

Citation Evidence Report

EB-1B Petition — Outstanding Professor or Researcher

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


Bo Yang

Asst Prof, Hong Kong Polytechnic University; DPhil, University of Oxford

[Google Scholar profile](#)

Generated 2026-06-08 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



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.

93.3% independent of 2,164 classified citing papers

Citation type	Count
Independent	2,020
Self-citation	27
Co-author	117
Same-institution	0

569 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 developed efficient semantic segmentation methods for large-scale point clouds, establishing a foundational framework that enabled subsequent advancements in 3D registration and unsupervised learning.

The researcher's primary contribution centers on the development of efficient semantic segmentation techniques for large-scale point clouds, anchored by the seminal 2020 paper RandLA-Net. This work serves as the cornerstone for a sustained research trajectory addressing the computational challenges inherent in processing massive 3D data structures.

This line of work appears to address the critical gap in handling the scale and complexity of 3D point cloud data. The progression from RandLA-Net to SpinNet (2021) and GrowSP (2023) suggests a deliberate expansion from segmentation to broader 3D understanding, including general surface descriptors for registration and unsupervised segmentation methods. The titles indicate a focus on creating robust, scalable solutions that generalize across different 3D processing tasks.

The significance of this contribution is evidenced by the substantial uptake of the core work, which has accumulated 2,864 citations. Notably, 93.3% of the citing papers originate from independent researchers, indicating that the methodology has been widely adopted and validated by the broader scientific community beyond the researcher's immediate circle. The continued relevance is further supported by the citation counts of the follow-up papers, demonstrating sustained impact in the field.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 1,244 · 139 flagged influential by Semantic Scholar

CORE PAPER

[RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds](#)

2020 · 2,864 citations (GS)

Field-normalised: 1,933 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	Point transformer v3: Simpler faster stronger	Moscow Institute of Thermal Technology, Peking University, The Chinese University of Hong Kong, Shenzhen	China, Hong Kong, Russia	—
2	Deep learning for LiDAR-only and LiDAR-fusion 3D perception: A survey	Tongji University	China	—
3	Deep learning for scene flow estimation on point clouds: A survey and prospective trends	Bournemouth University, Xi'an Jiaotong-Liverpool University	China, United Kingdom	—
4	Pointseg: A training-free paradigm for 3d scene segmentation via foundation models	Tencent	China	—
5	Link: Linear kernel for lidar-based 3d perception	Brown University, Nanjing University	China, United States	—
6	Lidarformer: A unified transformer-based multi-task network for lidar perception	Tusimple, University of Central Florida	PR China, United States	—
7	Deep-learning-based approaches for semantic segmentation of natural scene images: A review	Ankara University, Tampere University	Finland, Turkey	—
8	Learning to generate realistic lidar point clouds	University of Illinois at Urbana-Champaign	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
9	Transfer learning from synthetic to real lidar point cloud for semantic segmentation	Nanyang Technological University	Singapore	—
10	From slam to situational awareness: Challenges and survey	University of Luxembourg	Luxembourg	—
11	Efficient 3d deep lidar odometry	Shanghai Jiao Tong University	China	Influential
12	A survey on deep learning fundamentals	Harbin Institute of Technology	China	—
13	Transformer-based visual segmentation: A survey	Fudan University, Nanyang Technological University, Shanghai AI Laboratory	China, Singapore, United Kingdom	—
14	Rangevit: Towards vision transformers for 3d semantic segmentation in autonomous driving	Centre de Robotique, Ecole des Ponts, valeo.ai	France	—
15	Pointnext: Revisiting pointnet++ with improved training and scaling strategies	Amazon, King Abdullah University of Science and Technology (KAUST), Microsoft Research	Saudi Arabia, United States	Influential
16	Ulip: Learning a unified representation of language, images, and point clouds for 3d understanding	Apple, Salesforce, Salesforce AI	United States	—
17	Rethinking range view representation for lidar segmentation	Nanyang Technological University, National University of Singapore, Shanghai AI Laboratory	China, Singapore	—
18	Cross-modal unsupervised domain adaptation for 3D semantic segmentation via multi-scale fusion-then-distillation	PLA Army Engineering University	China	—
19	Searching efficient 3d architectures with sparse point-voxel convolution	Massachusetts Institute of Technology, MIT, Moscow Institute of Thermal Technology	China, Russia, United States	—
20	Octformer: Octree-based transformers for 3d point clouds	Peking University	China	—
21	Squeezesegv3: Spatially-adaptive convolution for efficient point-cloud segmentation	Facebook, Meta Research, University of California, Irvine Medical Center	United States	Influential
22	Frnet: Frustum-range networks for scalable lidar segmentation	Nanjing University of Aeronautics and Astronautics, Nanjing University of Information Science and Technology, Nanjing University of Posts and Telecommunications	China, Singapore	—
23	Enable deep learning on mobile devices: Methods, systems, and applications	Massachusetts Institute of Technology, MIT, Moscow Institute of Thermal Technology	Russia, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
24	Vehicle detection algorithms for autonomous driving: A review	Harbin Institute of Technology (Shenzhen), Henan University of Technology	China	—
25	Deep learning based 3D segmentation in computer vision: A survey	Hunan Normal University, Hunan University, The Australian National University	Australia, China	—
26	Vehicle detection for autonomous driving: A review of algorithms and datasets	University of Science and Technology of China	China	—
27	3d-mininet: Learning a 2d representation from point clouds for fast and efficient 3d lidar semantic segmentation	Universidad de Zaragoza	Spain	—
28	Deep learning on 3d semantic segmentation: A detailed review	National Technical University of Athens, Université de Strasbourg	France, Greece	Influential
29	Deep-learning-based point cloud semantic segmentation: A survey	North China University of Water Resources and Electric Power	China	—
30	Spiking point transformer for point cloud classification	Tencent Inc., University of Science and Technology of China, Zhejiang University	China	—

Showing the 30 most-cited of 808 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

[SpinNet: Learning a General Surface Descriptor for 3D Point Cloud Registration](#)

2021 · 537 citations (GS)

Field-normalised: 379 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2021 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Cmhanet: A cross-modal hybrid attention network for point cloud registration	China Electric Power Research Institute, Xi'an Jiaotong University	China	Influential
2	PointDifformer: Robust point cloud registration with neural diffusion and transformer	Continental Automotive Singapore Pte. Ltd., Nanyang Technological University	Singapore	Influential
3	LiDAR Point Cloud Registration Based on Pyramid Compatibility Graph Optimization With Coarse-to-Fine Feature Extraction	—	—	—
4	CGReg: Classification-Guided Point Cloud Registration via Equivariant Learning	Zhejiang Normal University, Zhejiang University of Finance and Economics	China	—

No.	Citing paper	Citing institution(s)	Country	S2
5	Weakly supervised 3d scene segmentation with region-level boundary awareness and instance discrimination	Chinese University of Hong Kong, City University of Hong Kong, The Chinese University of Hong Kong	China, Hong Kong	—
6	Gipso: Geometrically informed propagation for online adaptation in 3d lidar segmentation	University of Trento	Italy	—
7	Boosting few-shot 3d point cloud segmentation via query-guided enhancement	Harbin Institute of Technology	China	—
8	General 3D Vision-Language Model With Fast Rendering and Pre-Training Vision-Language Alignment	California Institute of Technology, Peking University, Tsinghua University	China, United States	—
9	Weaklabel3d-net: A complete framework for real-scene lidar point clouds weakly supervised multi-tasks understanding	Chinese University of Hong Kong, City University of Hong Kong, Wuhan University	China, Hong Kong	—
10	Dual-feature attention for robust point cloud registration: integrating transformation-variant and invariant features	Shenzhen University	China	—
11	A review of non-rigid transformations and learning-based 3D point cloud registration methods	University Hospital Heidelberg	Germany	—
12	Multiway non-rigid point cloud registration via learned functional map synchronization	NVIDIA, Princeton University, Stanford University	China, United States	—
13	p^ 3-net: Part mobility parsing from point cloud sequences via learning explicit point correspondence	Baidu Research, Beihang University, State Key Laboratory of Virtual Reality Technology and Systems	China	—
14	Oaaformer: Robust and efficient point cloud registration through overlapping-aware attention in transformer	Qingdao University of Science and Technology, Shandong University, Shandong University of Science and Technology	China, United States	—
15	3D point cloud registration with multi-scale architecture and unsupervised transfer learning	École Nationale Supérieure des Mines de Paris, Université Paris Sciences et Lettres	France	Influential
16	TSG-Seg: Temporal-selective guidance for semi-supervised semantic segmentation of 3D LiDAR point clouds	Nanyang Technological University, Nara Institute of Science and Technology, RIKEN Center for Advanced Intelligence Project	Japan, Singapore	—
17	Robust point cloud registration framework based on deep graph matching	Fudan University, Shandong Academy of Sciences Qilu University of Technology	China	—
18	Regformer: An efficient projection-aware transformer network for large-scale point cloud registration	China University of Mining and Technology, ETH Zurich, Shanghai Jiao Tong University	China, Switzerland	Influential
19	Lepard: Learning partial point cloud matching in rigid and deformable scenes	The University of Tokyo	Japan	—

No.	Citing paper	Citing institution(s)	Country	S2
20	Peal: Prior-embedded explicit attention learning for low-overlap point cloud registration	Hangzhou Dianzi University, Shanghai Jiao Tong University, Zhejiang Gongshang University	China	Influential
21	Frame averaging for invariant and equivariant network design	Facebook AI Research, McGill University, Meta AI	Canada, United States	—
22	Unsupervised deep probabilistic approach for partial point cloud registration	ETH Zurich, Fondazione Bruno Kessler, Peking University	China, Italy, Switzerland	—
23	Back to the feature: classical 3d features are (almost) all you need for 3d anomaly detection	Hebrew University of Jerusalem, The Hebrew University of Jerusalem	Israel	—
24	Freeze: Training-free zero-shot 6d pose estimation with geometric and vision foundation models	Fondazione Bruno Kessler, University of Trento	Italy	—
25	Turboreg: Turboclique for robust and efficient point cloud registration	Beijing Polytechnic, State Key Laboratory of Information Engineering in Surveying Mapping and Remote Sensing, University of Zaragoza	China	—
26	PANet: A point-attention based multi-scale feature fusion network for point cloud registration	Xidian University	China	—
27	Fastmac: Stochastic spectral sampling of correspondence graph	Peking University, Shanghai AI Lab, Shanghai Jiao Tong University	China, United States	—
28	Dual focus-attention transformer for robust point cloud registration	Fudan University, Qilu University of Technology Shandong Academy of Sciences, Shandong Academy of Sciences Qilu University of Technology	China, People's Republic of China	—
29	Riga: Rotation-invariant and globally-aware descriptors for point cloud registration	National University of Defense Technology, Shanghai Jiao Tong University, Technical University of Munich	China, Germany	Influential
30	Colorpcr: Color point cloud registration with multi-stage geometric-color fusion	Shanghai AI Laboratory, Tsinghua University, Xi'an Jiaotong University	China	—

Showing the 30 most-cited of 436 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

GrowSP: Unsupervised Semantic Segmentation of 3D Point Clouds

2023 · 105 citations (GS)

Field-normalised: 64 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2023 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 general radiance fields for 3D representation, subsequently extending this framework to dynamic physics learning and scene geometry decomposition.

The researcher established a foundational approach to 3D representation and rendering through the seminal 2021 paper 'GRF: Learning a General Radiance Field for 3D Representation and Rendering.' This core work serves as the anchor for a sustained line of inquiry into neural scene representations.

This line of work appears to address the challenge of generalizing 3D scene understanding beyond static geometry. The titles of subsequent publications suggest a logical progression from general radiance fields to more complex domains. Specifically, the 2023 paper 'NVFi: Neural Velocity Fields for 3D Physics Learning from Dynamic Videos' indicates an expansion into dynamic physical phenomena, while 'DM-NeRF: 3D Scene Geometry Decomposition and Manipulation from 2D Images' suggests advancements in structural analysis and manipulation capabilities.

The significance of this contribution is evidenced by substantial academic uptake. The core paper has accumulated 420 citations, while the follow-up works have garnered 36 and 99 citations respectively. Crucially, analysis of citing literature reveals that 93.3% of citations originate from independent researchers, demonstrating that this framework has been widely adopted and built upon by the broader scientific community rather than solely by the researcher's immediate collaborators.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 328 · 12 flagged influential by Semantic Scholar

CORE PAPER

[GRF: Learning a General Radiance Field for 3D Representation and Rendering](#)

2021 · 420 citations (GS)

Field-normalised: 257 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2021 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Enhance-NeRF: Multiple Performance Evaluation for Neural Radiance Fields	Nanjing Agricultural University, Nanjing University of Information Science and Technology	China	—
2	Advances in feed-forward 3d reconstruction and view synthesis: A survey	Caltech, Harvard University, Hillbot	China, Germany, Singapore	—
3	Zerorf: Fast sparse view 360deg reconstruction with zero pretraining	UC San Diego, University of California San Diego	United States	—
4	Neural fields in visual computing and beyond	Brown University, Google, Massachusetts Institute of Technology	Canada, Israel, United States	—
5	Neural rays for occlusion-aware image-based rendering	Max Planck Institute for Informatics, Texas A&M University, The University of Hong Kong	China, Germany, Hong Kong	—
6	Nerf-sr: High quality neural radiance fields using supersampling	Kuaishou Technology, Tsinghua University, University of Pennsylvania	China, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
7	Nope-nerf: Optimising neural radiance field with no pose prior	Mohammed Bin Zayed University of Artificial Intelligence, University of Oxford	United Arab Emirates, United Kingdom	—
8	Mine: Towards continuous depth mpi with nerf for novel view synthesis	ByteDance, Fudan University, National University of Singapore	China, Singapore	—
9	Multi-view consistent generative adversarial networks for 3d-aware image synthesis	DAMO Academy, Alibaba Group, National University of Singapore, University of Technology Sydney	Australia, China, Singapore	—
10	Multi-view consistent generative adversarial networks for compositional 3d-aware image synthesis	DAMO Academy, Alibaba Group, National University of Singapore, University of Technology Sydney	Australia, Singapore, United States	—
11	Flowcam: Training generalizable 3d radiance fields without camera poses via pixel-aligned scene flow	Harvard University, Massachusetts Institute of Technology, Stanford University	United States	—
12	I-DACS: Always maintaining consistency between poses and the field for radiance field construction without pose prior	Tongji University, University of Macau	China	—
13	3D hierarchical refinement and augmentation for unsupervised learning of depth and pose from monocular video	Shanghai Jiao Tong University	China	—
14	Mononerf: Learning generalizable nerfs from monocular videos without camera poses	Meta AI, UC San Diego, University of California, Irvine Medical Center	United States	—
15	pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction	Massachusetts Institute of Technology, Simon Fraser University	Canada, United States	—
16	Tensorf: Tensorial radiance fields	Adobe Research, ShanghaiTech University, University of California San Diego	China, Germany, United States	—
17	DL3dv-10k: A large-scale scene dataset for deep learning-based 3d vision	Adobe Inc., Google Inc., Purdue University	United States	—
18	Reconfusion: 3d reconstruction with diffusion priors	Columbia University, Google, Google DeepMind	United Kingdom, United States	—
19	Plenotrees for real-time rendering of neural radiance fields	UC Berkeley, University of California, Irvine Medical Center	United States	—
20	Generative novel view synthesis with 3d-aware diffusion models	NVIDIA, Stanford University	United States	—
21	Ibrnet: Learning multi-view image-based rendering	Cornell University, Google DeepMind, Google Research	United Kingdom, United States	—
22	Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps	ETH Zurich, Max Planck Institute for Intelligent Systems, Universität Tübingen	China, Germany, Switzerland	—
23	Mvimgnet: A large-scale dataset of multi-view images	ByteDance, Sun Yat-sen University, The Chinese University of Hong Kong, Shenzhen	China	—

No.	Citing paper	Citing institution(s)	Country	S2
24	Advances in neural rendering	Google Inc., Google Research, Massachusetts Institute of Technology	Germany, United States	—
25	Wonderworld: Interactive 3d scene generation from a single image	Cornell Tech, MIT, Stanford University	United States	—
26	Urban radiance fields	Google, Google Research, University of Toronto	Canada, United States	—
27	Multidiff: Consistent novel view synthesis from a single image	Max Planck Society, Meta, Technical University of Munich	Germany, Switzerland	—
28	Diffusion with forward models: Solving stochastic inverse problems without direct supervision	Massachusetts Institute of Technology, Princeton University	United States	—
29	DONeRF: Towards real-time rendering of compact neural radiance fields using depth oracle networks	Facebook, Graz University of Technology	Austria, United States	—
30	Nex: Real-time view synthesis with neural basis expansion	VISTEC	Thailand	—

Showing the 30 most-cited of 328 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 is Influential signal, Valenzuela et al. 2015), or *Background* (a passing mention).

FOLLOW-UP WORK

[NVFi: Neural Velocity Fields for 3D Physics Learning from Dynamic Videos](#)

2023 · 36 citations (GS)

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

FOLLOW-UP WORK

[DM-NeRF: 3D Scene Geometry Decomposition and Manipulation from 2D Images](#)

2023 · 99 citations (GS)

Field-normalised: 86 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2023 indexed by Semantic Scholar, by citation count.

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

Contribution 3

Claim — Contribution 3

The researcher pioneered adversarial learning for single-view 3D reconstruction, establishing a foundational framework that subsequent work extended to dense reconstruction and non-rigid physical dynamics.

The researcher established a seminal contribution in 3D computer vision through the 2017 paper '3D Object Reconstruction from a Single Depth View with Adversarial Learning.' This core work introduced a novel approach to reconstructing three-dimensional objects from limited depth data, leveraging adversarial learning techniques to enhance reconstruction quality and realism.

This line of work appears to address the challenge of generating complete 3D models from sparse or single-view inputs. The originality lies in applying adversarial training to this specific reconstruction problem. The researcher subsequently expanded this foundation in 2018 with 'Dense 3D Object Reconstruction from a Single Depth View,' suggesting an evolution toward higher-fidelity outputs. Additionally, the 2018 paper '3D-PhysNet' indicates a broadening of scope to incorporate intuitive physics for non-rigid deformations, demonstrating the versatility of the initial methodological framework.

The significance of this contribution is evidenced by substantial academic uptake. The core 2017 paper has accumulated 246 citations, while the 2018 dense reconstruction follow-up has garnered 205 citations. With 93.3% of the researcher's total classified citations originating from independent researchers, this indicates that the community widely recognizes and builds upon this specific line of inquiry, validating its impact beyond the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 70 · 6 flagged influential by Semantic Scholar

CORE PAPER

[3D Object Reconstruction from a Single Depth View with Adversarial Learning](#)

2017 · 246 citations (GS)

Field-normalised: 206 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2017 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Image synthesis with adversarial networks: A comprehensive survey and case studies	École de Technologie Supérieure, Massey University, Shanghai Jiao Tong University	Canada, China, New Zealand	—
2	Softpool++: An encoder–decoder network for point cloud completion	Google, Technical University of Munich, Technische Universität München	Germany, Switzerland	Influential
3	Reverse2complete: Unpaired multimodal point cloud completion via guided diffusion	Lancaster University, University of Science and Technology of China	China, United Kingdom	—
4	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	—
5	Artificial intelligence in the creative industries: a review	University of Bristol	United Kingdom	—
6	Pu-gan: a point cloud upsampling adversarial network	Huazhong University of Science and Technology, Hunan University, Tel Aviv University	China, Israel	—
7	GAN-based generation of realistic 3D volumetric data: A systematic review and taxonomy	University Medicine Essen, University Medicine Essen; Graz University of Technology, University of Minho	Germany, Germany; Austria, Portugal	—
8	PUFA-GAN: A frequency-aware generative adversarial network for 3D point cloud upsampling	City University of Hong Kong, De Montfort University, Peking University	China, United Kingdom	—
9	Escaping plato's cave: 3d shape from adversarial rendering	Google, University College London	United Kingdom, United States	—
10	A survey of image synthesis and editing with generative adversarial networks	Tsinghua University, University of Bath	China, United Kingdom	—

No.	Citing paper	Citing institution(s)	Country	S2
11	Convolutional generation of textured 3d meshes	ETH Zurich, KU Leuven	Belgium, Switzerland	—
12	Unsupervised point cloud representation learning with deep neural networks: A survey	Nanyang Technological University, Wenzhou University	China, Singapore	—
13	Towards point cloud completion: Point rank sampling and cross-cascade graph cnn	China University of Petroleum, Peking University	China	—
14	Dso: Aligning 3d generators with simulation feedback for physical soundness	Nanyang Technological University, University of Oxford	United Kingdom	—
15	Amodal3r: Amodal 3d reconstruction from occluded 2d images	Nanyang Technological University, Singapore Institute of Technology, University of Cambridge	Singapore, United Kingdom	—
16	Advancements in point cloud data augmentation for deep learning: A survey	Xi'an Jiaotong-Liverpool University	China	—
17	Pointnr: Diverse point cloud completion with geometry-aware transformers	Tsinghua University	China	—
18	Learning local displacements for point cloud completion	Google, Technical University of Munich, Technische Universität München	Germany, Switzerland	—
19	FinerPCN: High fidelity point cloud completion network using pointwise convolution	Xidian University	China	—
20	Point cloud completion via relative point position encoding and regional attention	Huazhong University of Science and Technology	China	—
21	CompleteDT: Point cloud completion with information-perception transformers	Beijing Institute of Optoelectronic Technology, Beijing Institute of Technology	China	—
22	Human body shape completion with implicit shape and flow learning	Centre Inria de l'Université Grenoble Alpes, Tsinghua University	China, France	—
23	MIX-NET: Deep learning-based point cloud processing method for segmentation and occlusion leaf restoration of seedlings	Agricultural Genomics Institute at Shenzhen, Ministry of Education of the People's Republic of China, Nanjing Normal University	China	—
24	SDA-Net: A global feature point cloud completion network based on serialization and dual attention	China University of Geosciences	China	—
25	Mendnet: Restoration of fractured shapes using learned occupancy functions	Clarkson University	United States	—
26	Structure-aware completion of plant 3D LiDAR point clouds via a multi-resolution GAN-inversion network	Bureau of Agriculture and Rural Affairs, China Agricultural University, Nanjing Institute of Agricultural Mechanization	China	—

No.	Citing paper	Citing institution(s)	Country	S2
27	RefComp: A Reference-guided Unified Framework for Unpaired Point Cloud Completion	ETH Zurich, Southern University of Science and Technology, SUSTech	China, Switzerland, United Kingdom	—
28	FEPoinTr: Point cloud completion networks independent of batch size	Guilin University of Electronic Technology, Guilin University of Technology	China	—
29	Implicit geometric regularization for learning shapes	Weizmann Institute of Science	Israel	—
30	Behave: Dataset and method for tracking human object interactions	Google Research, Max Planck Institute for Informatics, Tübingen AI Center	Germany, United States	—

Showing the 30 most-cited of 70 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

[3D-PhysNet: Learning the Intuitive Physics of Non-Rigid Object Deformations](#)

2018 · 41 citations (GS)

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

FOLLOW-UP WORK

[Dense 3D Object Reconstruction from a Single Depth View](#)

2018 · 205 citations (GS)

Field-normalised: 131 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2018 indexed by Semantic Scholar, by citation count.

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

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	102
Nanyang Technological University	Singapore	SCImago #137	97
Wuhan University	PR China	SCImago #80 · THE =122 · QS 186	89
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	78
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	74
Sun Yat-sen University	China	SCImago #40 · THE 201–250 · QS =276	64
University of Oxford	United Kingdom	SCImago #26 · THE 1 · QS 4	62
Chinese Academy of Sciences	China	SCImago #2	58

Institution	Country	World ranking	Citing papers
Xiamen University	China	SCImago #275 · THE 251–300 · QS 341	56
Peking University	China	SCImago #11 · THE 13 · QS 14	53
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	48
The Chinese University of Hong Kong	Hong Kong	SCImago #163 · THE =41 · QS =32	46
Stanford University	United States	SCImago #18 · THE =5 · QS 3	45
The Hong Kong Polytechnic University	Hong Kong	SCImago #256 · THE 80 · QS 54	42
National University of Defense Technology	China	SCImago #488	41

Geographic distribution of citing authors

Country	Citing papers
China	1,402
United States	429
United Kingdom	174
Germany	151
Singapore	146
Hong Kong	127
Canada	87
Australia	87
South Korea	67
Switzerland	64
France	47
Italy	34

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	RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds	1,244	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 2	GRF: Learning a General Radiance Field for 3D Representation and Rendering	328	8 CFR 204.5(i)(3) – Outstanding Researcher
Contribution 3	3D Object Reconstruction from a Single Depth View with Adversarial Learning	70	8 CFR 204.5(i)(3) – Outstanding Researcher