

From Far and Near: Perceptual Evaluation of Crowd Representations Across Levels of Detail

Xiaohan Sun
Trinity College Dublin
Ireland
sunx4@tcd.ie

Carol O’Sullivan
Trinity College Dublin
Ireland
Carol.OSullivan@tcd.ie

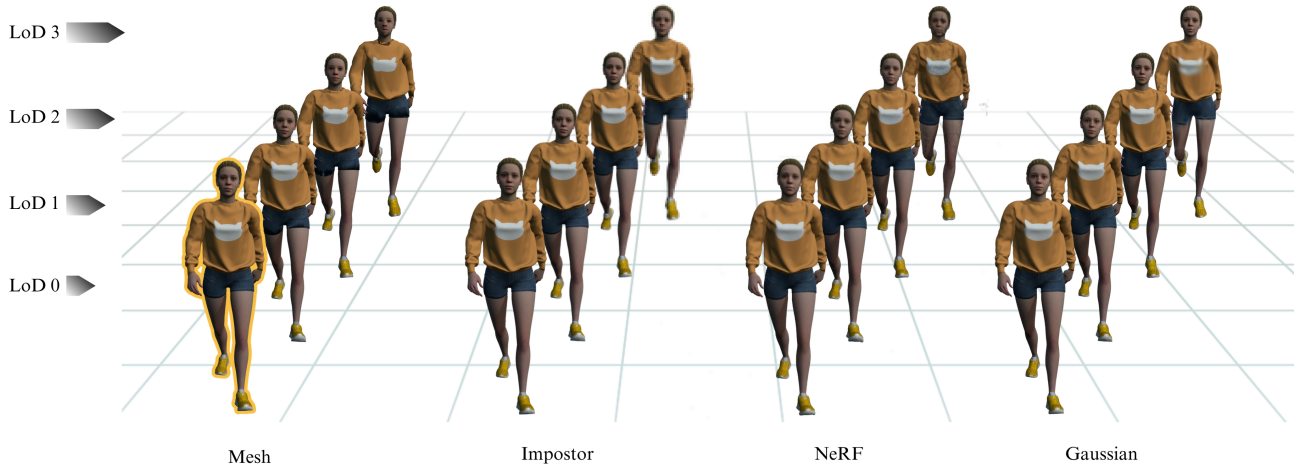


Figure 1. Comparison of four crowd rendering representations (geometric meshes, image-based impostors, Neural Radiance Fields, and 3D Gaussians) across four levels of detail (LoD 0-3). LoD 0 corresponds to the highest resolution (100 %), with detail halved at each subsequent level (50 %, 25 %, 12.5 %). The examples illustrate how fidelity, structural clarity, and smoothness degrade under simplification, revealing characteristic differences among representation types.

Abstract

In this paper, we investigate how users perceive the visual quality of crowd character representations at different levels of detail (LoD) and viewing distances. Each representation, including geometric meshes, image-based impostors, Neural Radiance Fields (NeRFs), and 3D Gaussians, exhibits distinct trade-offs between visual fidelity and computational performance. Our qualitative and quantitative results provide insights to guide the design of perceptually optimized LoD strategies for crowd rendering.

Keywords: Crowd rendering, Levels of Detail (LoD), Neural Rendering, Impostor, Perception

1. Introduction

Level of detail (LoD) techniques are widely used in crowd rendering to balance visual realism with performance

efficiency. Simplified geometric meshes and image-based impostors have previously been employed to render large, animated crowds in real time (e.g., [10]). With the recent emergence of neural rendering methods, such as Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS), a new question is raised: how do traditional LoD approaches compare with modern neural representations; and under what conditions might a neural representation be perceptually equivalent to a high quality mesh representation?

To answer this question, we compare the perceptual impact of LoD rendering for both traditional and neural crowd character representations. Through a controlled user study, we examine which LoD representations are perceptually indistinguishable from a high-resolution, ground-truth mesh at different viewing distances and levels of detail. Our contributions include:

1. **Perceptual thresholds for LoD placement:** We analyze the proportion of times each of the four representations—*Mesh*, *Impostor*, *NeRF*, and *3D Gaussian*—is

judged as most similar to a high-quality mesh. Each representation is evaluated at multiple LoDs and viewing distances, both with and without motion. The resulting perceptual data can inform threshold selection for LoD placement and switching.

2. **Perception-driven LoD pipeline and toolkit.** We introduce a representation-aware LoD rendering pipeline that integrates the four crowd character representations. For reproducibility, we also provide practical guidelines and tools for generating customized LoD assets.

2. Related Work

2.1. Data Representations for Crowd Rendering

Real-time crowd systems have traditionally relied on rigged, skinned *meshes*. In production, assets often share a common rig and topology with per-agent variation in materials and animation clips; engines pack textures into atlases, reuse animation via clip sampling and blending, and issue large instanced draw calls with palette skinning to keep submission overhead low at scale [3, 16].

Image-based approaches replace geometry with 2D view samples, thereby decoupling per-agent cost from polygon count. Early city-scale systems involved the precomputation of atlases over azimuth and animation phases [19]. Later, silhouette errors and texture distortions were reduced by using 2D polygonal proxies [7], or by animating limb imposters independently [2]. A recent neural rendering approach compresses appearance and motion with a CNN to enable constant-time, view-flexible character rendering with explicit material/lighting control [13].

Point-based rendering methods trade connected mesh polygons for point and particle primitives, and can deliver simple, prefilterable samples that are suitable for animation. Early works established the use of points as rendering primitives [9] and introduced surface splatting with anisotropic filters [29]. For animated content, Wand et al. [21] achieved real-time rendering of complex, moving geometry by building multi-resolution hierarchies of prefiltered point sets for keyframed scenes. Ellipsoidal splat formulations extend this idea by using anisotropic kernels with visibility-aware rasterization [29]. This precursor to the ellipsoidal kernels used in modern Gaussian splatting [8] is conceptually akin to point-based rendering, but with learned or optimized appearance parameters.

Neural scene representations provide promising new LoD options beyond polygons and image-based imposters. For example, Neural Radiance Fields (NeRF) can model volumetric radiance via a per-ray multilayer perceptron (MLP) [11]. These are typically converted to explicit or factorized structures for speedy deployment [25, 12]. Gaussian splatting can also deliver real-time rendering using

anisotropic splats and a fast, visibility-aware renderer [8], while space-time extensions tackle dynamic scenes and temporal coherence [24].

2.2. Level of Detail (LoD) for Animated Crowds

LoD rendering for animated crowds may be framed as the process of choosing the least costly representation per character that is perceptually adequate for given viewing conditions. LoD strategies that consider geometry, motion, materials, and behavior have been proposed, incorporating hierarchical reductions of skeletons, meshes, and animation samples to deliver interactive rates at scale. These approaches exploit batching and instancing, GPU-centric skinning and shading, animation compression, and visibility/overdraw control [14, 20, 3, 16].

At larger viewing distances, image-based impostors can be displayed based on the camera viewpoint and each character’s animation pose, thus reducing the overall rendering cost [19]. The viable viewing distance can be further extended by preserving silhouettes using 2D polygonal proxies [7], while per-joint impostors provide more flexibility for varied animations [2]. Point/particle proxies (and displaced-subdivision variants) have also been proposed and compared with image-based methods [15]. Hybrid 2.5D systems, such as *Geopostors* [4], strive to balance memory, draw calls, and temporal plausibility. Viewing distance bands are defined, beyond which geometry is replaced by impostors, and restored when the camera approaches. More recently, neural rendering methods for LoD control have been explored, e.g., by converting NeRF to explicit or factorized forms for faster evaluation, or via native, adaptive detail selection in Gaussian avatars and crowd pipelines [6, 17].

Comparative studies have been conducted to examine the relative advantages of different LoD crowd character representations. McDonnell et al. [10] systematically compared image-based impostors with low-resolution geometric meshes, with respect to appearance and motion fidelity. More recently, Sun et al. [18] examined the perceived quality of 3D Gaussian avatars based on motion, level of detail, and distance. In this paper, we extend this body of work by comparing the perceived quality of *Mesh*, *Impostor*, *NeRF*, and *3D Gaussian* representations of animated characters across different LoDs and viewing distances.

3. Implementation Methods

3.1. Datasets

We used a Mixamo [1] female character walking for 60 frames and captured each frame from 60 virtual cameras on a hemispherical rig (3,600 images total; see Figure 2). For all captures, we store standard pinhole camera intrinsics and rigid camera poses.



Figure 2. Examples of data.

To support different crowd representations, we construct representation-specific datasets from this shared capture. Mesh and impostor assets are derived directly from render-based captures, while learning-based methods additionally rely on calibrated multi-view inputs. Table 1 summarizes the preprocessing steps applied for each representation. We next describe how this shared dataset is instantiated for each crowd representation.

Table 1. Data preprocessing steps by representation used in our implementation pipeline.

Representation	Data preprocessing
Mesh	60 cameras \times 60 frames render capture; export LoDs via Blender <i>Decimate</i> .
Impostor	Union-alpha bounding box stabilization; pack 6×10 sprite sheets per LoD.
NeRF (Instant-NGP [12])	Camera intrinsics & poses from COLMAP; graphics convention export (x right, y up, $-z$).
3D Gaussian (3DGS [8])	Camera intrinsics & poses from COLMAP; sparse colored SfM to initialize splats.

3.2. Mesh

The mesh representation serves as our high-fidelity baseline, in which the captured character is rendered directly using standard skinned mesh animation. Our four Mesh LoDs are generated using Blender’s *Decimate* modifier. The base model (L0) and three simplified versions are created by applying decimation ratios of 1.0, 0.5, 0.25, and 0.125, yielding face counts of 27,048; 18,436; 10,811; and 5,864 respectively. Reducing the polygon count of animated characters to generate geometric mesh LoDs is a common practice to improve rendering performance in crowds [3, 5]. Edge-collapse or “collapse” type decimation (as used in Blender) provides control over face count reduction while preserving shape and animation fidelity as shown in Figure 3.




Figure 3. Examples of mesh LoD.

3.3. Impostor

For our second representation, the animated character is rendered as a precomputed impostor sequence packed into a 6×10 sprite sheet per LoD. At the highest level (L0), each tile consists of 1080×1080 pixels, which matches the resolution used to render the reference mesh. Each lower level halves the resolution to 540×540 , 270×270 , and 135×135 px (L1–L3). To avoid the per-frame “breathing” and root-motion loss typical of naive per-frame crops, we first compute a *union* alpha bounding box over all frames of the run cycle, crop every frame to that fixed window, and apply a single global scale per LoD before centering on a square canvas. The stabilized frames are then packed column-major into a texture atlas that Blender reads as a single image. At render time we use an unlit material and animate UVs over a sprite-sheet atlas: the UVs are cropped to a single tile and offset at each frame so that the impostor displays one precomputed pose at a time. This workflow follows standard impostor practice, replacing distant geometry by a textured quad while indexing an atlas of animation frames. It is widely used in crowd rendering for performance, with quality governed by atlas resolution and sampling order [19]. To match the baked appearance, we fix the impostor’s pose and world position to the capture setup and render from the same camera transform and FOV used during capture.

3.4. NeRF

We implement NeRF as our third representation, using *Instant-NGP*, which replaces the original fully connected NeRF with a multiresolution hash-grid encoding feeding a compact MLP, substantially reducing training and inference cost while retaining high visual quality [11, 12]. LoD is controlled via the hash table capacity and network width: we fix the number of hash levels L and features per level F , and from L0 to L3 successively halve the hash capacity (decreasing $\log_2 T$ by 1 per step) while modestly narrowing the MLP. Exact presets appear in Table 2. For animation, we reconstruct independent NeRFs for each of 60 frames and render them from a shared camera path, yielding a frame-wise dynamic sequence. To validate that L0 serves as a suitable upper bound, we evaluated it on 60 held-out views, obtaining PSNR = 36.18 and SSIM = 0.988, indi-

cating high-fidelity reconstruction quality overall.

3.5. Gaussian Splatting

3D Gaussian Splatting is our fourth representation, with a visibility-aware rasterizer and anisotropic splats [8]. Fidelity is controlled by *capping the Gaussian count N*. For large reductions we apply a light opacity prune ($\alpha < 0.01$). Our four presets are $N = \{120k, 30k, 7.5k, 1.9k\}$ for L0 to L3, resp. All LoDs use the same optimization recipe: an $L_1 + \text{SSIM}$ loss, spherical-harmonic color basis of degree 2, and Adam optimizer. Each continuation runs for 20k steps. LoD is applied with a simple count budget and a light opacity prune. This practical setup is compatible with recent approaches that address aliasing via scale-aware filtering and reduce model size through learned pruning [26, 28]. As with NeRF, we optimize a separate 3DGS per frame for 60 frames and render from a common camera, producing the dynamic sequence. L0 achieves PSNR = 36.69 and SSIM = 0.991 on 60 held-out views, confirming that our highest-detail preset provides a strong reference for comparisons.

4. Experimental Design

4.1. User Study

Stimuli. Four representations were displayed: Image-based Impostors (**I**), Meshes (**M**), NeRFs (**N**), and 3D Gaussians (**G**). Four levels of detail were created: (L0=100%, L1=50%, L2=25%, L3=12.5%) and five viewing distances (D0=100% pixels, D1=80%, D2=60%, D3=40%, D4=20%). In each trial, the four representations were displayed side by side in counterbalanced order. Participants completed two blocks: a) *Video* (with motion) and *Image* (with no motion).

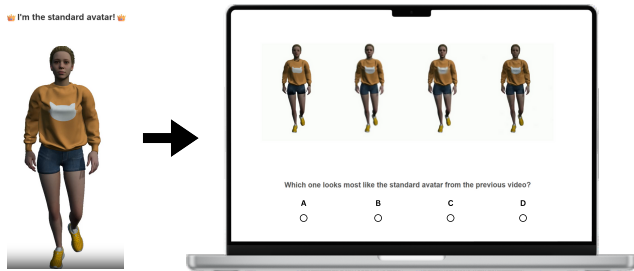


Figure 4. Example of video stimulus.

Participants and procedure. Twenty-four volunteers (17 male, 7 female; age 18–60+) completed a brief training session followed by two trial blocks (Video and Image). The experiment was conducted online via Qualtrics and was fully anonymous. Participants completed the study on their own devices (laptops, desktops, or tablets), with reported screen sizes ranging from under 15 inches to over 21 inches. Data from one participant using a mobile device were excluded to ensure consistent visual presentation.

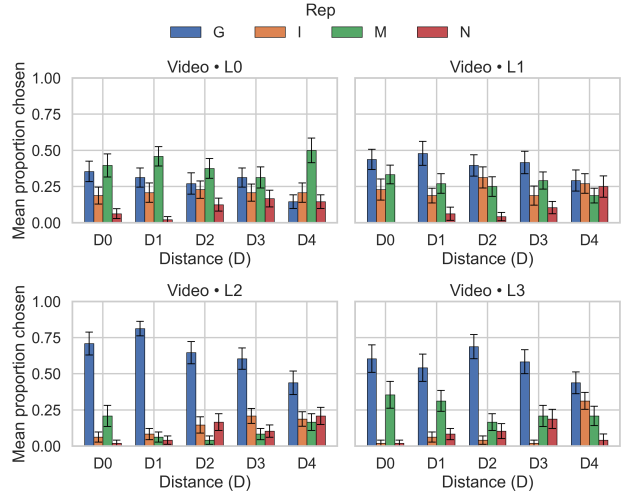


Figure 5. Video block: mean proportion chosen for each representation across LoD and Distance (D0–D4). Error bars show ± 1 SE across participants.

Screen size was recorded at the start of the experiment and showed no significant effect on perceptual judgments. For each trial, participants viewed a high-resolution mesh reference and then chose with one of the four displayed stimuli most closely matched the high-quality mesh (see Fig. 4). Trials were randomized, and the positions of the four representations were counterbalanced across participants.

Statistical analysis. Choices were analyzed as fully within-subjects repeated measures with factors *Representation* (G, I, M, N), *Distance* (D0–D4), and *LoD* (L0–L3), in both *Modes* (Image vs. Video). We averaged the repetitions per participant and condition to obtain selection proportions. We fitted an OLS with subject fixed effects and the full $\text{Representation} \times \text{Distance} \times \text{LoD} \times \text{Mode}$ factorial, and report a Type II ANOVA (Table 3). Because responses are binary at the trial level, we additionally fitted a trial-level binomial GLM (logit) with subject fixed effects, and obtained likelihood-ratio omnibus tests by comparing the full model to hierarchically reduced models (Table 4). Descriptive means (\pm SE) are shown in Figs. 5.

Results. The ANOVA revealed a strong main effect of *Representation*, as well as robust $\text{Representation} \times \text{Distance}$ and $\text{Representation} \times \text{LoD}$ interactions, The three-way interaction $\text{Representation} \times \text{Distance} \times \text{LoD}$ was also significant. No main effects of *Distance*, *LoD* or *Mode* (Image vs. Video) were observed and no interactions involving *Mode* reached significance (min $p = 0.257$). The confirmatory trial-level GLM echoed these patterns with significant omnibus LR tests for *Representation*, *Distance*, *LoD*, $\text{Representation} \times \text{LoD}$, $\text{LoD} \times \text{Distance}$, and the three-way interaction (Table 4).

Table 2. Instant-NGP LoD presets. L : hash levels; F : features/level; $T = 2^{\log_2 \text{hashmap_size}}$.

LoD	L	F	$\log_2 T$	Base res.	Density MLP (neurons, layers)	Dir. enc.	RGB net (neurons, layers)
LoD-0	12	2	18	16	128, 1×	SH deg. 4 + Identity	64, 2×
LoD-1	12	2	17	16	64, 1×	SH deg. 3 + Identity	32, 2×
LoD-2	12	2	16	16	32, 1×	SH deg. 2 + Identity	16, 2×
LoD-3	12	2	15	16	16, 1×	SH deg. 1 + Identity	16, 1×

Table 3. ANOVA showing the results of main and interaction effect tests with degrees of freedom, effect sizes (η^2) and (p) values. (Common denominator df for the OLS Type II ANOVA is 3657.)

Effect Tested	dof	F-Test	η^2	p
Representation	3	177.52	0.127	< 0.001
Representation × Distance	12	8.53	0.027	< 0.001
Representation × LoD	9	33.16	0.075	< 0.001
Representation × Mode	3	1.08	0.001	0.356
Representation × Distance × LoD	36	2.26	0.022	< 0.001
Representation × Distance × Mode	12	0.36	0.001	0.976
Representation × LoD × Mode	9	0.90	0.002	0.527
Representation × Distance × LoD × Mode	36	1.14	0.011	0.257

Table 4. Omnibus likelihood-ratio (LR) tests from the trial-level binomial GLM with subject fixed effects.

Effect	LR	df	p
Representation	1089.30	139	< 0.001
Distance	275.81	140	< 0.001
LoD	418.92	136	< 0.001
Mode	66.02	80	0.869
Representation × Distance	61.83	76	0.880
Representation × LoD	418.92	121	< 0.001
Representation × Mode	66.02	79	0.851
LoD × Distance	152.80	124	0.040
Representation × LoD × Distance	150.97	112	0.008

Discussion. Perceived fidelity is seen to depend on a combination of *Representation*, viewing *Distance*, and *LoD*. As expected, **Mesh** is most preferred at high detail and near distances, whereas **Gaussians** become increasingly indistinguishable from the Mesh as detail decreases and/or viewing distance increases. Both the omnibus ANOVA and the confirmatory GLM reveal strong effects of *Representation*, clear interactions with *LoD*, and a significant three-way interaction, whereas no significant main or interaction effect was found for *Mode* (Image vs. Video).

4.2. Quantitative Experiments

4.2.1 Visual Comparison

Image fidelity is calculated using an evaluation view camera that combines poses matching the user-study setup but is excluded from NeRF/3DGS training. For each LoD and representation, we compare with the Mesh reference in the same pose using PSNR, SSIM and LPIPS [23, 22, 27]. Table 5 reports metrics across LoDs; higher PSNR/SSIM and

lower LPIPS indicate closer agreement with the Mesh baseline.

Summary. At L0, both **3D Gaussians** and **NeRF** closely match the Mesh (Gaussians: 36.69 dB / 0.991 / 0.013; NeRF: 36.18 dB / 0.988 / 0.019 for PSNR / SSIM / LPIPS). With decreasing LoD, **NeRF** degrades smoothly in PSNR while maintaining high SSIM and low LPIPS; **Gaussians** show a larger PSNR drop but remain competitive on SSIM/LPIPS through mid LoDs. **Impostors** exhibit high SSIM and very low LPIPS at the matched view, despite lower PSNR, consistent with view-aligned textures that preserve structure but lack 3D parallax. Overall, frame-based metrics indicate that neural fields (NeRF/3DGS) are closest to Mesh at high detail, while impostors can appear perceptually similar at single views even with lower pixel-wise fidelity.

Table 5. Image Similarity at Evaluation View.

Rep.	LoD	PSNR ↑	SSIM ↑	LPIPS ↓
NeRF	L0	36.18	0.988	0.019
	L1	35.08	0.984	0.022
	L2	33.25	0.977	0.030
	L3	28.09	0.954	0.052
3DGS	L0	36.69	0.991	0.013
	L1	24.76	0.961	0.029
	L2	24.52	0.954	0.045
	L3	24.01	0.938	0.090
Impostor	L0	23.95	0.965	0.006
	L1	23.40	0.970	0.010
	L2	21.35	0.970	0.017
	L3	19.40	0.971	0.030

Table 6. On-disk asset size and offline per-frame training time by representation and LoD. Training time is reported for learning-based methods only (20k steps on RTX 4090).

Representation	LoD0		LoD1		LoD2		LoD3	
	Size	Train	Size	Train	Size	Train	Size	Train
NeRF (Instant-NGP)	21.2 MB	~240 s	17.0 MB	~207 s	15.5 MB	~214 s	14.0 MB	~205 s
3D Gaussian (3DGS)	19.0 MB	~320 s	8.00 MB	~122 s	2.00 MB	~78 s	0.768 MB	~58 s
Impostor	328 KB	—	100 KB	—	32.0 KB	—	12.0 KB	—
Mesh	5.19 MB	—	4.37 MB	—	3.86 MB	—	3.50 MB	—

4.2.2 Memory Usage and Training Time

We quantify deployment memory as the on-disk size of the exact LoD assets used in the study. This metric is engine-agnostic and comparable across heterogeneous toolchains.

Table 6 reports the footprint per representation and LoD. The entries correspond to (i) **Mesh**: the rigged mesh package (geometry, armature, and textures), (ii) **Impostor**: billboard atlas images, (iii) **3D Gaussian**: optimized Gaussian caches, and (iv) **NeRF**: trained model snapshots.

For learning-based representations, we additionally report the offline training time required to generate each LoD asset. This cost reflects authoring and preprocessing overhead and is incurred prior to deployment, with no impact on runtime performance. Training time is reported per frame. Traditional representations (meshes and impostors) do not require learning-based optimization and therefore have no associated training cost.

Summary. Deployment footprint varies significantly by representation. **Impostors** are the most memory-efficient by a wide margin and scale down aggressively with decreasing LoD, while also incurring no learning-based training cost. **3D Gaussian** representations exhibit strong compressibility across LoDs, offering a tunable volumetric alternative that requires substantially less memory than meshes at lower LoDs, at the expense of offline training time that decreases with representation complexity. In contrast, **Mesh** assets yield only modest memory savings under simplification, reflecting constraints imposed by geometry and rig data, but benefit from the absence of training overhead. **NeRFs** require the largest deployment footprint across all LoDs and achieve comparatively smaller proportional memory reductions, while also incurring consistent offline training costs per LoD. Overall, these results highlight distinct trade-offs between memory efficiency and authoring cost across representations, which are critical to consider when designing perceptually optimized LoD strategies for animated crowds.

5. Conclusions and Future Work

Our results show that perceived similarity to a mesh ground truth is driven primarily by a character’s *Representation* and its interactions with *Distance* and *LoD*, and does not appear to be affected by the presence or absence of motion.

Based on these results, we propose the following practical guidelines: **Impostors** are the cheapest and most scalable choice for far/low-detail rendering (i.e., minimal geometry cost and excellent batching at distance), although flexibility is sacrificed for large precomputed atlases and limited animation blending. **3D Gaussians** become increasingly indistinguishable from meshes as pixel density drops or detail is reduced, thus offering a strong mid-to-high LoD alternative with real-time rendering. Geometric **Meshes** retain a clear advantage at nearer views and higher detail, but for further distances, more efficient representations are highly competitive.

Our findings have direct implications for immersive and virtual reality applications, where perceptually informed LoD selection can help balance visual realism and performance in crowd-populated environments, contributing to stable frame rates and user comfort. We conclude by outlining directions for future work.

Beyond appearance: motion, dynamics, and generalization. To enable a controlled comparison, our study focuses on visual appearance using a single character, fixed lighting, and a simple animation. Within this setting, motion does not significantly affect perceived similarity across representations. An interesting direction for future work is to extend this evaluation to richer motion patterns, diverse character appearances and lighting conditions, and physical interactions such as clothing dynamics and inter-agent contact. Longer sequences and more dynamic scenarios may further reveal how perceptual trends observed here carry over to increasingly complex crowd settings, and motivate the exploration of temporally consistent and dynamic neural rendering approaches.

Runtime performance and system-level evaluation.

While this work emphasizes perceptual quality and asset-level efficiency, end-to-end runtime performance remains a critical factor for real-time deployment and depends on hardware, renderer implementation, and system-level optimizations. Future work should incorporate standardized performance benchmarks to more explicitly quantify the trade-offs between perceptual fidelity and rendering throughput in practical deployment scenarios.

In summary, we present a unified perceptual evaluation of Mesh, Impostor, NeRF, and 3D Gaussian representations and provide practical guidance for managing LoD in crowd rendering systems. Our results and accompanying pipeline offer a foundation for designing perceptually driven crowd representations that scale effectively across viewing conditions.

Acknowledgement

This work was conducted with the financial support of the Research Ireland Centre for Research Training in Digitally-Enhanced Reality (d-real) under Grant No. 18/CRT/6224. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

- [1] Adobe. Mixamo. <https://www.mixamo.com/>. Accessed 2025-10-13. 2
- [2] Á. Beacco, N. Pelechano, and C. Andújar. Efficient rendering of animated characters through optimized per-joint impostors. *Computer Animation and Virtual Worlds*, 23(3–4):331–339, 2012. 2
- [3] A. Beacco, N. Pelechano, and C. Andújar. A survey of real-time crowd rendering. *Computer Graphics Forum*, 35(8):32–50, 2016. 2, 3
- [4] S. Dobbyn, J. Hamill, K. O’Conor, and C. O’Sullivan. Geopostors: A real-time geometry/impostor crowd rendering system. In *Proceedings of the Symposium on Interactive 3D Graphics and Games (I3D)*, pages 95–102, 2005. 2
- [5] Y. Dong and C. Peng. Real-time large crowd rendering with efficient character and instance management on gpu. *International Journal of Computer Games Technology*, 2019(1):1792304, 2019. 3
- [6] X. Dongye, H. Guo, L. Luo, H. Jiang, Y. Bao, Z. Tian, and D. Weng. Lodavatar: Hierarchical embedding and adaptive levels of detail with gaussian splatting for enhanced human avatars. arXiv:2410.20789, 2024. 2
- [7] L. Kavan, S. Dobbyn, S. Collins, J. Žára, and C. O’Sullivan. Polypostors: 2d polygonal impostors for 3d crowds. In *Proceedings of the Symposium on Interactive 3D Graphics and Games (I3D)*, pages 149–155, 2008. 2
- [8] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023. 2, 3, 4
- [9] M. Levoy and T. Whitted. The use of points as a display primitive. Technical Report TR 85-022, Department of Computer Science, University of North Carolina at Chapel Hill, 1985. 2
- [10] R. McDonnell, S. Dobbyn, and C. O’Sullivan. Lod human representations: A comparative study. In *V-CROWDS: Intl. Workshop on Crowd Simulation*, pages 101–115, 2005. 1, 2
- [11] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2, 3
- [12] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (SIGGRAPH)*, 41(4):1–15, 2022. 2, 3
- [13] M. Ostrek, N. J. Mitra, and C. O’Sullivan. Neuropostors: Neural geometry-aware 3d crowd character impostors. In *Pattern Recognition: 27th International Conference, ICPR 2024, Kolkata, India, December 1–5, 2024, Proceedings, Part XXII*, volume 15322 of *Lecture Notes in Computer Science*, pages 432–448. Springer, 2024. 2
- [14] C. O’Sullivan, J. Cassell, J. Dingliana, S. Dobbyn, B. McNamee, C. Peters, and T. Giang. Levels of detail for crowds and groups. *Computer Graphics Forum*, 21(4):733–741, 2002. 2
- [15] I. Rudomín and E. Millán. Point-based rendering and displaced subdivision for interactive animation of crowds of clothed characters. In *VRIPHYS 2004: Virtual Reality Interaction and Physical Simulation Workshop*, pages 139–148, 2004. 2
- [16] G. Ryder, M. Chalmers, and R. McCreedy. A survey of real-time rendering techniques for crowds. In *Eurographics State of the Art Reports*. 2005. 2
- [17] X. Sun, Y. Xu, J. Dingliana, and C. O’Sullivan. Crowd-splat: Exploring gaussian splatting for crowd rendering. arXiv:2501.17792, 2025. Includes LoD integration for crowd-scale rendering. 2
- [18] X. Sun, Y. Xu, J. Dingliana, and C. O’Sullivan. Evaluating crowdsplat: Perceived level of detail for gaussian crowds. In *2025 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pages 720–724. IEEE, 2025. 2
- [19] F. Tecchia, C. Loscos, and Y. Chrysanthou. Image-based crowd rendering. *IEEE Computer Graphics and Applications*, 22(2):36–43, 2002. 2, 3
- [20] L. Toledo, O. D. Gyves, and V. Rudomín. Hierarchical level of detail for varied animated crowds. *The Visual Computer*, 30(6):949–961, 2014. 2
- [21] M. Wand and W. Straßer. Multi-resolution rendering of complex animated scenes. *Computer Graphics Forum*, 21(3):483–491, 2002. Proc. Eurographics 2002. 2
- [22] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 5
- [23] S. Wolf and M. Pinson. Reference algorithm for computing peak signal-to-noise ratio (psnr) of a video sequence with a

constant delay. Contribution com 9–c 6, ITU-T Study Group 9 (contribution from NTIA/ITS, USA), Geneva, Switzerland, Feb. 2009. Defines a standardized PSNR computation procedure for video. [5](#)

- [24] G. Wu, T. Yi, J. Fang, L. Xie, X. Zhang, W. Wei, W. Liu, Q. Tian, and X. Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *CVPR*, 2024. [2](#)
- [25] A. Yu, R. Li, M. Tancik, L. Hao, R. Ng, and A. Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. [2](#)
- [26] Z. Yu, B. Kerbl, C. Xie, N. Rüegg, S. I. Liu, and A. Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *CVPR*, 2024. [4](#)
- [27] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. [5](#)
- [28] Z. Zhang, T. Song, Y. Lee, L. Yang, C. Peng, R. Chellappa, and D. Fan. Lp-3dgs: Learning to prune 3d gaussian splatting. In *NeurIPS*, 2024. [4](#)
- [29] M. Zwicker, H. Pfister, J. van Baar, and M. Gross. Surface splatting. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*, pages 371–378. ACM, 2001. [2](#)