

Background & Motivation

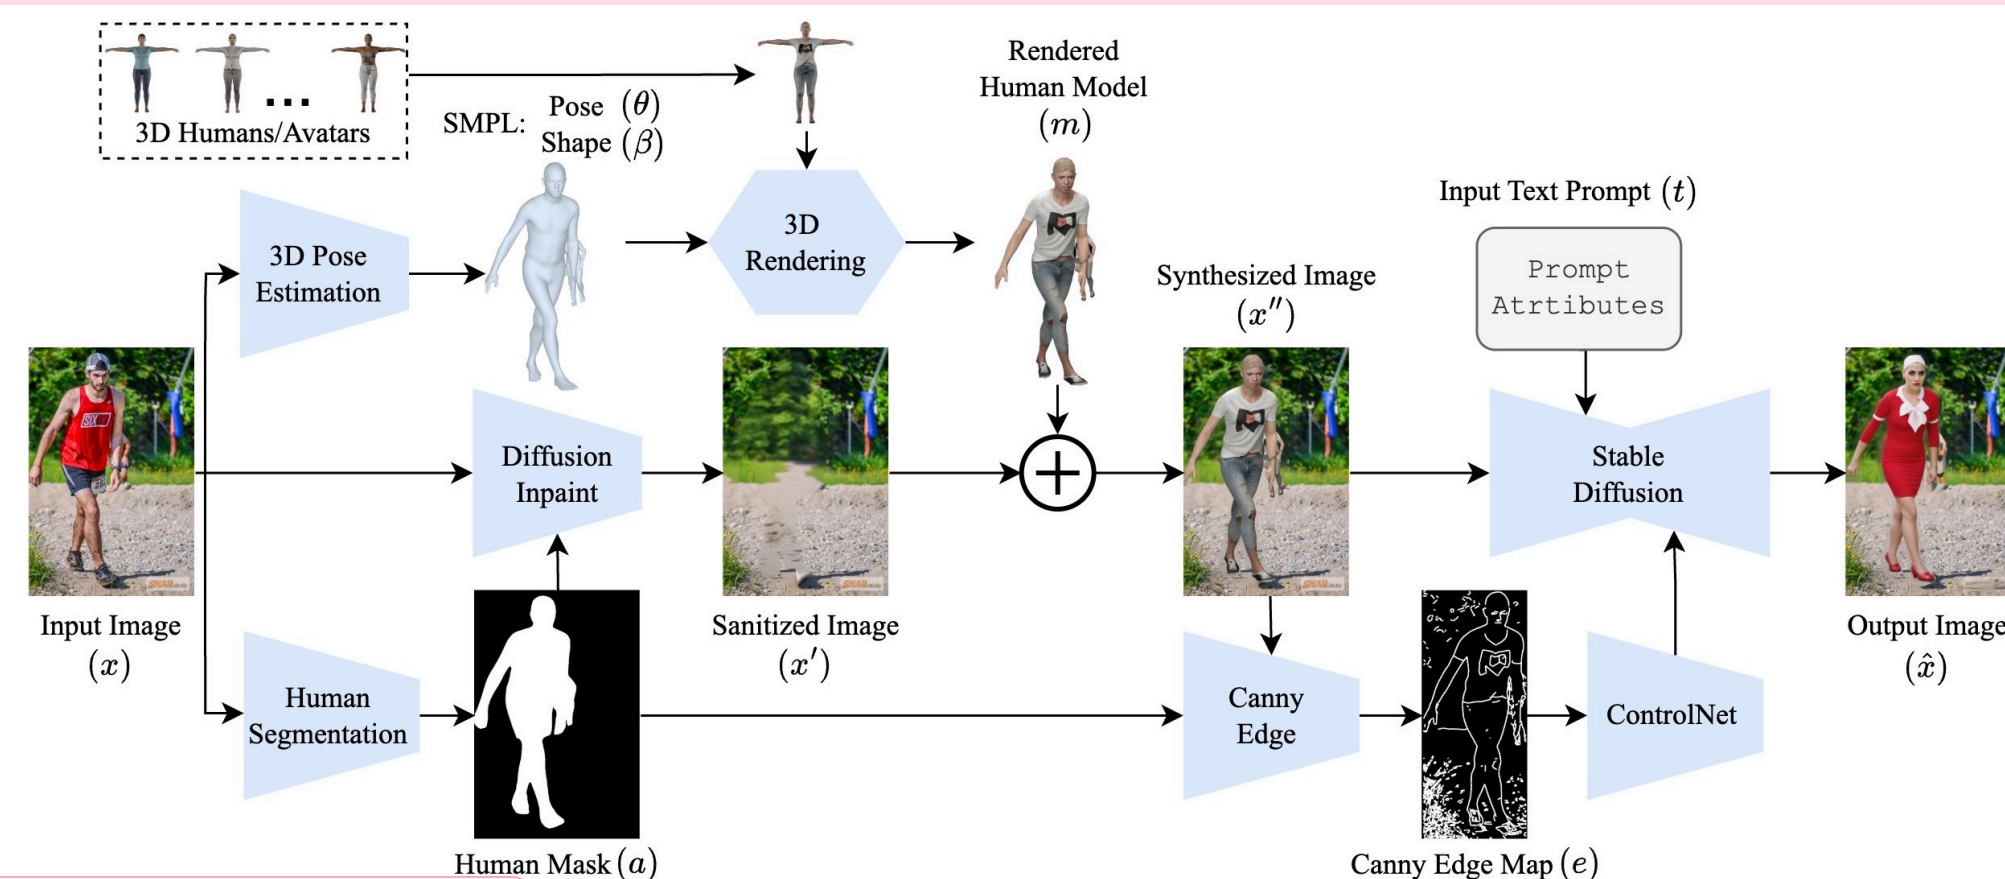
Privacy regulations (GDPR, CCPA) mandate public datasets with permissive licenses (e.g., CC BY 4.0) containing humans be **pseudonymized** before use in machine learning pipelines.

Existing anonymization methods fall short:

- **Masking & blurring:** removes identity but destroys pose and degrades downstream utility
- **GAN-based synthesis:** limited controllability, morphed outputs, poor photorealism
- **Diffusion editing:** risks retaining identifiable texture features from the original subject

We propose **Rendering Refined Stable Diffusion (RefSD)**: a pipeline for privacy-compliant in-place human pseudonymization that preserves scene context and pose.

Rendering-Refined Stable Diffusion (RefSD)



1. Human Removal

Detect & segment human subjects. Remove the original human. Inpaint the background.

2. 3D Avatar Rendering

Extract 3D pose using SMPL. Render a synthetic, identity-free avatar. Composite rendered avatar into scene.

3. Diffusion Refinement

Textured avatar as structural guidance. Text prompts to control attributes. Refine with Stable Diffusion.

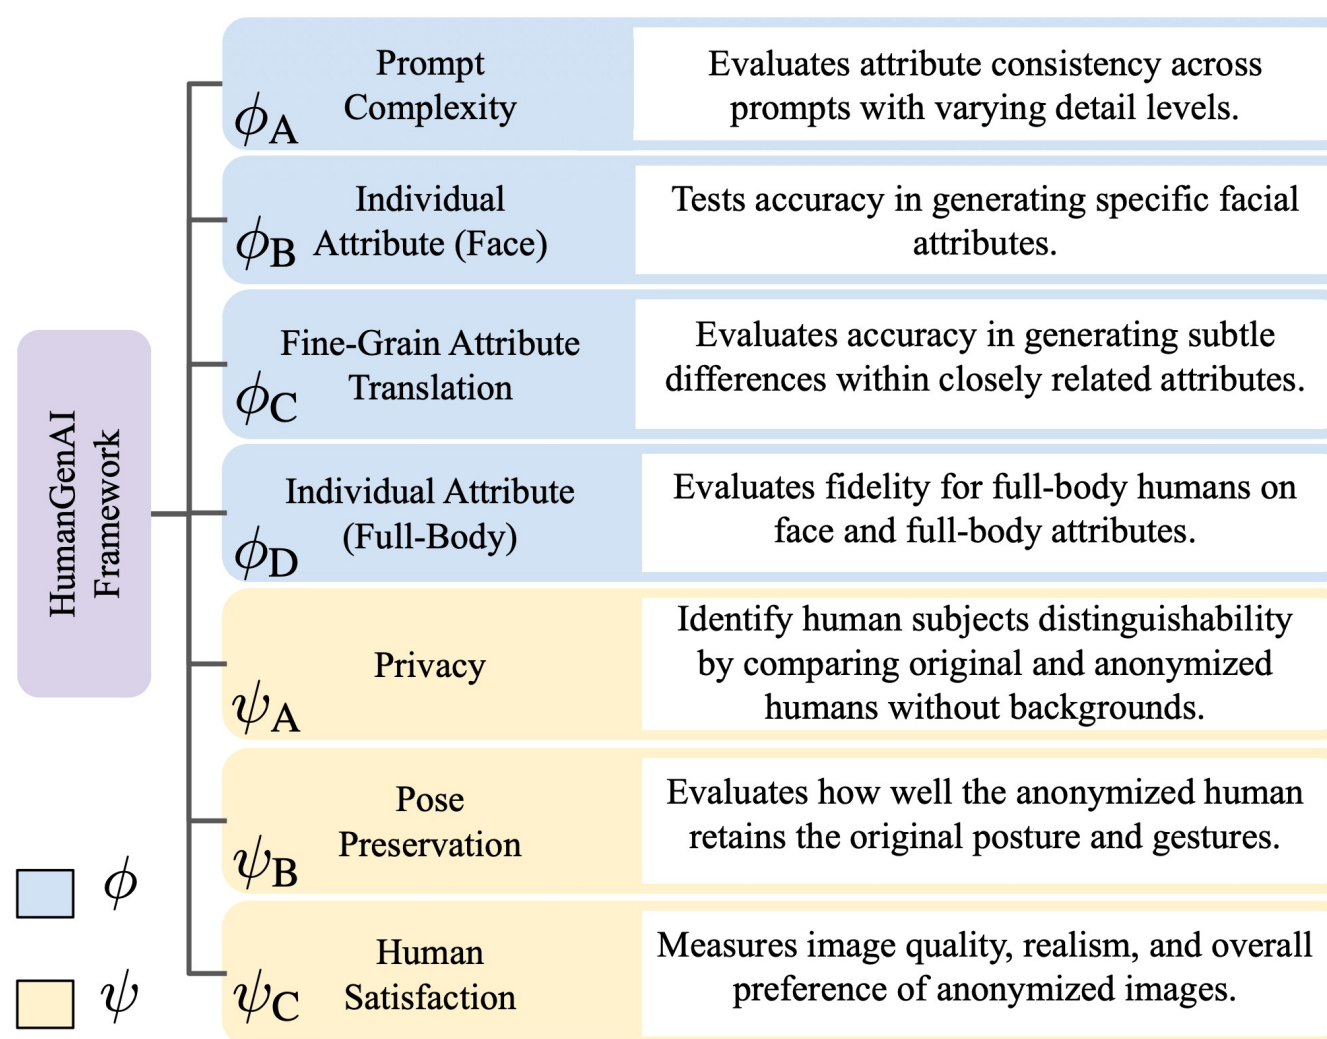
Qualitative Results



HumanGenAI: A New Evaluation


We introduce **HumanGenAI** to measure whether generated humans are realistic, privacy-preserving, and attribute-aligned. It evaluates two major aspects:

1. **Generated Attribute Fidelity (ϕ):** Tests whether generated humans match prompt-specified attributes.
2. **Generic Property Assurance (ψ):** Measures privacy and visual quality using human judgments.

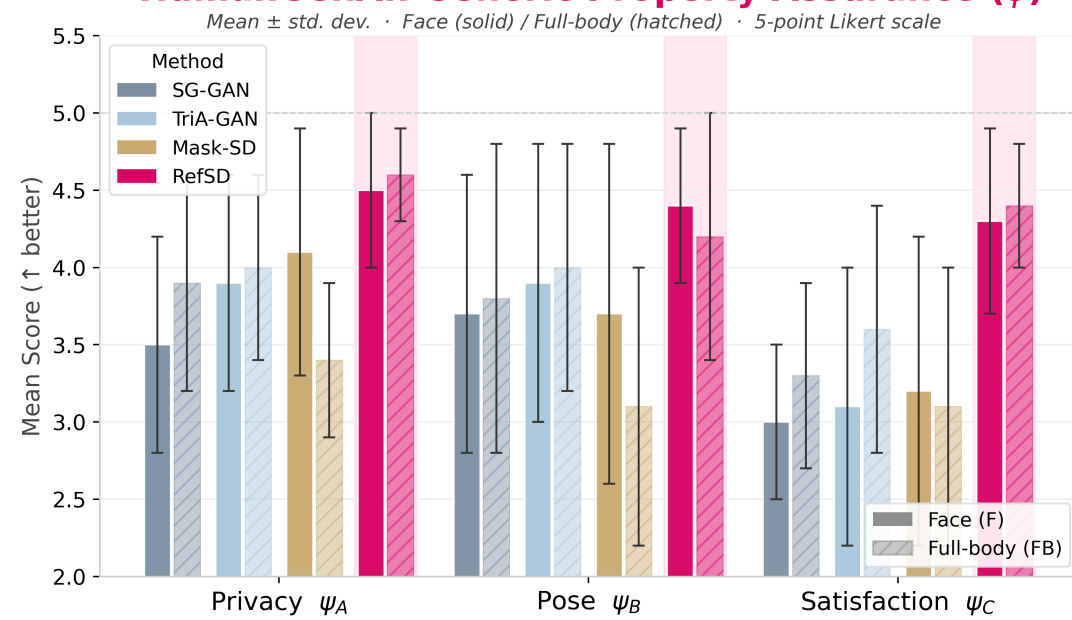


Experiments & Results

HumanGenAI: Fine Grain Attribute Translation (ϕ_C)



HumanGenAI: Generic Property Assurance (ψ)



Privacy: Person Re-Identification

Method	R@1 ↓	R@5 ↓	R@10 ↓	mAP ↓
Original	93.3	97.6	98.5	82.8
WeakBlur	39.2	49.1	54.0	33.7
StrongBlur	31.5	38.4	41.5	28.1
Mask-SD	28.4	32.0	33.8	24.6
FADM +	67.8	-	-	53.7
DP2: SG-GAN	40.3	46.3	49.2	34.5
DP2: TriA-GAN	40.1	44.6	46.5	35.1
MaskOut *	27.4	29.7	30.7	23.8
RefSD (Ours)	27.5	29.9	31.0	23.9


Utility: Classification

Config	Emotion	Age	Gender	Ethnicity
S (synth)	36.3 / 39.6	48.2 / 48.4	53.1 / 52.9	67.6 / 68.2
R (real only)	41.5 / 41.5	58.4 / 57.0	61.9 / 60.6	78.2 / 77.5
S+R (joint)	44.3 / 42.0	59.9 / 58.5	73.0 / 63.4	79.9 / 77.5
S-R (pre+fine)	45.3 / 42.2	58.1 / 55.7	64.4 / 65.1	78.8 / 77.6

S (synth) = pseudonymized RefSD images. S→R: pre-train on synthetic, fine-tune on real.

Note: See paper for remaining results.

Failure Cases



Conclusion

- ✓ RefSD achieves re-ID on par with complete masking.
- ✓ Pseudonymized images complement or exceed real data in downstream training.
- ✓ Modular design supports plug-in upgrades to renderer and diffusion models.
- ✓ HumanGenAI provides a new standardized benchmark for pseudonymization research.