



Selective Computation for Reliable and Scalable Vision Systems

Allen Tu · Ph.D. Student, Computer Science
University of Maryland, College Park



Thesis

Reliable perception emerges when computation, data, and modeling capacity are **selectively applied** where they improve performance.

My dissertation studies **selective computation** as a **unifying principle for efficiency and robustness** through three decisions:

- 1. What to Compute
- 2. What to Trust
- 3. Where to Generate

About Me



Advisor: Prof. Tom Goldstein

Collaborators: Profs. Matthias Zwicker, Vishal M. Patel, Rama Chellappa

Experience: UMD · STR · Amazon Science

CV & Publications: tualen.github.io

Contact: atu1@umd.edu

Research Areas: Reliable and Scalable Vision Systems
3D/4D Reconstruction · Biometrics · Generative Modeling

Research Programs:

- **IARPA BRIAR:** Biometric recognition from incomplete face, body, and gait information in severe operational conditions.
- **IARPA WRIVA:** Unconstrained 3D reconstruction and novel view synthesis in challenging real-world environments.

Community: SPAR-3D Workshop Organizer, CVPR 2026

Ongoing Research and Feedback


How can selective computation make generative modeling and 3D/4D reconstruction both efficient and robust?

- **Scalable Diffusion Reference Conditioning:** Use richer reference context to improve cross-view consistency.
- **Uncertainty-Aware Diffusion Priors:** Guide sparse-view reconstruction where geometry or appearance is uncertain.
- **3D/4D Reconstruction for Physical AI:** Build scalable reconstruction for embodied AI in dynamic environments.
- **Adaptive Modeling:** Design new 3D/4D representations for efficient, high-fidelity, and capable reconstruction.

1. What to Compute

PUP 3D-GS: Principled Uncertainty Pruning for 3D Gaussian Splatting


CVPR 2025 · pup3dgs.github.io



Prune 90% of primitives from any pretrained 3D Gaussian Splatting model using a mathematically principled sensitivity score while retaining more salient foreground information and higher visual fidelity than previous techniques at a substantially higher compression ratio.

Speedy-Splat: Fast 3D Gaussian Splatting with Sparse Pixels and Sparse Primitives

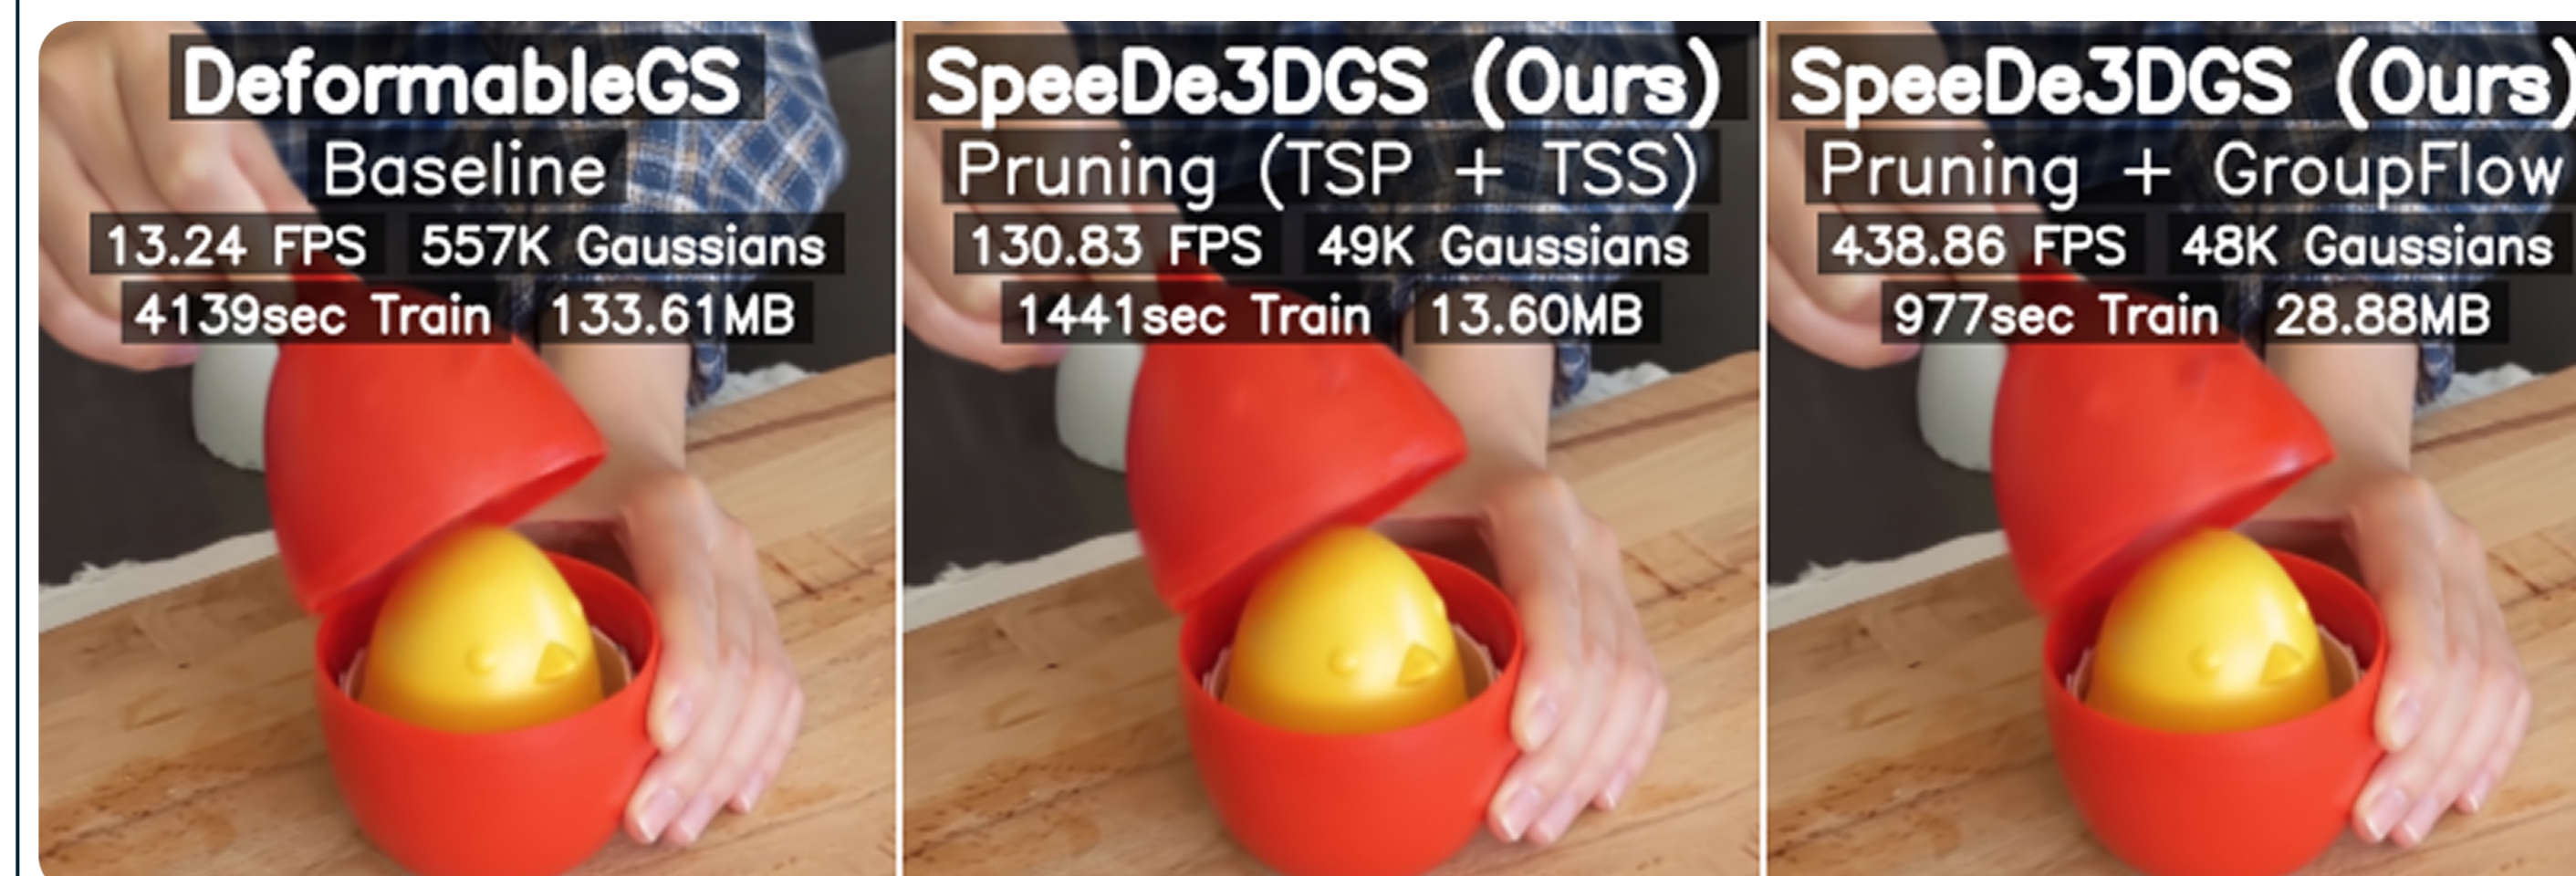
CVPR 2025 · speedysplat.github.io



Accelerate 3D Gaussian Splatting rendering speed by 2× for free by accurately localizing primitives during rasterization and over 6× in total by pruning the scene by over 90% during training, producing higher speedups than previous technique while maintaining competitive image quality.

Speede3DGS: Speedy Deformable 3D Gaussian Splatting with Temporal Pruning and Motion Grouping

CVPR 2026 · speede3dgs.github.io

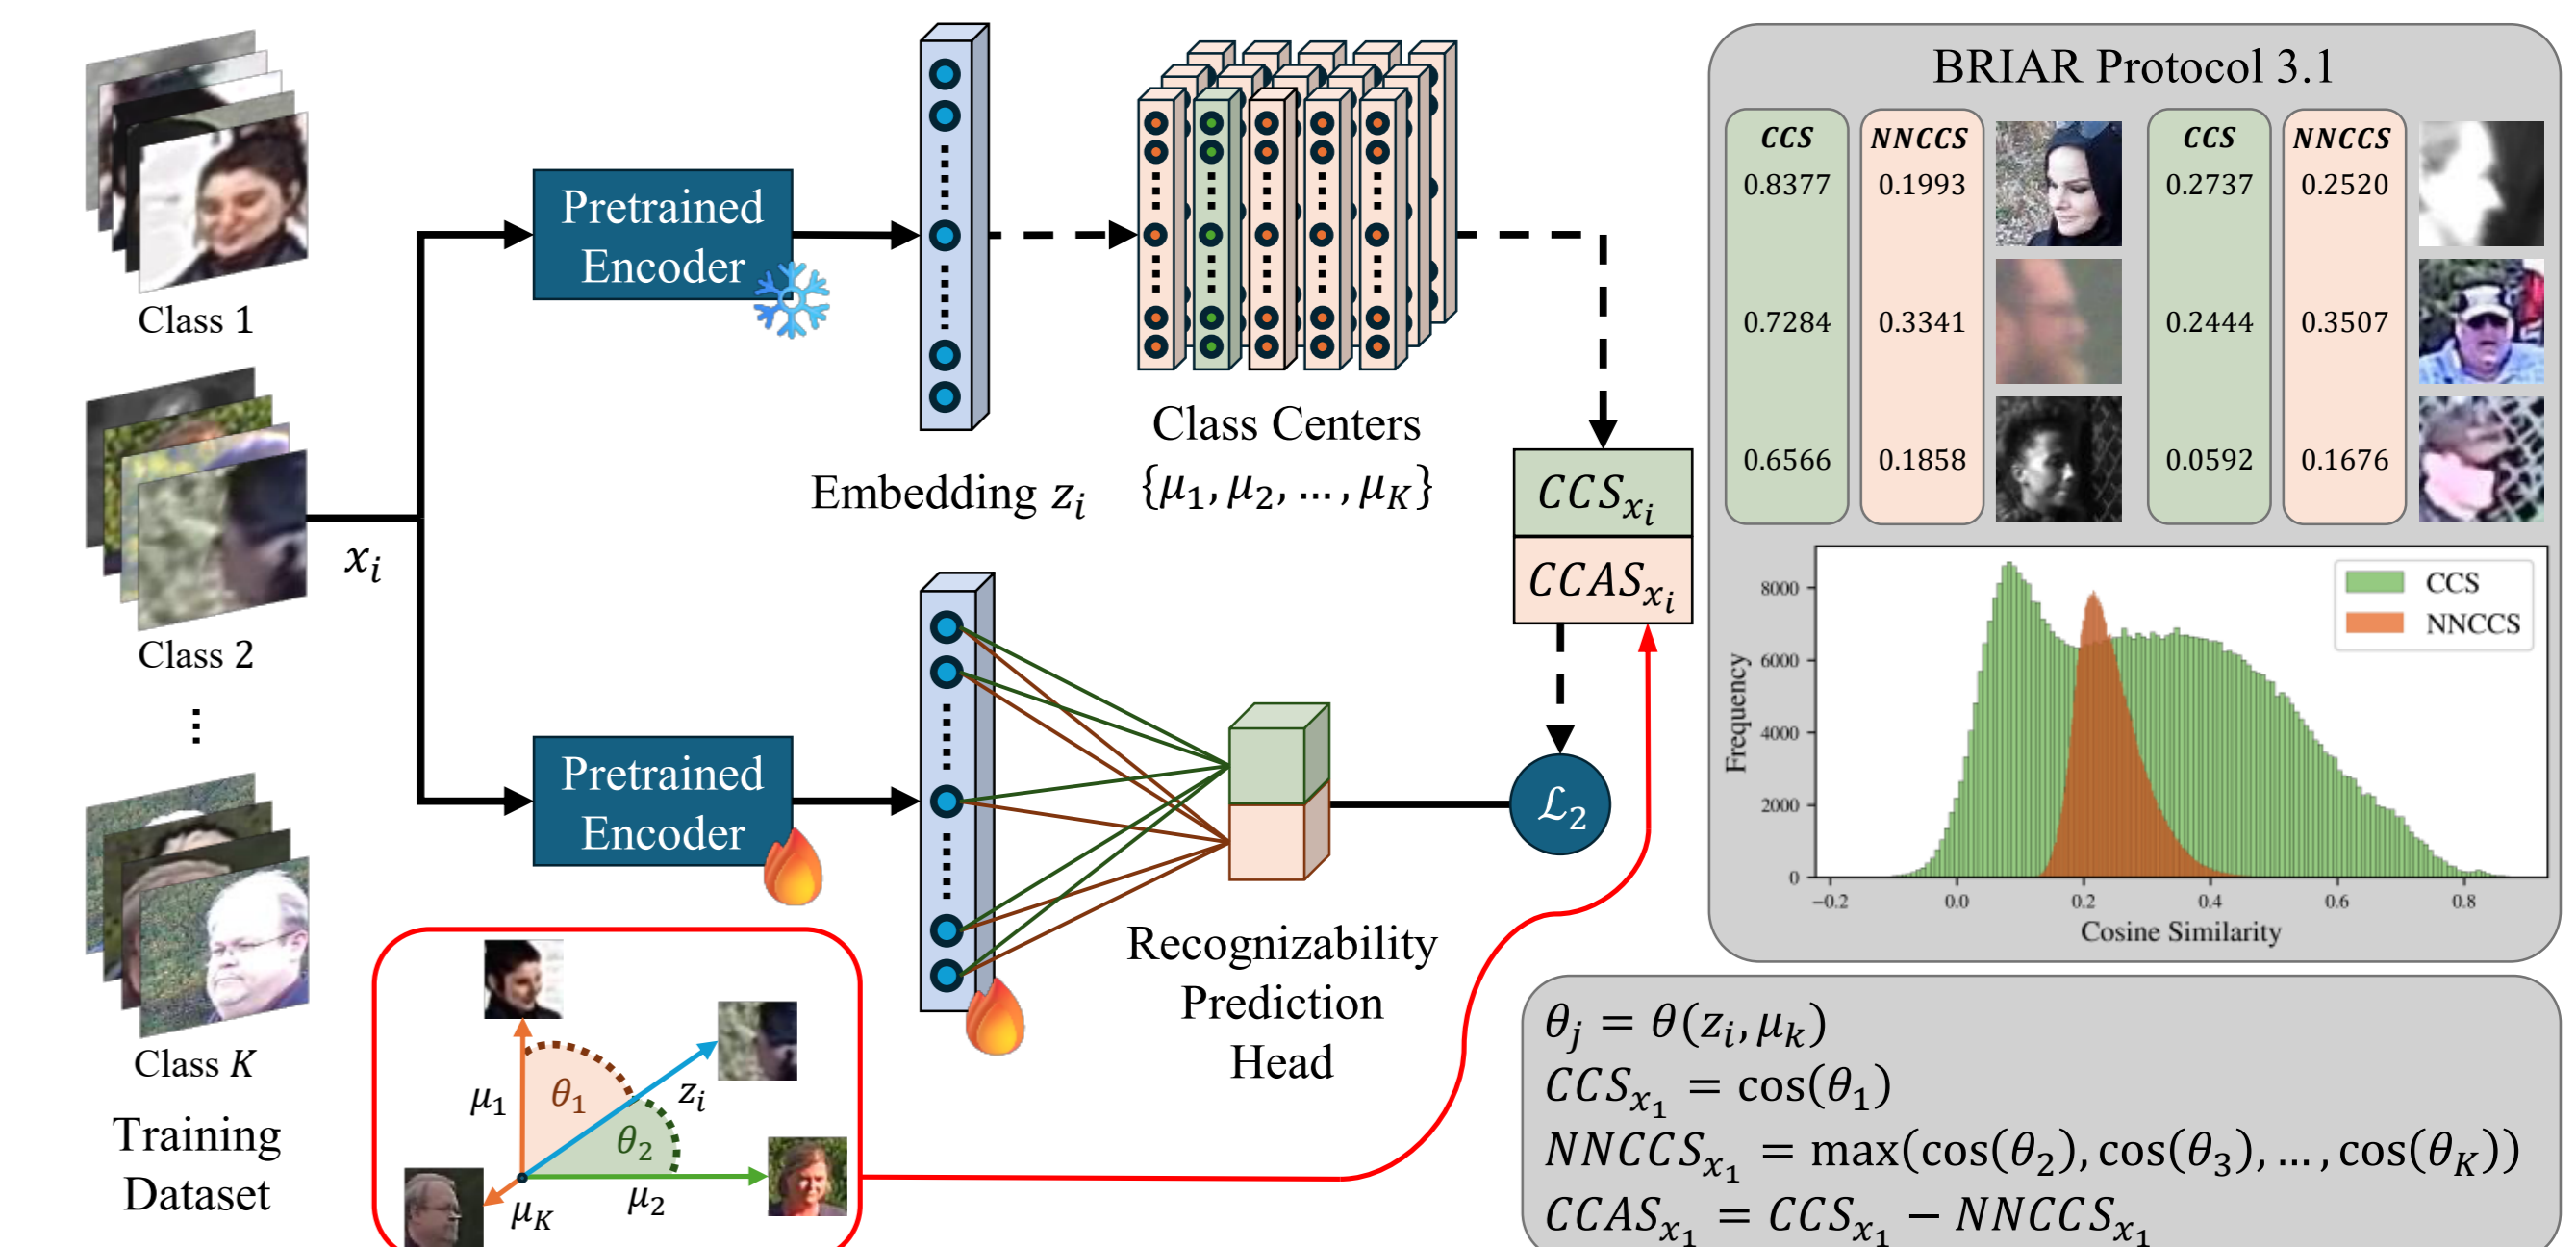


Boost DeformableGS rendering speed from 20 to 276 FPS using temporal sensitivity pruning and groupwise SE(3) motion distillation, all while preserving the superior image quality of per-Gaussian neural motion.

2. What to Trust

TransFIRA: Transfer Learning for Face Image Recognizability Assessment

FG 2026 (Oral) · transfira.github.io



Learn explainable, encoder-specific recognizability assessment from embedding geometry, enabling principled filtering, recognizability-aware template aggregation, state-of-the-art face image quality assessment performance, and generalization to body recognizability assessment.

3. Where to Generate

SplatSuRe: Selective Super-Resolution for Multi-view Consistent 3D Gaussian Splatting

CVPR 2026 · splatsure.github.io



Enhance low-resolution 3D Gaussian Splatting by injecting super-resolution only in regions lacking high-frequency supervision from other views, yielding sharper, more faithful, multi-view consistent reconstructions.

UniBRep: Learning Unified Geometry and Topology for Image-conditioned B-Rep Generation

