Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

Minghua Liu

Hillbot

Google Scholar profile

Generated 2026-06-04 by CiteMap. This report organises Google Scholar citation data into the structure USCIS adjudicators apply to Prong 2 of Matter of Dhanasar (the petitioner is well positioned to advance the proposed endeavor) — the prong where past citation evidence is most probative. 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

2,227	2,556	30	21
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.

94.1% independent of 2,021 classified citing papers

Citation type	Count
Independent	1,902
Self-citation	13
Co-author	106
Same-institution	0

206 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 established a foundational simulated part-based interactive environment, subsequently advancing the field with training-free articulated object generation and 3D feature fields for part segmentation.

The researcher's core contribution rests on the 2020 paper 'Sapien: A simulated part-based interactive environment,' which appears to have introduced a significant framework for simulating interactive environments based on object parts. This work serves as the foundation for a sustained line of inquiry into 3D object understanding and manipulation.

This line of work appears to address the challenge of modeling complex, articulated objects in 3D space. The progression from the core simulation environment to follow-up papers like 'Freeart3d' (2025) and 'Partfield' (2025) suggests an evolution toward more efficient, training-free generation methods and advanced feature field learning for segmentation, indicating a deepening technical approach to part-based 3D analysis.

The significance of this contribution is evidenced by the core paper's 879 citations, marking it as a highly influential resource. Furthermore, analysis of citing papers from 2021 reveals that 94.1% of citations originated from independent researchers, demonstrating broad adoption and impact across the wider scientific community beyond the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 534 · **56 flagged influential by Semantic Scholar**

CORE PAPER

Sapien: A simulated part-based interactive environment

2020 · 879 citations (GS)

Field-normalised: 744 Semantic Scholar citations place it in the top 1% of Engineering papers from 2020 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Voxposer: Composable 3d value maps for robotic manipulation with language models	MIT, Stanford University, University of Pennsylvania	United States	—
2	DragMesh: Interactive 3D Generation Made Easy	Peking University	China	—
3	Reartgs: Reconstructing and generating articulated objects via 3d gaussian splatting with geometric and motion constraints	Hefei Institutes of Physical Science Chinese Academy of Sciences, Hefei University of Technology, Shanghai Jiao Tong University	China	—
4	GSWorld: Closed-Loop Photo-Realistic Simulation Suite for Robotic Manipulation	Meta, UC San Diego, UC San Diego, UC Los Angeles	United States	—
5	Advances in 4D Representation: Geometry, Motion, and Interaction	Simon Fraser University, University of Alberta	Canada	—
6	Pact: Part-decomposed single-view articulated object generation	The Chinese University of Hong Kong, Shenzhen	China	—
7	Particulate: Feed-Forward 3D Object Articulation	Nanyang Technological University, University of Cambridge, University of Oxford	United Kingdom	Influential
8	MonoArt: Progressive Structural Reasoning for Monocular Articulated 3D Reconstruction	Nanyang Technological University	Singapore	Influential

No.	Citing paper	Citing institution(s)	Country	S2
9	ArticFlow: Generative Simulation of Articulated Mechanisms	Columbia University	United States	—
10	MotionAnymesh: Physics-Grounded Articulation for Simulation-Ready Digital Twins	Hefei University of Technology, Shanghai Jiao Tong University, The Hong Kong Polytechnic University	China	**Influential**
11	OAfford: One-Shot 3D Object-to-Object Affordance Grounding for Generalizable Robotic Manipulation	University of Virginia	United States	—
12	Advances and challenges in foundation agents: From brain-inspired intelligence to evolutionary, collaborative, and safe systems	Nanyang Technological University, The Hong Kong University of Science and Technology (Guangzhou), University of California, Irvine Medical Center	Canada, China, Singapore	—
13	Objaverse: A universe of annotated 3d objects	Allen Institute for AI, Allen Institute for Artificial Intelligence, Vercept	United States	—
14	Toward general-purpose robots via foundation models: A survey and meta-analysis	Carnegie Mellon University, Georgia Institute of Technology, Georgia Tech	United States	—
15	A survey of embodied ai: From simulators to research tasks	Institute for Infocomm Research, A*STAR, National University of Singapore, University of Washington	Singapore	—
16	Understanding world or predicting future? a comprehensive survey of world models	Tsinghua University	China	—
17	A Survey on Vision–Language–Action Models for Embodied AI	Chinese University of Hong Kong, Huawei, The Chinese University of Hong Kong	China	—
18	Dense policy: Bidirectional autoregressive learning of actions	Shanghai Innovation Institute, Shanghai Jiao Tong University, Xidian University	China	—
19	DeliveryBench: Can Agents Earn Profit in Real World?	University of California, Irvine Medical Center, University of Michigan	United States	—
20	Rtmv: A ray-traced multi-view synthetic dataset for novel view synthesis	NVIDIA, University of Maryland, College Park	United States	—
21	A-sdf: Learning disentangled signed distance functions for articulated shape representation	Huawei, Max Planck Institute for Informatics, UC San Diego	China, Germany, United States	—
22	Learning Disentangled Representations for Controllable Synthesis	University of Toronto, Vanguard University of Southern California	Canada, United States	—
23	Experience grounds language	Carnegie Mellon University	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
24	Embodiedbench: Comprehensive benchmarking multi-modal large language models for vision-driven embodied agents	Northwestern University, Toyota Technological Institute at Chicago, University of Illinois Urbana-Champaign	Canada, United States	—
25	⊠⊠ ProcTHOR: Large-Scale Embodied AI Using Procedural Generation	Allen Institute for AI, Allen Institute for AI, University of Washington, Allen Institute for Artificial Intelligence	United States	—
26	Behavior-1k: A benchmark for embodied ai with 1,000 everyday activities and realistic simulation	Stanford University	United States	**Influential**
27	Dreamart: Generating interactable articulated objects from a single image	BIGAI, Peking University, Peking University; Tsinghua University	China	—
28	Generating Objects with Part-Articulation from a Single Image	Beijing Institute for General Artificial Intelligence, Peking University, Peking University; Tsinghua University	China	—
29	From specialist to generalist: A comprehensive survey on world models	Chinese Academy of Engineering, National University of Defense Technology, Northwestern Polytechnical University	China, United States	—
30	RoboTransfer: Controllable Geometry-Consistent Video Diffusion for Manipulation Policy Transfer	GigaAI, Horizon Robotics, Institute of Automation	China	—

Showing the 30 most-cited of 534 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 · 20 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 independent citing papers resolved for this paper in the current crawl.

FOLLOW-UP WORK

Partfield: Learning 3d feature fields for part segmentation and beyond

2025 · 72 citations (GS)

Field-normalised: 64 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2025 indexed by Semantic Scholar, by citation count.

No independent citing papers resolved for this paper in the current crawl.

Contribution 2

Claim — Contribution 2

The researcher pioneered efficient neural architectures for dense point cloud completion and single-image 3D reconstruction, establishing foundational methods widely adopted by the independent computer vision community.

The researcher's significant contribution centers on advancing geometric deep learning, anchored by the 2020 paper 'Morphing and sampling network for dense point cloud completion.' This core work introduced a novel approach to reconstructing complete 3D shapes from sparse inputs, addressing a fundamental challenge in 3D computer vision. The titles suggest a methodological innovation in handling morphing and sampling processes to achieve dense, high-fidelity outputs.

This line of work demonstrates clear originality and evolution. Building on the foundation of point cloud completion, the researcher extended these principles to broader 3D generation tasks. The 2023 follow-up papers, 'One-2-3-45' and 'Zero123++,' indicate a shift toward rapid, optimization-free 3D mesh generation and consistent multi-view diffusion from single images. This progression suggests a strategic expansion from specific completion tasks to generalizable, efficient 3D synthesis pipelines, highlighting the researcher's ability to scale initial insights into broader, high-impact applications.

The significance of this contribution is evidenced by substantial citation metrics and broad independent adoption. The core 2020 paper has accumulated 477 citations, while the subsequent 2023 works have garnered 777 and 560 citations respectively, indicating rapid and sustained influence. Crucially, analysis of citing literature reveals that 94.1% of citations originate from independent researchers, confirming that this work has become a standard reference point for the wider scientific community rather than merely circulating within the researcher's immediate network.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 1,355 · **128 flagged influential by Semantic Scholar**

CORE PAPER

Morphing and sampling network for dense point cloud completion

2020 · 477 citations (GS)

Field-normalised: 360 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2020 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Comprehensive review of deep learning-based 3D point cloud completion processing and analysis	Fudan University, Shanghai AI Laboratory, Stanford University	China, United States	Influential
2	Snowflake point deconvolution for point cloud completion and generation with skip-transformer	Kuaishou Technology, Tsinghua University, VAST	China, United States	Influential
3	Fbnet: Feedback network for point cloud completion	Hikvision, Hong Kong University of Science and Technology, Shanghai University	China, Hong Kong	—
4	SymmCompletion: High-fidelity and high-consistency point cloud completion with symmetry guidance	HKUST, Hong Kong University of Science and Technology, Sichuan University	China, Hong Kong	—
5	Learnable Chamfer Distance for point cloud reconstruction	Zhejiang University	China	—
6	Softpool++: An encoder–decoder network for point cloud completion	Google, Technical University of Munich, Technische Universität München	Germany, Switzerland	Influential

No.	Citing paper	Citing institution(s)	Country	S2
7	Casfusionnet: A cascaded network for point cloud semantic scene completion by dense feature fusion	Huazhong University of Science and Technology	China	—
8	Reverse2complete: Unpaired multimodal point cloud completion via guided diffusion	Lancaster University, University of Science and Technology of China	China, United Kingdom	—
9	DSPF: Dual-Stage Preservation and Fusion for Source-Free Domain Adaptive Point Cloud Completion	Nanjing Agricultural University, Southeast University, Sun Yat-sen University	China, United States	—
10	Deep learning for 3d point cloud enhancement: A survey	Chang'an University, Huazhong University of Science and Technology, Northwestern Polytechnical University	China	—
11	VQ-DcTr: Vector-quantized autoencoder with dual-channel transformer points splitting for 3D point cloud completion	Fudan University	China	—
12	High-fidelity point cloud completion with low-resolution recovery and noise-aware upsampling	Chinese Academy of Sciences, Tencent America	China, United States	**Influential**
13	Kt-net: knowledge transfer for unpaired 3d shape completion	Tsinghua University, University of Science and Technology of China, Wuhan University	China	—
14	Multi-stage refinement network for point cloud completion based on geodesic attention	Xi'an University of Architecture and Technology	China	—
15	Learning to train a point cloud reconstruction network without matching	Kingston and St George's University, Sheffield Emergency Care Forum, University of Bath	United Kingdom	—
16	Edge-guided generative network with attention for point cloud completion	Xinjiang University	China	—
17	TopologyFormer: structure transformer assisted topology reconstruction for point cloud completion	Chongqing University of Posts and Telecommunications	China	—
18	Point completion by a Stack-Style Folding Network with multi-scaled graphical features	Jiangsu University, University of Electronic Science and Technology of China	China	—
19	Convolutional neural network-based efficient dense point cloud generation using unsigned distance fields	University of Vaasa	Finland	—
20	DuInNet: Dual-Modality Feature Interaction for Point Cloud Completion	Information Engineering University, National University of Defense Technology	China	—
21	Carvenet: Carving point-block for complex 3D shape completion	Alibaba Group, Nanyang Technological University, Tianjin University	Canada, China, P. R. China	—

No.	Citing paper	Citing institution(s)	Country	S2
22	Deep learning for 3d reconstruction, augmentation, and registration: A review paper	University of Tartu, Yıldız Technical University	Estonia, Turkey	—
23	Learning-based 3d reconstruction methods for non-collaborative surfaces—a metrological evaluation	Fondazione Bruno Kessler, University of Trento, University of Udine	Italy	—
24	Adaptive Abstractions for Robust Hierarchical Manipulation Planning	Massachusetts Institute of Technology	United States	—
25	MPED: Quantifying point cloud distortion based on multiscale potential energy discrepancy	Nanjing University, Shanghai Jiao Tong University	China	—
26	Point cloud compression and objective quality assessment: A survey	Shanghai Jiao Tong University, University of Canterbury, University of Missouri–Kansas City	China, New Zealand, United States	—
27	Novel Point Cloud Distance Metric With Adjustable Error Sensitivity and Shape Compensation	Sungkyunkwan University	South Korea	—
28	Advancing 3D point cloud understanding through deep transfer learning: A comprehensive survey	Qatar University, University of Dubai, University of Medea	Algeria, India, Qatar	—
29	Unsupervised point cloud representation learning with deep neural networks: A survey	Nanyang Technological University, Wenzhou University	China, Singapore	—
30	Superpoint network for point cloud oversegmentation	Nanjing University, Nanjing University of Science and Technology	China	—

Showing the 30 most-cited of 394 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

One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization

2023 · 777 citations (GS)

Field-normalised: 674 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2023 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	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	—
2	Materialfusion: Enhancing inverse rendering with material diffusion priors	Carnegie Mellon University, Tel Aviv University	Israel, United States	—
3	SuperCarver: Texture-Consistent 3D Geometry Super-Resolution for High-Fidelity Surface Detail Generation	City University of Hong Kong, Tencent Games, Texas A & M University	China, United States	—

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	**Influential**
7	Track, Inpaint, Resplat: Subject-driven 3D and 4D Generation with Progressive Texture Infilling	Snap Inc., University of Toronto	Canada, United States	—
8	High-fidelity 3D Object Generation from Single Image with RGBN-Volume Gaussian Reconstruction Model	University College London, University of Utah, Zhejiang University	China, United Kingdom, United States	—
9	Unigs: Modeling unitary 3d gaussians for novel view synthesis from sparse-view images	Hong Kong University of Science and Technology, International Digital Economy Academy, The Chinese University of Hong Kong, Shenzhen	China, Hong Kong	—
10	Detailgen3d: Generative 3d geometry enhancement via data-dependent flow	The Chinese University of Hong Kong, Tsinghua University, Tsinghua University, VAST, Sun Yat-sen University	China	—
11	THOM: Generating Physically Plausible Hand-Object Meshes From Text	POSTECH, Ulsan National Institute of Science and Technology, University of Birmingham	South Korea, United Kingdom	—
12	Motion avatar: Generate human and animal avatars with arbitrary motion	La Trobe University, Monash University, The Australian National University	Australia, United Kingdom	—
13	A review on 3D Gaussian splatting for sparse view reconstruction	PLA Academy of Military Sciences	China	—
14	Stable part diffusion 4d: Multi-view rgb and kinematic parts video generation	Stability AI, University of Illinois Urbana-Champaign	United Kingdom, United States	—
15	A survey of multimodal controllable diffusion models	Baidu Inc., Nanjing University, Zhejiang University	China	—
16	Cat3d: Create anything in 3d with multi-view diffusion models	Google, Google DeepMind, Google Research	United Kingdom, United States	—
17	Vd3d: Taming large video diffusion transformers for 3d camera control	KAUST, Snap Inc., University of Toronto	Canada, Italy, United States	—
18	Mvsplat360: Feed-forward 360 scene synthesis from sparse views	ByteDance, ETH Zurich, Monash University	Australia, Singapore, Switzerland	—
19	Zeronvs: Zero-shot 360-degree view synthesis from a single image	Cornell Tech, Google Research, Stanford University	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
20	Viewdiff: 3d-consistent image generation with text-to-image models	Meta, Technical University of Munich	Germany, United States	—
21	Gamba: Marry gaussian splatting with mamba for single-view 3d reconstruction	ByteDance, Nanyang Technological University, National University of Singapore	China, Singapore	—
22	Matrix3d: Large photogrammetry model all-in-one	Apple, Everest Innovation Technology, Nanjing University	China, United States	—
23	SparseGS: Real-time 360 sparse view synthesis using Gaussian splatting	University of California, Irvine Medical Center	United States	—
24	Megascenes: Scene-level view synthesis at scale	Adobe Research, Cornell University, New York University	United States	—
25	Stereodiffusion: Training-free stereo image generation using latent diffusion models	Technical University of Denmark	Denmark	—
26	Scene splatter: Momentum 3d scene generation from single image with video diffusion model	Tencent Inc., Tsinghua University	China	—
27	Cavia: Camera-controllable multi-view video diffusion with view-integrated attention	Apple, Google, The University of Texas at Austin	United States	—
28	Extranerf: Visibility-aware view extrapolation of neural radiance fields with diffusion models	Cornell University, Google Research, National Tsing Hua University	United States	—
29	Efficient-NeRF2NeRF: streamlining text-driven 3D editing with multiview correspondence-enhanced diffusion models	Apple, Facebook AI Research, University at Buffalo, State University of New York	United States	—
30	Movis: Enhancing multi-object novel view synthesis for indoor scenes	BIGAI, Peking University, Peking University; Tsinghua University	China	—

Showing the 30 most-cited of 509 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

Zero123++: a single image to consistent multi-view diffusion base model

2023 · 560 citations (GS)

Field-normalised: 560 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2023 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Materialfusion: Enhancing inverse rendering with material diffusion priors	Carnegie Mellon University, Tel Aviv University	Israel, United States	—
2	SuperCarver: Texture-Consistent 3D Geometry Super-Resolution for High-Fidelity Surface Detail Generation	City University of Hong Kong, Tencent Games, Texas A & M University	China, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
3	Kiss3dgen: Repurposing image diffusion models for 3d asset generation	Guangzhou Quwan Network Technology, HKUST(GZ), HKUST(GZ), HKUST	China	—
4	Tar3d: Creating high-quality 3d assets via next-part prediction	Nankai University, Shanghai AI Lab, The Chinese University of Hong Kong	China, Hong Kong	—
5	Dimer: Disentangled mesh reconstruction model	HKUST(GZ), HKUST(GZ), China and HKUST, China, HKUST(GZ), HKUST	China	Influential
6	CharacterShot: Controllable and Consistent 4D Character Animation	Nanyang Technological University, National University of Singapore, Shanghai AI Lab	China, Singapore	—
7	Dreamcar: Leveraging car-specific prior for in-the-wild 3d car reconstruction	City University of Macau, Intel Inc., Li Auto Inc.	Australia, China	—
8	Track, Inpaint, Resplat: Subject-driven 3D and 4D Generation with Progressive Texture Infilling	Snap Inc., University of Toronto	Canada, United States	—
9	High-fidelity 3D Object Generation from Single Image with RGBN-Volume Gaussian Reconstruction Model	University College London, University of Utah, Zhejiang University	China, United Kingdom, United States	—
10	Detailgen3d: Generative 3d geometry enhancement via data-dependent flow	The Chinese University of Hong Kong, Tsinghua University, Tsinghua University, VAST, Sun Yat-sen University	China	—
11	TopoMesh: High-Fidelity Mesh Autoencoding via Topological Unification	ByteDance, The Hong Kong University of Science and Technology, Tsinghua University	China	—
12	GaMO: Geometry-aware Multi-view Diffusion Outpainting for Sparse-View 3D Reconstruction	National Yang Ming Chiao Tung University	Taiwan	—
13	ACT-R: Adaptive Camera Trajectories for Single View 3D Reconstruction	Simon Fraser University	Canada	—
14	MMPart: Harnessing Multi-Modal Large Language Models for Part-Aware 3D Generation	Iran University of Science and Technology	Iran	—
15	Advances in 4d generation: A survey	University of Nevada Reno, Zhejiang University	China, United States	—
16	Motion avatar: Generate human and animal avatars with arbitrary motion	La Trobe University, Monash University, The Australian National University	Australia, United Kingdom	—
17	Shape-for-motion: Precise and consistent video editing with 3d proxy	City University of Hong Kong, Tencent	China	—
18	Motion 3-to-4: 3D Motion Reconstruction for 4D Synthesis	Adobe Research, Huazhong University of Science and Technology, Westlake University	China, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
19	DualMat: PBR Material Estimation via Coherent Dual-Path Diffusion	Goertek, Meta, OPPO	Australia, China, United States	—
20	Anymate: A dataset and baselines for learning 3d object rigging	Stanford University	United States	—
21	Stable part diffusion 4d: Multi-view rgb and kinematic parts video generation	Stability AI, University of Illinois Urbana-Champaign	United Kingdom, United States	—
22	SOPHY: Learning to Generate Simulation-Ready Objects with Physical Materials	Technical University of Crete, University of Massachusetts Amherst	—	—
23	ArticFlow: Generative Simulation of Articulated Mechanisms	Columbia University	United States	—
24	Learning-based 3d reconstruction methods for non-collaborative surfaces—a metrological evaluation	Fondazione Bruno Kessler, University of Trento, University of Udine	Italy	—
25	Stable virtual camera: Generative view synthesis with diffusion models	Stability AI, University of California, Irvine Medical Center, University of Oxford	United Kingdom, United States	—
26	Collaborative video diffusion: Consistent multi-video generation with camera control	CUHK, Stanford University, The Chinese University of Hong Kong	China, United States	—
27	Camco: Camera-controllable 3d-consistent image-to-video generation	NVIDIA, The University of Texas at Austin, University of Illinois Urbana-Champaign	United States	—
28	Advances in 3d generation: A survey	City University of Hong Kong, South China University of Technology, Tencent	China	—
29	Matrix3d: Large photogrammetry model all-in-one	Apple, Everest Innovation Technology, Nanjing University	China, United States	—
30	Prometheus: 3d-aware latent diffusion models for feed-forward text-to-3d scene generation	The Chinese University of Hong Kong, University of Tübingen, Xiamen University	China, Germany	—

Showing the 30 most-cited of 452 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).

Contribution 3

Claim — Contribution 3

The researcher developed foundational methods for task and path planning in multi-agent pickup and delivery systems, establishing a widely adopted framework for coordinating autonomous agents.

CLAIM: The researcher's seminal 2019 paper, 'Task and path planning for multi-agent pickup and delivery,' represents a core contribution to the field of multi-agent systems, specifically addressing the coordination of agents in logistics contexts. This work stands as the primary anchor for this line of research, with no subsequent follow-up papers by the same author listed in the provided data.

ORIGINALITY: The title suggests the work addresses the complex challenge of simultaneously assigning tasks and determining paths for multiple agents, a critical bottleneck in automated logistics. By integrating task allocation with path planning, the research appears to offer a unified approach to solving coordination problems that were previously treated separately or with less efficiency in multi-agent environments.

SIGNIFICANCE: The impact of this work is evidenced by its substantial citation count of 272. Furthermore, citation analysis reveals that 94.1% of citing papers originate from independent researchers, indicating that the methodology has been widely adopted and validated by the broader scientific community rather than just the researcher's immediate circle. This high degree of independent uptake underscores the work's foundational role in advancing multi-agent planning techniques.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 44 **· 2 flagged influential by Semantic Scholar**

CORE PAPER

Task and path planning for multi-agent pickup and delivery

2019 · 272 citations (GS)

Field-normalised: 169 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2019 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	The atlas for the aspiring network scientist	IT University of Copenhagen	Denmark	—
2	A comprehensive review on leveraging machine learning for multi-agent path finding	University of Cambridge	United Kingdom	—
3	Research challenges and opportunities in multi-agent path finding and multi-agent pickup and delivery problems	Ben-Gurion University of the Negev, Technion – Israel Institute of Technology	Israel	—
4	Lifelong multi-agent path finding in large-scale warehouses	Amazon Robotics, Carnegie Mellon University, University of Southern California	United States	—
5	Group-based distributed auction algorithms for multi-robot task assignment	Delft University of Technology	Netherlands	—
6	PRIMAL: Pathfinding Via Reinforcement and Imitation Multi-Agent Learning - Lifelong	National University of Singapore	Singapore	—
7	Learn to follow: Decentralized lifelong multi-agent pathfinding via planning and learning	AIRI, Federal Research Center for Computer Science and Control of Russian Academy of Sciences	Russia	—
8	Multi-agent pathfinding with continuous time	Omron (Japan)	Japan	—
9	Multi-agent path finding–an overview	Ben-Gurion University of the Negev	Israel	—
10	A review of graph-based multi-agent pathfinding solvers: From classical to beyond classical	Chinese Academy of Sciences, Harbin Institute of Technology	China	—
11	Robotic mobile fulfillment systems: A survey on recent developments and research opportunities	Universidade Federal da Paraíba	Brasil	—

No.	Citing paper	Citing institution(s)	Country	S2
12	Efficient large-scale multi-drone delivery using transit networks	Aeronautics and Astronautics Research Center, Stanford University, Technion	Israel, Portugal, United States	—
13	Reinforcement-learning-assisted multi-UAV task allocation and path planning for IIoT	Beihang University, North China University of Technology, Zhengzhou University	China	—
14	Improved Heuristics for Multi-Agent Path Finding with Conflict-Based Search.	Ben-Gurion University of the Negev, Carnegie Mellon University, Simon Fraser University	Canada, Israel, United States	—
15	CBSS: A new approach for multiagent combinatorial path finding	Carnegie Mellon University, Texas A&M University	United States	—
16	Scaling lifelong multi-agent path finding to more realistic settings: Research challenges and opportunities	Carnegie Mellon University	United States	—
17	Multi-goal multi-agent pickup and delivery	Carnegie Mellon University, Simon Fraser University, University of Southern California	Canada, United States	**Influential**
18	Guidance graph optimization for lifelong multi-agent path finding	Carnegie Mellon University, University of Southern California	United States	—
19	Efficient task planning for heterogeneous AGVs in warehouses	Shenzhen Institute of Information Technology | Harbin Institute of Technology	China	—
20	Which MAPF model works best for automated warehousing?	Carnegie Mellon University, UC San Diego, University of Southern California	United States	—
21	Resolving conflicting constraints in multi-agent reinforcement learning with layered safety	Massachusetts Institute of Technology, University of California, Irvine Medical Center	United States	—
22	Integrated task allocation and path coordination for large-scale robot networks with uncertainties	Chinese University of Hong Kong, Shanghai Jiao Tong University	China	—
23	Double-deck multi-agent pickup and delivery: Multi-robot rearrangement in large-scale warehouses	Simon Fraser University	Canada	—
24	Arbitrarily scalable environment generators via neural cellular automata	Carnegie Mellon University, University of Southern California	United States	—
25	Online guidance graph optimization for lifelong multi-agent path finding	Carnegie Mellon University, Monash University, Tsinghua University	Australia, China, United States	—
26	A TSP-based online algorithm for multi-task multi-agent pickup and delivery	Osaka Metropolitan University	Japan	—
27	Multi-robot pickup and delivery via distributed resource allocation	University of Bologna	Italy	—

No.	Citing paper	Citing institution(s)	Country	S2
28	Solving simultaneous target assignment and path planning efficiently with time-independent execution	Tokyo Institute of Technology, University of Cambridge	Japan, United Kingdom	—
29	A coordinated scheduling approach for task assignment and multi-agent path planning	Kunming University of Science and Technology	China	—
30	Cooperative multi-agent path finding: Beyond path planning and collision avoidance	Technion, Technion – Israel Institute of Technology	Israel	—

Showing the 30 most-cited of 44 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).

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	179
Peking University	China	SCImago #11 · THE 13 · QS 14	127
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	104
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	95
Nanyang Technological University	Singapore	SCImago #137	94
Stanford University	United States	SCImago #18 · THE =5 · QS 3	86
University of California, Irvine Medical Center	United States	—	77
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	77
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	60
The University of Hong Kong	Hong Kong	SCImago #195 · THE 33 · QS 11	56
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	53
UC San Diego	United States	—	52
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	51
NVIDIA	United States	—	43
Beihang University	China	SCImago #160 · THE 251–300 · QS =388	42

Geographic distribution of citing authors

Country	Citing papers
China	1,136
United States	734
Singapore	157
United Kingdom	157

Country	Citing papers
Germany	131
Hong Kong	122
Canada	105
Australia	57
South Korea	56
Switzerland	46
Japan	44
Italy	31

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	Sapien: A simulated part-based interactive environment	534	Dhanasar — Prong 2 (well-positioned)
Contribution 2	Morphing and sampling network for dense point cloud completion	1,355	Dhanasar — Prong 2 (well-positioned)
Contribution 3	Task and path planning for multi-agent pickup and delivery	44	Dhanasar — Prong 2 (well-positioned)