Citation Evidence Report

EB-1B Petition — Outstanding Professor or Researcher

8 CFR § 204.5(i)(3) · Authorship + Original Contributions

Isabella Liu

University of California, San Diego

Google Scholar profile

Generated 2026-05-21 by CiteMap. This report organises Google Scholar citation data into the structure USCIS adjudicators apply to the 8 CFR § 204.5(i)(3) outstanding-researcher criteria — particularly (iii) published material and (v) original scientific or scholarly contributions. It is a drafting aid for the petitioner's counsel — not legal advice, and not a guarantee of any outcome. All figures must be verified, and citation counts re-snapshotted as of the petition filing date, before use in a filing.

A. Overview & Filtering Statement

450	479	12	9
Citing papers mapped	Citation edges	Home papers mapped	h-index (GS)

Filtering statement — methodology & limits

Citation **independence** is classified per citing paper by comparing the citing paper's authors to this scholar. *Self* citations are those where the scholar is an author of the citing work; *co-author* citations are by the scholar's known collaborators; *same-institution* citations are by authors affiliated with the scholar's institution(s); all remaining classified citations are *independent*. Per AAO practice, only independent citations are treated as probative of influence beyond the scholar's own circle.

Known limitations — counsel must verify. (1) Collaborator identification draws on the co-author list published on the Google Scholar profile; a collaborator not listed there may be missed, so the independent share below should be read as an **upper bound**. (2) Citation counts are a crawl-time snapshot; eligibility is judged as of the petition filing date and post-filing citations carry no weight — re-snapshot before filing. (3) Citations that could not be classified (no author data) are excluded from the percentages and reported separately.

B. Citation Independence

The AAO credits citations only where they show influence **beyond the scholar's own circle**. Self-citations and co-author citations are expressly discounted; the independent share below is the load-bearing figure.

93.1% independent of 346 classified citing papers

Citation type	Count
Independent	322
Self-citation	4
Co-author	20
Same-institution	0

104 citing papers could not be classified (no author data) and are excluded from the percentages above.

C. Significant Contributions & Their Citation Evidence

Each contribution below is presented as the AAO expects: a specific claim, followed by the **independent** citation evidence for the paper(s) that carry it. Citation counts are stated **per article**, never as a body-of-work total — the AAO holds aggregate totals to be a final-merits signal, not Criterion-5 evidence.

Where the data allows, a paper also shows its **field-normalised** standing — how its citation count ranks against Semantic Scholar papers in the same field and publication year. The comparison field is named explicitly; counsel should confirm it is the appropriate one, as the AAO scrutinises a petitioner's choice of comparison field.

Contribution 1

Claim — Contribution 1

The researcher pioneered tensorial inverse rendering methods, establishing a foundational framework for decomposing scene properties that has significantly influenced subsequent research in 3D reconstruction and illumination analysis.

The researcher's contribution centers on the development of tensorial inverse rendering, as introduced in the 2023 paper 'Tensoir: Tensorial inverse rendering.' This work serves as the cornerstone of a research line that addresses the complex challenge of disentangling material, geometry, and lighting from visual data. By proposing a tensorial approach, the researcher appears to have offered a novel mathematical framework for inverse rendering problems that were previously difficult to solve with high fidelity.

The originality of this line of work is further evidenced by subsequent publications that build upon these foundational concepts. The 2024 release of 'Openillumination,' a multi-illumination dataset for evaluating inverse rendering on real objects, suggests an effort to ground theoretical advances in empirical validation. Additionally, the 2025 paper 'Imls-splatting' indicates an extension of these principles toward efficient mesh reconstruction via point representation, demonstrating the versatility and applicability of the initial tensorial framework to broader 3D vision tasks.

The significance of this contribution is underscored by its substantial uptake in the academic community. With 233 citations for the core paper and 40 for the associated dataset, the work has clearly resonated with peers. Notably, 97.1% of the citing papers originate from independent researchers, indicating that the methodology has been widely adopted and validated by the broader scientific community rather than merely within the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 197 · **26 flagged influential by Semantic Scholar**

CORE PAPER

Tensoir: Tensorial inverse rendering

2023 · Computer Vision and Pattern Recognition (CVPR), 2023 · 233 citations (GS)

Field-normalised: 183 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2023 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Relightable 3d gaussians: Realistic point cloud relighting with brdf decomposition and ray tracing	Fudan University, Nanjing University	China	Methodology
2	Gs-ir: 3d gaussian splatting for inverse rendering	South China University of Technology, Tencent, The Chinese University of Hong Kong, Shenzhen	China	Methodology
3	Benchmarking neural radiance fields for autonomous robots: An overview	China Aerodynamics Research and Development Center, City College of New York, Hangzhou Dianzi University	China, United Kingdom, United States	Background
4	Digital twin catalog: A large-scale photorealistic 3d object digital twin dataset	Meta, Stanford University	United States	—
5	Gir: 3d gaussian inverse rendering for relightable scene factorization	Baidu Inc., Beihang University, Peking University	China	Methodology
6	Dreammat: High-quality pbr material generation with geometry-and light-aware diffusion models	Tencent Games, Texas A&M University, The University of Hong Kong	China, Hong Kong, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
7	Hravatar: High-quality and relightable gaussian head avatar	International Digital Economy Academy, Tsinghua Shenzhen International Graduate School, Tsinghua University	China	—
8	Irgs: Inter-reflective gaussian splatting with 2d gaussian ray tracing	Fudan University	China	—
9	Bilateral guided radiance field processing	The Chinese University of Hong Kong	China	—
10	MaterialRefGS: Reflective gaussian splatting with multi-view consistent material inference	NYU Abu Dhabi, Tsinghua University, Wayne State University	China, United Arab Emirates, United States	—
11	Animatable and relightable gaussians for high-fidelity human avatar modeling	NNKosmos Technology, Tsinghua University	China	Methodology
12	Rogr: Relightable 3d objects using generative relighting	Google DeepMind, Google Research, Technical University of Munich	Germany, United Kingdom, United States	—
13	SAGD: Boundary-enhanced segment anything in 3D Gaussian via Gaussian decomposition	Chinese Academy of Sciences, Shandong University, The Hong Kong Polytechnic University	China, Hong Kong	Background
14	Nerf-casting: Improved view-dependent appearance with consistent reflections	Carnegie Mellon University, Google, Google DeepMind	United Kingdom, United States	—
15	A Diffusion Approach to Radiance Field Relighting using Multi-Illumination Synthesis	Adobe Research, Inria, Inria, Université Côte d'Azur, Université Laval	Canada, France, France, Canada	—
16	Mirror-nerf: Learning neural radiance fields for mirrors with whitted-style ray tracing	Alibaba Group, Zhejiang University	China	—
17	Transparentgs: Fast inverse rendering of transparent objects with gaussians	Nanjing University, Nankai University, Zhejiang University	China	—
18	Ref-gs: Directional factorization for 2d gaussian splatting	Huazhong University of Science and Technology, Westlake University, Zhejiang University	China	—
19	Tensosdf: Roughness-aware tensorial representation for robust geometry and material reconstruction	Nanjing University, Shandong University, The Hong Kong Polytechnic University	China	—
20	Intrinsicavatar: Physically based inverse rendering of dynamic humans from monocular videos via explicit ray tracing	ETH Zürich, ETH Zürich, University of Tübingen, Tübingen AI Center, University of Tübingen	Germany, Switzerland, Switzerland, Germany	Methodology
21	Lirm: Large inverse rendering model for progressive reconstruction of shape, materials and view-dependent radiance fields	Meta Reality Labs, University of Maryland, College Park	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
22	Recent Trends in 3D Reconstruction of General Non-Rigid Scenes	Google, Max Planck Institute for Informatics, Saarland University	China, Germany, United Kingdom	—
23	Mv-colight: Efficient object compositing with consistent lighting and shadow generation	Nanjing University, Shanghai Artificial Intelligence Laboratory, Shanghai Jiao Tong University	China, Hong Kong	—
24	Reflective gaussian splatting	Fudan University, University of Surrey	China, United Kingdom	—
25	SpecNeRF: Gaussian directional encoding for specular reflections	Meta Reality Labs, The Hong Kong University of Science and Technology	China	Background
26	Illuminerf: 3d relighting without inverse rendering	Google, Google DeepMind, Google Research	United Kingdom, United States	Methodology
27	iVR-GS: Inverse volume rendering for explorable visualization via editable 3D Gaussian splatting	University of Notre Dame	United States	—
28	Diffusion posterior illumination for ambiguity-aware inverse rendering	Massachusetts Institute of Technology, Max-Planck-Institut, Reality Labs	Germany, United States	Methodology
29	Gaussian splatting with discretized sdf for relightable assets	Nanjing University, Nankai University	China	—
30	Deferredgs: Decoupled and editable gaussian splatting with deferred shading	Chinese Academy of Sciences, RWTH Aachen University, University of Surrey	China, Germany, United Kingdom	Methodology

Showing the 30 most-cited of 170 independent citing papers.

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology* / *Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

Citing-text excerpts — how the field used this work

METHODOLOGY Relightable 3d gaussians: Realistic point cloud relighting with brdf decomposition and ray tracing

"TensoIR [23] performs inverse rendering based on tensor factorization and neural fields."

METHODOLOGY Gs-ir: 3d gaussian splatting for inverse rendering

"We perform relighting experiments using the recovered geometry, material, and illumination from our GS-IR method."

METHODOLOGY Gir: 3d gaussian inverse rendering for relightable scene factorization

"Qualitative evaluation of relighting on TensoIR dataset[12]."

METHODOLOGY Animatable and relightable gaussians for high-fidelity human avatar modeling

"In the current version, we further introduce physically-based rendering (PBR) [19], [20] into our avatar representation for creating both animatable and relightable human avatars."

METHODOLOGY Intrinsicavatar: Physically based inverse rendering of dynamic humans from monocular videos via explicit ray tracing

"Given this progress, physically based inverse rendering of static scenes under unknown natural illumination has been demonstrated [33, 88]."

FOLLOW-UP WORK

Imls-splatting: Efficient mesh reconstruction from multi-view images via point representation

2025 · ACM Transactions on Graphics (TOG) 44 (4), 1-11, 2025 · 1 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Mesh Splatting for End-to-end Multiview Surface Reconstruction	Hong Kong Baptist University	China	—

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology* / *Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

FOLLOW-UP WORK

Openillumination: A multi-illumination dataset for inverse rendering evaluation on real objects

2024 · Advances in Neural Information Processing Systems (NeurIPS), 2024 · 40 citations (GS)

Field-normalised: 37 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2024 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Digital twin catalog: A large-scale photorealistic 3d object digital twin dataset	Meta, Stanford University	United States	—
2	Flash cache: Reducing bias in radiance cache based inverse rendering	Carnegie Mellon University, Google, Google DeepMind	United Kingdom, United States	—
3	Objects with lighting: A real-world dataset for evaluating reconstruction and rendering for object relighting	Adobe, Intel Labs, NVIDIA	China, Hong Kong, United States	Background
4	LightCity: An Urban Dataset for Outdoor Inverse Rendering and Reconstruction under Multi-illumination Conditions	Zhejiang University	China	—
5	PIR: Photometric Inverse Rendering with Shading Cues Modeling and Surface Reflectance Regularization	FNii-Shenzhen, SSE, CUHKSZ, Shenzhen University, Sun Yat-sen University	China	—
6	OLATverse: A Large-scale Real-world Object Dataset with Precise Lighting Control	Max-Planck-Institut, Max Planck Institute for Informatics, Nanjing University	China, Germany	**Influential**
7	POLAR: A Portrait OLAT Dataset and Generative Framework for Illumination-Aware Face Modeling	PICO, Shanghai Jiao Tong University	China	—
8	Seg-invRender: fusing semantic segmentation based on NeRF for inverse rendering considering shadows	Wuhan University	China	—
9	An illumination-robust feature extractor augmented by relightable 3d reconstruction	Hong Kong University of Science and Technology, Peking University	China, Hong Kong	**Methodology**
10	MLI-NeRF: Multi-Light Intrinsic-Aware Neural Radiance Fields	Stony Brook University, Universitat Autonoma de Barcelona	Spain, United States	—
11	UAVLight: A Benchmark for Illumination-Robust 3D Reconstruction in Unmanned Aerial Vehicle (UAV) Scenes	Beijing University of Chemical Technology, Meituan, The Hong Kong University of Science and Technology (Guangzhou)	China	—

No.	Citing paper	Citing institution(s)	Country	S2
12	Baking Relightable NeRF for Real-time Direct/Indirect Illumination Rendering	Seoul National University	South Korea	Methodology
13	A review on 3D Gaussian splatting for sparse view reconstruction	PLA Academy of Military Sciences	China	—
14	Gaussianobject: High-quality 3d object reconstruction from four views with gaussian splatting	Huawei Inc., Shanghai Jiao Tong University, University of Toronto	Canada, China	—
15	Stanford-orb: a real-world 3d object inverse rendering benchmark	Stanford University	United States	—
16	Baking gaussian splatting into diffusion denoiser for fast and scalable single-stage image-to-3d generation and reconstruction	Adobe Research, Johns Hopkins University, The Hong Kong University of Science and Technology	China, United States	—
17	Sparse-view 3D reconstruction: Recent advances and open challenges	Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences	China	—
18	Humanolat: A large-scale dataset for full-body human relighting and novel-view synthesis	Max Planck Institute for Informatics, NVIDIA	Germany, United States	—
19	Auggs: Self-augmented gaussians with structural masks for sparse-view 3d reconstruction	Peking University	China	—
20	Incorporating dense metric depth into neural 3D representations for view synthesis and relighting	Carnegie Mellon University, Toyota Research Institute	United States	—
21	Learning Latent Proxies for Controllable Single-Image Relighting	CUHK, Hong Kong University of Science and Technology, The Hong Kong Polytechnic University	China, Hong Kong	—
22	LumiSculpt: enabling consistent portrait lighting in video generation	Ant Group, Chinese Academy of Sciences	China	—
23	RoboLight: A Dataset with Linearly Composable Illumination for Robotic Manipulation	KTH Royal Institute of Technology	Sweden	—
24	EvalMVX: A Unified Benchmarking for Neural 3D Reconstruction under Diverse Multiview Setups	Beijing University of Posts and Telecommunications, Independent Researcher, Peking University	China, United States	—
25	Reni++ a rotation-equivariant, scale-invariant, natural illumination prior	Friedrich-Alexander-Universität Erlangen-Nürnberg, University of York	Germany, United Kingdom	Methodology
26	A Multi-flash Stereo Camera for Photo-realistic Capture of Small Scenes	Carnegie Mellon University, Toyota Research Institute	United States	—

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology* / *Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

Citing-text excerpts — how the field used this work

METHODOLOGY An illumination-robust feature extractor augmented by relightable 3d reconstruction

"But the collection of data under various illumination conditions [22], [20], [21] in the world is still struggled by (1) the requirement of a large number of images; (2) the complexity and difficulty of artificial construction and precise control of illuminations; and (3) the limitation of specific..."

METHODOLOGY Baking Relightable NeRF for Real-time Direct/Indirect Illumination Rendering

"OpenIllumination benchmark [12] is employed and we select 4 random scenes (bird, metal bucket, pumpkin, and sponge)."

METHODOLOGY Reni++ a rotation-equivariant, scale-invariant, natural illumination prior

"We did not evaluate our model on the OpenIllumination dataset [80] due to its use of non-natural illumination conditions that RENI++ is not trained to represent."

Contribution 2

Claim — Contribution 2

The researcher pioneered high-quality 3D mesh generation and articulated object rigging, establishing a foundational framework for template-free, diffusion-based 3D asset creation widely adopted by independent scholars.

The researcher's core contribution centers on the development of advanced methods for 3D mesh generation and articulated object manipulation, anchored by the seminal 2024 paper "Meshformer." This work introduced a 3D-guided reconstruction model that appears to have set a new standard for generating high-quality meshes, serving as the technical foundation for subsequent innovations in the field.

This line of work addresses the challenge of creating complex, articulated 3D assets without relying on rigid templates or extensive manual rigging. The progression from "Meshformer" to the 2025 follow-ups, "Freeart3d" and "RigAnything," suggests a deliberate expansion from static mesh generation to dynamic, training-free articulated object generation and autoregressive rigging. The titles indicate a shift toward more flexible, template-free approaches that leverage 3D diffusion and autoregressive techniques to handle diverse 3D assets.

The significance of this research is evidenced by its rapid uptake within the academic community. With 74 citations for the core paper and substantial early citations for the follow-ups (36 and 12 respectively), the work has clearly influenced ongoing research. Notably, 97.1% of the 346 classified citations originate from independent researchers, demonstrating that this framework has been widely adopted and built upon by the broader scientific community rather than just the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 88 · **12 flagged influential by Semantic Scholar**

CORE PAPER

Meshformer: High-quality mesh generation with 3d-guided reconstruction model

2024 · Advances in Neural Information Processing Systems (NeurIPS), 2024 · 74 citations (GS)

Field-normalised: 64 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2024 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Structured 3d latents for scalable and versatile 3d generation	Microsoft Research, Tsinghua University, University of Science and Technology of China	China	—
2	Hi3dgen: High-fidelity 3d geometry generation from images via normal bridging	ByteDance, The Chinese University of Hong Kong, Shenzhen, Tsinghua University	China	—
3	Sparseflex: High-resolution and arbitrary-topology 3d shape modeling	The Chinese University of Hong Kong, Tsinghua University, VAST	China	—

No.	Citing paper	Citing institution(s)	Country	S2
4	3dtopia-xl: Scaling high-quality 3d asset generation via primitive diffusion	Nanyang Technological University, Peking University, Shanghai AI Laboratory	China, Singapore	—
5	Dora: Sampling and benchmarking for 3d shape variational auto-encoders	ByteDance, The Hong Kong University of Science and Technology	China	—
6	Cadcrafter: Generating computer-aided design models from unconstrained images	A*STAR, Institute for Infocomm Research, A*STAR, Moxin (Huzhou) Technology Co., LTD., Zhejiang University	China, Singapore, United States	—
7	Generating physically stable and buildable brick structures from text	Carnegie Mellon University	United States	—
8	Kiss3dgen: Repurposing image diffusion models for 3d asset generation	Guangzhou Quwan Network Technology, HKUST(GZ), HKUST(GZ), HKUST	China	—
9	Genxd: Generating any 3d and 4d scenes	Microsoft, Microsoft Corporation, National University of Singapore	Singapore, United States	—
10	Freesplatter: Pose-free gaussian splatting for sparse-view 3d reconstruction	Tencent PCG, The University of Hong Kong	Hong Kong	—
11	Vertexregen: Mesh generation with continuous level of detail	Meta, UC San Diego	United States	—
12	Tar3d: Creating high-quality 3d assets via next-part prediction	Nankai University, Shanghai AI Lab, The Chinese University of Hong Kong	China, Hong Kong	Influential
13	Efficient autoregressive shape generation via octree-based adaptive tokenization	Carnegie Mellon University, Roblox, Stanford University	United States	—
14	Brightdreamer: Generic 3d gaussian generative framework for fast text-to-3d synthesis	HKUST(GZ), Nanyang Technological University	China, Singapore	—
15	Unitex: Universal high fidelity generative texturing for 3d shapes	HKUST, Hong Kong University of Science and Technology, Light Illusion	China, Hong Kong	—
16	PBR-SR: Mesh PBR Texture Super Resolution from 2D Image Priors	Intel Labs, Technical University of Munich	China, Germany	—
17	Cue3d: Quantifying the role of image cues in single-image 3d generation	University of Illinois at Urbana-Champaign	United States	—
18	Dimer: Disentangled mesh reconstruction model	HKUST(GZ), HKUST(GZ), China and HKUST, China, HKUST(GZ), HKUST	China	Influential
19	Collaborative multi-modal coding for high-quality 3d generation	Nanyang Technological University, Shanghai Artificial Intelligence Laboratory	China, Singapore	—
20	Meshgen: Generating pbr textured mesh with render-enhanced auto-encoder and generative data augmentation	Beijing National Research Center for Information Science and Technology (BNRist), Tsinghua University, Beijing Normal University,	China, Hong Kong, Singapore	Influential

No.	Citing paper	Citing institution(s)	Country	S2
		Hong Kong University of Science and Technology		
21	CharacterShot: Controllable and Consistent 4D Character Animation	Nanyang Technological University, National University of Singapore, Shanghai AI Lab	China, Singapore	—
22	Hyper3d: Efficient 3d representation via hybrid triplane and octree feature for enhanced 3d shape variational auto-encoders	Mohammed Bin Zayed University of Artificial Intelligence, Sensory Universe, The Chinese University of Hong Kong	Australia, China, United Arab Emirates	**Influential**
23	Neurally integrated finite elements for differentiable elasticity on evolving domains	NVIDIA, University of Toronto	Canada, France, United States	—
24	3d arena: An open platform for generative 3d evaluation	Hugging Face	United States	—
25	Deepwheel: Generating a 3d synthetic wheel dataset for design and performance evaluation	KAIST	South Korea	—
26	Dreamcar: Leveraging car-specific prior for in-the-wild 3d car reconstruction	City University of Macau, Intel Inc., Li Auto Inc.	Australia, China	—
27	FastAvatar: Towards Unified and Fast 3D Avatar Reconstruction with Large Gaussian Reconstruction Transformers	AKool, Shanghai Jiao Tong University, Tongji University	China	—
28	Track, Inpaint, Resplat: Subject-driven 3D and 4D Generation with Progressive Texture Infilling	Snap Inc., University of Toronto	Canada, United States	**Influential**
29	HiFi-Mesh: High-Fidelity Efficient 3D Mesh Generation via Compact Autoregressive Dependence	Imperial Vision Technology Co. Ltd, Macao Polytechnic University, Shanghai Jiao Tong University	China	—
30	ShapeGen: Towards High-Quality 3D Shape Synthesis	The Chinese University of Hong Kong, Tsinghua University, VAST	China	—

Showing the 30 most-cited of 48 independent citing papers.

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology / Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

FOLLOW-UP WORK

Freeart3d: Training-free articulated object generation using 3d diffusion

2025 · Proceedings of the SIGGRAPH Asia 2025 Conference Papers, 1-13, 2025 · 12 citations (GS)

Field-normalised: 14 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	DragMesh: Interactive 3D Generation Made Easy	Peking University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
2	HOICraft: In-Situ VLM-based Authoring Tool for Part-Level Hand-Object Interaction Design in VR	KAIST, New York University	South Korea, United States	—
3	Pact: Part-decomposed single-view articulated object generation	The Chinese University of Hong Kong, Shenzhen	China	Influential
4	ArtLLM: Generating Articulated Assets via 3D LLM	Hong Kong University of Science and Technology, ShanghaiTech University, Tencent Hunyuan	China, Hong Kong, Singapore	—
5	Particulate: Feed-Forward 3D Object Articulation	University of Cambridge, University of Oxford, University of Oxford, Nanyang Technological University	United Kingdom	Influential
6	MonoArt: Progressive Structural Reasoning for Monocular Articulated 3D Reconstruction	Nanyang Technological University	Singapore	—
7	MorphAny3D: Unleashing the Power of Structured Latent in 3D Morphing	Nanjing University, Peking University	China	—
8	EgoFun3D: Modeling Interactive Objects from Egocentric Videos using Function Templates	Simon Fraser University	Canada	—
9	Lookalike3D: Seeing Double in 3D	Technical University of Munich	Germany	—
10	ArticFlow: Generative Simulation of Articulated Mechanisms	Columbia University	United States	—
11	MotionAnymesh: Physics-Grounded Articulation for Simulation-Ready Digital Twins	Hefei University of Technology, Shanghai Jiao Tong University, The Hong Kong Polytechnic University	China	—

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology / Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

FOLLOW-UP WORK

RigAnything: Template-Free Autoregressive Rigging for Diverse 3D Assets

2025 · ACM Transactions on Graphics (TOG), 2025 · 36 citations (GS)

Field-normalised: 36 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	SV-GS: Sparse View 4D Reconstruction with Skeleton-Driven Gaussian Splatting	University of Minnesota	United States	—
2	SPRig: Self-Supervised Pose-Invariant Rigging from Mesh Sequences	The University of Edinburgh, The University of Hong Kong, University of Pennsylvania	Hong Kong, United Kingdom, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
3	Puppeteer: Rig and animate your 3d models	ByteDance, Institute for Infocomm Research, A*STAR, Nanyang Technological University	China, Singapore	—
4	Physanimator: Physics-guided generative cartoon animation	Netflix, UCLA	United States	—
5	Anymate: A dataset and baselines for learning 3d object rigging	Stanford University	United States	—
6	Physrig: Differentiable physics-based skinning and rigging framework for realistic articulated object modeling	Stability AI, University of Illinois Urbana Champaign, University of Illinois Urbana-Champaign	United Kingdom, United States	—
7	Auto-connect: Connectivity-preserving rigformer with direct preference optimization	Beijing Normal University, Hong Kong University of Science and Technology, National University of Singapore	China, Hong Kong, Singapore	Influential
8	Stable part diffusion 4d: Multi-view rgb and kinematic parts video generation	Stability AI, University of Illinois Urbana-Champaign	United Kingdom, United States	—
9	Gaussian See, Gaussian Do: Semantic 3D Motion Transfer from Multiview Video	NVIDIA, Technion - Israel Institute of Technology	Israel, United States	—
10	3D asset generation: a survey of evolution towards autoregressive and agent-driven paradigms	Beihang University, Nanyang Technological University	China, Singapore	—
11	SkinCells: Sparse Skinning using Voronoi Cells	Meta	Canada, Switzerland, United States	—
12	RigAnyFace: Scaling Neural Facial Mesh Auto-Rigging with Unlabeled Data	Penn State University, Roblox	United States	—
13	Make-It-Poseable: Feed-forward Latent Posing Model for 3D Humanoid Character Animation	University of Science and Technology of China	China	—
14	Rig-Reconstruct-Render (R33D): Collaborative Representation for Editable and Skeleton-Drivable 3D Asset Generation	Shanghai Jiao Tong University	China	—
15	AniGen: Unified Fields for Animatable 3D Asset Generation	The Chinese University of Hong Kong, The University of Hong Kong, Tsinghua University	China, Hong Kong	Influential
16	UNICA: A Unified Neural Framework for Controllable 3D Avatars	Nanjing University	China	—
17	RigMo: Unifying Rig and Motion Learning for Generative Animation	Carnegie Mellon University, Nanyang Technological University, Snap Inc.	Singapore, United States	—
18	Skin Tokens: A Learned Compact Representation for Unified Autoregressive Rigging	Tsinghua University, VAST	China	Influential
19	Stroke3D: Lifting 2D strokes into rigged 3D model via latent diffusion models	Harvard University, Zhejiang University	China, United States	Influential

No.	Citing paper	Citing institution(s)	Country	S2
20	PoissonNet: A Local-Global Approach for Learning on Surfaces	Adobe Research, Brown University, Université de Montréal	Canada, United States	—
21	SOPHY: Learning to Generate Simulation-Ready Objects with Physical Materials	Technical University of Crete, University of Massachusetts Amherst	—	—
22	iTACO: Interactable Digital Twins of Articulated Objects from Casually Captured RGBD Videos	AI, Shanghai Jiao Tong University, Simon Fraser University	Canada, China	—
23	ActionMesh: Animated 3D Mesh Generation with Temporal 3D Diffusion	Meta AI, SpAIitial, University College London	United Kingdom, United States	—
24	MimiCAT: Mimic with Correspondence-Aware Cascade-Transformer for Category-Free 3D Pose Transfer	Institute for Infocomm Research, A*STAR, National University of Singapore, The Chinese University of Hong Kong	China, Singapore	—
25	GaussiAnimate: Reconstruct and Rig Animatable Categories with Level of Dynamics	Macau University of Science and Technology, Nanjing University, The University of Hong Kong	China, Hong Kong	—
26	ViPS: Video-informed Pose Spaces for Auto-Rigged Meshes	Adobe Research, Columbia University, University College London	Canada, United Kingdom, United States	**Influential**
27	Animator-Centric Skeleton Generation on Objects with Fine-Grained Details	Nanyang Technological University, Tencent, Tsinghua Shenzhen International Graduate School	China, Singapore	—
28	Advances in 4D Representation: Geometry, Motion, and Interaction	Simon Fraser University, University of Alberta	Canada	—
29	Scanmove: Motion prediction and transfer for unregistered body meshes	University of Lille, Univ. Lille, Univ. Lille, CNRS, Centrale Lille, Institut Mines-Télécom	France	—

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology / Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

Contribution 3

Claim — Contribution 3

The researcher developed Activezero, a mixed-domain learning framework for active stereovision that eliminates the need for manual annotation, subsequently extending this zero-annotation approach to depth completion.

The researcher established a foundational contribution in computer vision through the 2022 paper 'Activezero,' which introduced mixed-domain learning for active stereovision with zero annotation. This work serves as the core of a specialized research line focused on reducing data labeling burdens in stereo vision systems.

This line of work appears to address the significant challenge of acquiring annotated data for stereo vision tasks. By proposing a method that requires zero annotation, the researcher offered a novel alternative to traditional supervised approaches. The subsequent 2023 follow-up, 'Activezero++,' suggests an expansion of this methodology to include confidence-based depth completion, indicating a progressive refinement and broadening of the initial zero-annotation framework.

The significance of this contribution is evidenced by its adoption within the broader scientific community. With 11 citations for the core paper and 7 for the follow-up, the work has garnered attention from peers. Notably, 97.1% of the researcher's total citing papers originate from independent researchers, suggesting that this specific line of work has resonated beyond the researcher's immediate institutional circle and influenced independent academic inquiry.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 8 · **1 flagged influential by Semantic Scholar**

CORE PAPER

Activezero: Mixed domain learning for active stereovision with zero annotation

2022 · Computer Vision and Pattern Recognition (CVPR), 2022 · 11 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	A survey on deep stereo matching in the twenties	University of Bologna	Italy	Influential
2	Asgrasp: Generalizable transparent object reconstruction and 6-dof grasp detection from rgb-d active stereo camera	Peking University, Samsung R&D Institute China	China	Background
3	Active stereo without pattern projector	University of Bologna	Italy	Background
4	Active Stereo in the Wild through Virtual Pattern Projection: L. Bartolomei et al.	University of Bologna	Italy	—
5	OAfford: One-Shot 3D Object-to-Object Affordance Grounding for Generalizable Robotic Manipulation	University of Virginia	United States	—

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology / Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

FOLLOW-UP WORK

Activezero++: mixed domain learning stereo and confidence-based depth completion with zero annotation

2023 · IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2023 · 7 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	A survey on deep stereo matching in the twenties	University of Bologna	Italy	—
2	Towards real-world aerial vision guidance with categorical 6d pose tracker	Hunan University, Nanyang Technological University	China, Singapore	—
3	Robust depth completion based on Semantic Aggregation: Z. Fu et al.	East China Normal University	China	—

Independent citing papers only; self- and co-author citations excluded. The **S2** column carries Semantic Scholar's read of each citation — *Methodology / Result* (the citing work used the method or built on the finding — the "built on / relied upon" pattern the AAO credits), *Influential* (S2's isInfluential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	28
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	27
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	21
Nanyang Technological University	Singapore	SCImago #137	20
Nanjing University	China	SCImago #178 · THE =62 · QS =103	18
Adobe Research	United States	—	17
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	14
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	14
Peking University	China	SCImago #11 · THE 13 · QS 14	14
The University of Hong Kong	Hong Kong	SCImago #195 · THE 33 · QS 11	13
Stanford University	United States	SCImago #18 · THE =5 · QS 3	12
Hong Kong University of Science and Technology	Hong Kong	SCImago #483 · THE =58 · QS 44	9
Meta	United States	—	9
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	8
University of Illinois Urbana-Champaign	United States	QS =70	8

Geographic distribution of citing authors

Country	Citing papers
China	188
United States	118
United Kingdom	36
Germany	28
Singapore	26
Hong Kong	26
South Korea	16
Canada	15
Switzerland	14
France	11
Japan	10
Australia	9

Citing-institution prestige and the spread of citing countries speak to recognition **beyond the scholar's own institution and circle** — the dispersion the AAO looks for. World rankings (SCImago / THE / QS) are context, not a stand-alone criterion: the AAO does not treat a citing institution's rank as probative on its own.

F. AAO Precedent Considerations

Pre-filing self-check (AAO denial patterns)

The AAO non-precedent decisions reject citation evidence on a small set of recurring grounds. Confirm the petition addresses each before filing:

- ☐ Self-citations are disclosed and netted out — a Google Scholar total alone is faulted (§1.1).
- ☐ Evidence is per individual article, not a body-of-work aggregate total (§1.2).
- ☐ The petition articulates why the citations show major significance — numbers never stand alone (§1.5).
- ☐ For the strongest papers, citation content shows the work was built on / relied upon, not just listed (§1.6, §2.2).
- ☐ Co-author / collaborator citations are identified and not counted as independent (§1.7).
- ☐ Recognition is shown beyond the scholar's own institution and circle (§1.8).
- ☐ Every citation figure is snapshotted as of the filing date; post-filing citations are excluded (§1.9).
- ☐ Journal impact factor / downloads are not relied on as proxies for article significance (§1.10, §1.12).
- ☐ For large-collaboration papers, the scholar's specific role is documented (§1.13).
- ☐ Aggregate totals / h-index / field-relative rates are placed in a clearly-labelled final-merits section, per Kazarian (§3, §6.1.7).

Disclaimer

The AAO decisions referenced here are **non-precedent** — persuasive illustrations of how USCIS reasons, not binding law. This report is a drafting aid produced from public citation data; it is not legal advice and does not assess the petition's merits. All analysis must be reviewed by qualified immigration counsel.

G. Citation Evidence Index

Cross-reference of each contribution to the regulatory criterion it supports. Counsel should map these to the petition's exhibit numbers.

Contribution	Core paper	Indep. cites	Supports
Contribution 1	Tensoir: Tensorial inverse rendering	197	8 CFR 204.5(i)(3) — Outstanding Researcher
Contribution 2	Meshformer: High-quality mesh generation with 3d-guided reconstruction model	88	8 CFR 204.5(i)(3) — Outstanding Researcher
Contribution 3	Activezero: Mixed domain learning for active stereovision with zero annotation	8	8 CFR 204.5(i)(3) — Outstanding Researcher