Fidelity Score (DFS) ↓

- Quantify the fidelity of distillation
- FID ↓ between:
 - (1) *rendered images of edited scene*, and
 - (2) *edited images by diffusion*



Take-Aways

- **ConsistDreamer** enables 3D-consistent instruction-guided scene editing based on 2D-diffusion-distillation
- **ConsistDreamer**'s three synergistic strategies lift 2D diffusion models to generate 3D-consistent images
- **ConsistDreamer** achieves state-of-the-art editing results across various scenes and editing tasks