

Citation Evidence Report

EB-2 NIW Petition — National Interest Waiver

Matter of Dhanasar · Prong 2 (well-positioned)

Junyuan (Jason) Hong

Massachusetts General Hospital & Harvard Med; National University of Singapore

[Google Scholar profile](#)

Generated 2026-05-21 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

1,824 Citing papers mapped	1,939 Citation edges	59 Home papers mapped	21 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.

91.7% independent of 950 classified citing papers

Citation type	Count
Independent	871
Self-citation	9
Co-author	68
Same-institution	2

874 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 data-free knowledge distillation for heterogeneous federated learning, establishing a foundational framework for privacy-preserving model aggregation that subsequent work extended to robustness and efficiency.

The researcher's core contribution rests on the 2021 paper 'Data-free knowledge distillation for heterogeneous federated learning,' which appears to address the challenge of aggregating models across diverse devices without accessing raw client data. This work establishes a methodological baseline for handling heterogeneity in federated settings while preserving data privacy.

Originality in this line of work is suggested by the chronological progression from the core paper to follow-up studies. The 2022 paper on 'Efficient Split-Mix Federated Learning' and the 2023 paper on 'Federated robustness propagation' indicate that the researcher expanded the initial framework to address on-demand customization and adversarial robustness. This trajectory suggests a systematic effort to enhance the practicality and security of the original data-free distillation approach.

The significance of this contribution is evidenced by the core paper's 1,284 citations, indicating substantial uptake in the field. Furthermore, citation analysis reveals that 95.9% of citing papers originate from independent researchers, demonstrating that the work has influenced a broad, external community rather than relying on self-citation or institutional clustering.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 349 · 37 flagged influential by Semantic Scholar

CORE PAPER

[Data-free knowledge distillation for heterogeneous federated learning](#)

2021 · International conference on machine learning, 12878-12889, 2021 · 1,284 citations (GS)

Field-normalised: 956 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	Heterogeneous Federated Learning: State-of-the-art and Research Challenges	Hong Kong Baptist University, Nanyang Technological University, Wuhan University	China, Singapore	Influential
2	FedALA: Adaptive Local Aggregation for Personalized Federated Learning	Louisiana State University, Queen's University Belfast, Shanghai Jiao Tong University	China, United Kingdom, United States	—
3	End-Edge-Cloud Collaborative Computing for Deep Learning: A Comprehensive Survey	Beijing Institute of Technology, University of Chinese Academy of Sciences, University of Oslo	Canada, China, Norway	—
4	Federated Continual Learning via Knowledge Fusion: A Survey	JD iCity, Nanyang Technological University, Southwestern University of Finance and Economics	China, Singapore	—
5	An Aggregation-Free Federated Learning for Tackling Data Heterogeneity	Agency for Science, Technology and Research (A*STAR)	Singapore	—
6	Generalizable Heterogeneous Federated Cross-Correlation and Instance Similarity Learning	Wuhan University	China	—
7	pFedLoRA: Model-Heterogeneous Personalized Federated Learning with LoRA Tuning	—	—	—

No.	Citing paper	Citing institution(s)	Country	S2
8	Deep Model Fusion: A Survey	JD Explore Academy, National University of Defense Technology	China	—
9	A survey on knowledge distillation: Recent advancements	Ryerson University, Seneca Polytechnic	Canada	—
10	Towards Personalized Federated Learning	Hong Kong University of Science and Technology, Nanyang Technological University, Shandong University	China, Hong Kong, Singapore	Influential
11	The Impact of Adversarial Attacks on Federated Learning: A Survey	—	—	—
12	Preserving Privacy in Federated Learning with Ensemble Cross-Domain Knowledge Distillation	Universitas Islam Indonesia, University at Buffalo	Indonesia, United States	Influential
13	Federated learning survey: A multi-level taxonomy of aggregation techniques, experimental insights, and future frontiers	CESI, LabRi-SBA Laboratory	Algeria, France	—
14	Fine-tuning global model via data-free knowledge distillation for non-iid federated learning	JD Explore Academy, Peking University, University of Sydney	Australia, China	Influential
15	Towards building the federatedgpt: Federated instruction tuning	Adobe Research, Amazon, Duke University	United States	—
16	Fedtgp: Trainable global prototypes with adaptive-margin-enhanced contrastive learning for data and model heterogeneity in federated learning	Queen's University Belfast, Shanghai Jiao Tong University, Tsinghua University	China, United Kingdom	Influential
17	Feddisco: Federated learning with discrepancy-aware collaboration	Carnegie Mellon University, Shanghai Jiao Tong University	China, United States	—
18	Personalized federated learning with feature alignment and classifier collaboration	Tsinghua Shenzhen International Graduate School, Tsinghua University	China	—
19	Feddat: An approach for foundation model finetuning in multi-modal heterogeneous federated learning	Ludwig Maximilian University of Munich, Siemens Technology, University of Oxford	Germany, United Kingdom	—
20	Generalized federated learning via sharpness aware minimization	Mississippi State University, University of South Florida	United States	—
21	Fedgh: Heterogeneous federated learning with generalized global header	Nankai University, Nanyang Technological University, University of Science and Technology of China	China, Singapore	—
22	Knowledge distillation and dataset distillation of large language models: Emerging trends, challenges, and future directions	Augusta University, Carnegie Mellon University, Harvard University	United States	—
23	Label-efficient self-supervised federated learning for tackling data heterogeneity in medical imaging	Stanford University, The University of Hong Kong, University of California, Santa Cruz	Hong Kong, United States	—

No.	Citing paper	Citing institution(s)	Country	S2
24	Applications of generative AI (GAI) for mobile and wireless networking: A survey	Kyonggi University, MEPCO SCHLENK Engineering College, Pusan National University	India, Ireland, South Korea	—
25	Privacy-preserving federated learning and uncertainty quantification in medical imaging	Moffitt Cancer Center, Rowan University, University of South Florida	United States	—
26	Personalized subgraph federated learning	KAIST, University of North Carolina at Chapel Hill	South Korea, United States	—
27	Towards efficient replay in federated incremental learning	Ant Group, Huazhong University of Science and Technology	China	—
28	A comprehensive survey of federated transfer learning: challenges, methods and applications	Beihang University, Beijing Academy of Blockchain and Edge Computing, Shandong University	China	—
29	A data-free approach to mitigate catastrophic forgetting in federated class incremental learning for vision tasks	University of Southern California	United States	—
30	Personalized federated learning via feature distribution adaptation	Northeastern University	United States	—

Showing the 30 most-cited of 307 independent citing papers.

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

FOLLOW-UP WORK

[Federated robustness propagation: Sharing adversarial robustness in heterogeneous federated learning](#)

2023 · Proceedings of the AAAI Conference on Artificial Intelligence 37 (7), 7893-7901, 2023 · 81 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Towards communication-efficient adversarial federated learning for robust edge intelligence	Kyung Hee University, University of Houston	South Korea, United States	Influential
2	A systematic literature review of robust federated learning: Issues, solutions, and future research directions	Deakin University, Qilu University of Technology	Australia, China	—
3	Threats and defenses in the federated learning life cycle: A comprehensive survey and challenges	Nantong University, The University of Sydney, University of St Andrews	Australia, China, United Kingdom	—
4	Logit calibration and feature contrast for robust federated learning on non-iid data	Kyung Hee University	South Korea	Influential
5	Federated hybrid training and self-adversarial distillation: Towards robust edge networks	Kyung Hee University, University of Houston	South Korea, United States	—
6	Distributionally adaptive meta reinforcement learning	Massachusetts Institute of Technology, UC Berkeley,	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
		University of California, Berkeley		
7	Making batch normalization great in federated deep learning	The Ohio State University	United States	—
8	Characterizing internal evasion attacks in federated learning	Carnegie Mellon University	United States	—
9	Sylva: Tailoring Personalized Adversarial Defense in Pre-trained Models via Collaborative Fine-tuning	Beijing Institute of Technology, Sun Yat-sen University, Xi'an Jiaotong University	China	—
10	Combating exacerbated heterogeneity for robust models in federated learning	Hong Kong Baptist University, Shanghai Jiao Tong University, The University of Sydney	Australia, China, Hong Kong	—
11	On the security & privacy in federated learning	Ikerlan Technology Research Centre, Radboud University	Netherlands, Spain	—
12	Towards robust federated learning via logits calibration on non-iid data	Kyung Hee University	South Korea	—
13	Improving machine learning robustness via adversarial training	University of South Florida	United States	—
14	Robust and privacy-preserving collaborative learning: A comprehensive survey	Chongqing University, Nanyang Technological University, Zhejiang Lab	China, Singapore	—
15	Certified federated adversarial training	IBM Research Europe, Imperial College London	Switzerland, United Kingdom	—
16	AD-FL: adversarial defense in federated learning via attention denoising	COSMOPlat Institute of Industrial Intelligence (Qingdao) Co., Ltd., Nanjing University of Aeronautics and Astronautics, Qingdao Penghai Software Co., Ltd.	China	—
17	FedProphet: Memory-Efficient Federated Adversarial Training via Robust and Consistent Cascade Learning	Accenture, Duke University	United States	Influential
18	Lorica: A Synergistic Fine-Tuning Framework for Advancing Personalized Adversarial Robustness	Beijing Institute of Technology, Sun Yat-sen University, Xi'an Jiaotong University	China	—

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

FOLLOW-UP WORK

[Efficient Split-Mix Federated Learning for On-Demand and In-Situ Customization](#)

2022 · International Conference on Learning Representations, 2022 · 89 citations (GS)

Field-normalised: 76 Semantic Scholar citations place it in the top 5% of Computer Science papers from 2022 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Heterogeneous Federated Learning: State-of-the-art and Research Challenges	Hong Kong Baptist University, Nanyang Technological University, Wuhan University	China, Singapore	—
2	Personalized federated learning under mixture of distributions	NEC Laboratories America, UCLA, University of California, Los Angeles	United States	Influential
3	Dfrd: Data-free robustness distillation for heterogeneous federated learning	East China Normal University, Peking University	China	—
4	Efficient personalized federated learning via sparse model-adaptation	Alibaba Group	China	—
5	A survey of what to share in federated learning: Perspectives on model utility, privacy leakage, and communication efficiency	Hong Kong University of Science and Technology	Hong Kong	—
6	Fedhm: Efficient federated learning for heterogeneous models via low-rank factorization	Huazhong University of Science and Technology, Huazhong University of Science & Technology, Lehigh University	China, United States	—
7	Dfdg: data-free dual-generator adversarial distillation for one-shot federated learning	East China Normal University, Naval Medical University, Shanghai Normal University	China	—
8	FedMHO: Heterogeneous one-shot federated learning towards resource-constrained edge devices	Huazhong University of Science and Technology, Mohamed bin Zayed University of Artificial Intelligence	China, United Arab Emirates	—
9	Exploring the practicality of federated learning: A survey towards the communication perspective	Technische Universität Berlin, University of Notre Dame, VinUniversity	Germany, United States, Vietnam	—
10	DarkDistill: Difficulty-Aligned Federated Early-Exit Network Training on Heterogeneous Devices	Beihang University, City University of Hong Kong	China	—
11	Adapterfl: Adaptive heterogeneous federated learning for resource-constrained mobile computing systems	East China Normal University, Nanyang Technological University	China, Singapore	—
12	Nefl: Nested model scaling for federated learning with system heterogeneous clients	Korea Advanced Institute of Science and Technology, Samsung Electronics, University of Texas at Austin	South Korea, United States	—
13	FedMHO: Heterogeneous One-Shot Federated Learning Towards Resource-Constrained Clients	Huazhong University of Science and Technology, Mohamed bin Zayed University of Artificial Intelligence	China, United Arab Emirates	—
14	Memory-adaptive Depth-wise Heterogeneous Federated Learning	Carnegie Mellon University, Lehigh University, Samsung Research America	United States	—
15	DynFed: Adaptive Federated Learning via Quantization-Aware Knowledge Distillation	Beijing Institute of Technology, Tsinghua University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
16	GNN at the edge: Cost-efficient graph neural network processing over distributed edge servers	Sun Yat-sen University	China	—
17	Recurrent early exits for federated learning with heterogeneous clients	Brave Software, Flower Labs, Samsung	United Kingdom	—
18	Federated learning of large models at the edge via principal sub-model training	Inha University, Intel Labs, University of Southern California	China, South Korea, United States	Influential
19	Personalised federated learning on heterogeneous feature spaces	Criteo	France	—
20	When computing power network meets distributed machine learning: An efficient federated split learning framework	Nankai University, Tianjin University, University of Oregon	China, United States	—
21	Nebula: An edge-cloud collaborative learning framework for dynamic edge environments	Huawei, Shanghai Jiao Tong University	China	—
22	CC-FedAvg: Computationally customized federated averaging	Harbin Institute of Technology, Harbin Institute of Technology (Shenzhen)	China	—
23	ECLM: Efficient edge-cloud collaborative learning with continuous environment adaptation	Huawei, Shanghai Jiao Tong University	China	—
24	ASFL: An Adaptive Model Splitting and Resource Allocation Framework for Split Federated Learning	Southern University of Science and Technology, The University of British Columbia	Canada, China	—

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

Contribution 2

Claim — Contribution 2

The researcher pioneered privacy-preserving prompt engineering for large language models, establishing a foundational framework that subsequent independent studies have widely adopted to assess and mitigate generative privacy risks.

The researcher’s core contribution centers on the 2023 paper ‘DP-OPT: Make Large Language Model Your Privacy-Preserving Prompt Engineer,’ which appears to introduce a novel approach to integrating differential privacy into the prompt engineering process for large language models. This work serves as the foundational pillar for a broader line of inquiry into the intersection of generative AI and data security.

This line of work addresses the emerging challenge of protecting user data privacy when interacting with powerful generative models. The progression from the core paper to follow-up studies such as ‘LLM-PBE’ and ‘Shake to Leak’ suggests a systematic expansion of this framework. The researcher appears to have moved from proposing a specific privacy-preserving mechanism to broadly assessing data privacy vulnerabilities in large language models and diffusion models, indicating a comprehensive effort to map and mitigate privacy risks across different generative architectures.

The significance of this contribution is evidenced by its substantial uptake in the academic community. The core paper has accumulated 97 citations, while the follow-up work ‘LLM-PBE’ has garnered 125 citations, suggesting growing interest in the field. Notably, 95.9% of the 950 citing papers classified for this scholar originate from independent researchers, indicating that

this line of work has resonated beyond the researcher’s immediate circle and has become a recognized reference point for independent scholars investigating privacy in generative AI.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 117 · 5 flagged influential by Semantic Scholar

CORE PAPER

[DP-OPT: Make Large Language Model Your Privacy-Preserving Prompt Engineer](#)

2023 · ICLR 2024 (Spotlight), 2023 · 97 citations (GS)

Field-normalised: 69 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	Privacy Preserving Prompt Engineering: A Survey	University of Arkansas Fayetteville	United States	—
2	Graph Neural Prompting with Large Language Models	Amazon, University of Notre Dame	United States	—
3	How to DP-fy Your Data: A Practical Guide to Generating Synthetic Data With Differential Privacy	Google DeepMind, Google Research, NYU	Chile, United Kingdom, United States	—
4	NOIR: Privacy-Preserving Generation of Code with Open-Source LLMs	Hamad Bin Khalifa University, Kent State University, New Jersey Institute of Technology	Qatar, United States	—
5	Complex QA and language models hybrid architectures, Survey	Aix Marseille Univ, Aix Marseille Univ, CNRS, Naval Group	France	—
6	A survey on responsible LLMs: inherent risk, malicious use, and mitigation strategy	Huazhong University of Science and Technology, Tsinghua University	China	—
7	Beyond data privacy: New privacy risks for large language models	Alibaba, Purdue University	China, United States	—
8	Bridging today and the future of humanity: Ai safety in 2024 and beyond	University of California, Irvine	United States	—
9	Generative ai for self-adaptive systems: State of the art and research roadmap	KU Leuven, Peking University, Southwest University	Belgium, China, Japan	—
10	Open llms are necessary for current private adaptations and outperform their closed alternatives	CISPA Helmholtz Center for Information Security	Germany	—
11	Privacy-preserving retrieval-augmented generation with differential privacy	University of California, San Diego	United States	—
12	Grounding foundation models through federated transfer learning: A general framework	Huazhong University of Science and Technology, The Hong Kong University of Science and Technology, WeBank	China	—
13	Private prediction for large-scale synthetic text generation	Google	United States	—
14	Privacy auditing of large language models	Google DeepMind, Princeton University	United Kingdom, United States	—
15	PAPILLON: Privacy preservation from Internet-based and local language model ensembles	Columbia University, Stanford University	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
16	Large language models as decision aids in neuro-oncology: a review of shared decision-making applications	Jena University Hospital-Friedrich Schiller University Jena, University of California, Riverside, University of Florida	Germany, United States	—
17	No free lunch theorem for privacy-preserving llm inference	Huazhong University of Science and Technology, WeBank	China	—
18	Casper: Prompt sanitization for protecting user privacy in web-based large language models	New Jersey Institute of Technology, Pyte	United States	—
19	Remoterag: A privacy-preserving llm cloud rag service	The Chinese University of Hong Kong, University of Science and Technology of China	China	—
20	Papillon: Privacy preservation from internet-based and local language model ensembles	Columbia University, Stanford University	United States	—
21	Prmpt: Sanitizing Sensitive Prompts for LLMs	Langroid Incorporated, University of California, San Diego, University of Michigan	Canada, United States	—
22	Protecting users from themselves: Safeguarding contextual privacy in interactions with conversational agents	IBM Research	Japan, United States	—
23	SafeSynthDP: leveraging large language models for privacy-preserving synthetic data generation using differential privacy	University of Alberta	Canada	—
24	Federated domain-specific knowledge transfer on large language models using synthetic data	Peking University, Shenzhen International Graduate School, Tsinghua University, South China University of Technology	China	—
25	LLM access shield: domain-specific LLM framework for privacy policy compliance	Hong Kong Applied Science and Technology Research Institute	China	—
26	Towards harnessing the collaborative power of large and small models for domain tasks	AsiaInfo Technologies, Institute of Computing Technology, Chinese Academy of Sciences, The Hong Kong Polytechnic University	China	—
27	GradOT: Training-free Gradient-preserving Off-site-tuning for Large Language Models	Ant Group, Cleveland Clinic Lerner Research Institute, Soochow University	China, United Kingdom, United States	—
28	The tug of war within: Mitigating the fairness-privacy conflicts in large language models	Renmin University of China, Shanghai Artificial Intelligence Laboratory	China	—
29	Clustering and median aggregation improve differentially private inference	Google Research, University of Southern California	United States	—
30	Deprompt: Desensitization and evaluation of personal identifiable information in large language model prompts	Hainan University, Xidian University	China	—

Showing the 30 most-cited of 46 independent citing papers.

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* – ones that substantively build on the work (S2’s isInfluential signal, Valenzuela et al. 2015) – the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

FOLLOW-UP WORK

LLM-PBE: Assessing Data Privacy in Large Language Models

2024 · VLDB (Best Paper Nomination), 2024 · 125 citations (GS)

Field-normalised: 59 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	Understanding World or Predicting Future? A Comprehensive Survey of World Models	Tsinghua University	China	—
2	A Survey: Towards Privacy and Security in Mobile Large Language Models	Georgia State University, Kennesaw State University, Nexa AI	China, United States	—
3	A survey on privacy risks and protection in large language models	Nanyang Technological University, New York University, The Hong Kong University of Science and Technology (Guangzhou)	China, Singapore, United States	—
4	A comprehensive survey and guide to multi-modal large language models in vision-language tasks	AppCubic, Indiana University, JTB Technology Corp.	United Kingdom, United States	—
5	Privaci-bench: Evaluating privacy with contextual integrity and legal compliance	Hong Kong University of Science and Technology, Huawei Technologies, National University of Singapore	Hong Kong, Singapore, United Kingdom	—
6	The scales of justitia: A comprehensive survey on safety evaluation of llms	Beihang University, Beijing University of Posts and Telecommunications, China Academy of Information and Communications Technology	China, United States	—
7	When {LLMs} go online: The emerging threat of {Web-Enabled}{LLMs}	KAIST	South Korea	—
8	A survey on responsible LLMs: inherent risk, malicious use, and mitigation strategy	Huazhong University of Science and Technology, Tsinghua University	China	—
9	Toward holistic evaluation of recommender systems powered by generative models	Amazon Web Services, Autónoma University of Madrid, Duke University	Italy, Spain, United States	—
10	Safelawbench: Towards safe alignment of large language models	Hong Kong University of Science and Technology, Peking University	China, Hong Kong	—
11	Big Help or Big Brother? Auditing Tracking, Profiling, and Personalization in Generative {AI} Assistants	UC Davis, UCL, Universidad Carlos III de Madrid	Italy, Spain, United Kingdom	—
12	A survey on model extraction attacks and defenses for large language models	Duke University, Florida State University, Northwestern University	United States	—

No.	Citing paper	Citing institution(s)	Country	S2
13	Toward a human-centered evaluation framework for trustworthy llm-powered gui agents	Johns Hopkins University, Northeastern University, University of Notre Dame	United States	—
14	Tokens for learning, tokens for unlearning: Mitigating membership inference attacks in large language models via dual-purpose training	Emory University	United States	—
15	Less or more: Towards glanceable explanations for LLM recommendations using ultra-small devices	Meta Reality Labs, Purdue University, University of California San Diego	United States	—
16	Towards more realistic extraction attacks: An adversarial perspective	McGill University	Canada	—
17	Generalist multimodal ai: A review of architectures, challenges and opportunities	Pacific Northwest National Laboratory	United States	—
18	Pig: Privacy jailbreak attack on llms via gradient-based iterative in-context optimization	Chinese Academy of Sciences, Guangzhou University	China	—
19	Membership inference attacks as privacy tools: Reliability, disparity and ensemble	IBM Research, Rensselaer Polytechnic Institute	Japan, United States	—
20	Mind the third eye! benchmarking privacy awareness in mllm-powered smartphone agents	Columbia University, Hong Kong University of Science and Technology (Guangzhou), Shandong University	China, United States	—
21	Privacyscalpel: Enhancing llm privacy via interpretable feature intervention with sparse autoencoders	Huawei	Germany	—
22	Multi-PA: A Multi-perspective Benchmark on Privacy Assessment for Large Vision-Language Models	Chinese Academy of Sciences	China	—
23	Ip leakage attacks targeting llm-based multi-agent systems	Hong Kong University of Science and Technology, The Hong Kong University of Science and Technology	China, Hong Kong	—
24	Small language models: Architectures, techniques, evaluation, problems and future adaptation	American International University-Bangladesh, Cornell University	Bangladesh, United States	—
25	Beyond data privacy: New privacy risks for large language models	Alibaba, Purdue University	China, United States	—
26	Can textual gradient work in federated learning?	Nanyang Technological University, The University of British Columbia, University of Pennsylvania	Canada, Singapore, United States	—
27	ACAI for SBOs: AI Co-creation for Advertising and Inspiration for Small Business Owners	Google DeepMind, University of Oxford	United Kingdom	—
28	From reviews to dialogues: Active synthesis for zero-shot llm-based conversational recommender system	Cornell University, Netflix Inc., UC San Diego	United States	—
29	HarmMetric Eval: Benchmarking Metrics and Judges for LLM Harmfulness Assessment	Zhejiang University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
30	SoK: Semantic Privacy in Large Language Models	CSIRO Data61, Northeastern University, University of Technology Sydney	Australia, China, United States	—

Showing the 30 most-cited of 55 independent citing papers.

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

FOLLOW-UP WORK

[Shake to Leak: Fine-tuning Diffusion Models Can Amplify the Generative Privacy Risk](#)

2024 · 2nd IEEE Conference on Secure and Trustworthy Machine Learning, 2024 · 25 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Bridging today and the future of humanity: Ai safety in 2024 and beyond	University of California, Irvine	United States	—
2	Replication in visual diffusion models: A survey and outlook	Baidu Inc., University of Technology Sydney, Zhejiang University	Australia, China	—
3	SoK: Blockchain-Based Decentralized AI (DeAI)	Ethereum Foundation, Flock.io, Imperial College London	Switzerland, United Kingdom	Influential
4	Synthetic face datasets generation via latent space exploration from brownian identity diffusion	École Polytechnique Fédérale de Lausanne, Idiap Research Institute	Switzerland	Influential
5	Exactly minimax-optimal locally differentially private sampling	KAIST, McMaster University	Canada, South Korea	—
6	Trustworthy text-to-image diffusion models: A timely and focused survey	Chalmers University of Technology, Lancaster University, The University of Warwick	Sweden, United Kingdom	—
7	A common pool of privacy problems: Legal and technical lessons from a large-scale web-scraped machine learning dataset	Carnegie Mellon University, Georgetown University, University of Washington	United States	Influential
8	A survey on privacy attacks against digital twin systems in AI-robotics	Mississippi State University	United States	—
9	Unveiling synthetic faces: How synthetic datasets can expose real identities	Idiap Research Institute	Switzerland	—
10	Differentially private kernel density estimation	Ensemble AI, Northwestern University, UC Berkeley	United States	—
11	Dual-model defense: Safeguarding diffusion models from membership inference attacks through disjoint data splitting	VinAI Research	—	—
12	LoyalDiffusion: A diffusion model guarding against data replication	University of Southern California	United States	—
13	Noise as a Probe: Membership Inference Attacks on Diffusion Models Leveraging Initial Noise	Nanjing University of Posts and Telecommunica-	Australia, China	—

No.	Citing paper	Citing institution(s)	Country	S2
		tions, Southeast University, The University of Queensland		
14	AugGen: Synthetic Augmentation using Diffusion Models Can Improve Recognition	EPFL, Idiap Research Institute, UNIL	Switzerland	—
15	Evaluation without Generation: Non-Generative Assessment of Harmful Model Specialization with Applications to CSAM	Boston University, MIT, Thorn	United States	—
16	Guided Safe Diffusion: Prohibiting Diffusion Models from Generating Inappropriate Content	Kingston and St George's University, Sheffield Emergency Care Forum, University of Bath	United Kingdom	—

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

Contribution 3

Claim – Contribution 3

The researcher advanced resilient and communication-efficient federated learning, subsequently extending this framework to address robustness in knowledge distillation and secure model watermarking.

The researcher’s core contribution centers on the 2022 paper ‘Resilient and communication efficient learning for heterogeneous federated systems,’ which established a foundation for optimizing distributed machine learning in diverse environments. This work appears to address the dual challenges of maintaining model performance and minimizing communication overhead in heterogeneous federated settings.

Building on this foundation, the researcher’s subsequent publications suggest a strategic expansion into security and robustness. The 2023 follow-up, ‘Revisiting Data-Free Knowledge Distillation with Poisoned Teachers,’ indicates an application of these principles to mitigate risks in data-free distillation scenarios. Another 2023 paper, ‘Safe and Robust Watermark Injection with a Single OoD Image,’ further extends this line of inquiry toward secure model identification, suggesting a cohesive research trajectory focused on trustworthy and efficient AI systems.

The significance of this work is evidenced by its substantial uptake in the academic community. The core paper has garnered 61 citations, while the follow-up studies have accumulated 21 and 5 citations respectively. Notably, analysis of 950 citing papers reveals that 95.9% originate from independent researchers, indicating that this line of work has achieved broad recognition and influence beyond the researcher’s immediate institutional circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 31

CORE PAPER

[Resilient and communication efficient learning for heterogeneous federated systems](#)

2022 · Proceedings of Thirty-ninth International Conference on Machine Learning ..., 2022 · 61 citations (GS)

Field-normalised: 46 Semantic Scholar citations place it in the top 10% of Computer Science papers from 2022 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	pFedLoRA: Model-Heterogeneous Personalized Federated Learning with LoRA Tuning	—	—	—

No.	Citing paper	Citing institution(s)	Country	S2
2	Fedgh: Heterogeneous federated learning with generalized global header	Nankai University, Nanyang Technological University, University of Science and Technology of China	China, Singapore	—
3	Federated model heterogeneous matryoshka representation learning	Nankai University, Nanyang Technological University, The University of British Columbia	Canada, China, Singapore	—
4	Knowledge distillation in federated learning: A survey on long lasting challenges and new solutions	City University of Macau, University of Illinois at Chicago, University of Technology Sydney	Australia, China, United States	—
5	pfedes: Model heterogeneous personalized federated learning with feature extractor sharing	Nankai University, Nanyang Technological University	China, Singapore	—
6	Fedcache 2.0: Federated edge learning with knowledge caching and dataset distillation	Beihang University, Beijing Jiaotong University, Chinese Academy of Sciences	China	—
7	Data-Free Continual Learning of Server Models in Model-Heterogeneous Cloud-Device Collaboration	Beihang University, McGill University, Shandong University	Canada, China	—
8	When foresight pruning meets zeroth-order optimization: Efficient federated learning for low-memory devices	East China Normal University	China	—
9	Bridging today and the future of humanity: AI safety in 2024 and beyond	University of California, Irvine	United States	—
10	Hetefedrec: Federated recommender systems with model heterogeneity	Beihang University, Shandong University, The Hong Kong University of Science and Technology	Australia, China	—
11	Fedmoe: Personalized federated learning via heterogeneous mixture of experts	Beijing University of Posts and Telecommunications, University of Cambridge	China, United Kingdom	—
12	Hide your model: A parameter transmission-free federated recommender system	Deakin University, Griffith University, The University of Queensland	Australia	—
13	Towards resilient federated learning in cyberedge networks: Recent advances and future trends	Real-Time and Embedded Computing Systems Research Centre, TU Berlin, University of Cambridge	Germany, Portugal, United Kingdom	—
14	Robust model aggregation for heterogeneous federated learning: Analysis and optimizations	CSIRO, Institute of High Performance Computing, A*STAR, Nanjing University of Science and Technology	Australia, China, Singapore	—
15	Da-pfl: Dynamic affinity aggregation for personalized federated learning	Dongguan University of Technology, Harbin Institute of Technology (Shenzhen), Harbin Institute of Technology (Shenzhen); National University of Defense Technology	China	—

No.	Citing paper	Citing institution(s)	Country	S2
16	Toward Enhancing Representation Learning in Federated Multi-Task Settings	Huawei	Canada	—
17	Operator-Theoretic Framework for Gradient-Free Federated Learning	Primetals Technologies Austria GmbH, Software Competence Center Hagenberg GmbH, University of Rostock	Austria, Germany	—
18	Heterogeneous Federated Learning with Prototype Alignment and Upscaling	SAKAK Inc., The Cyber University of Korea	South Korea	—
19	pFedAFM: Adaptive Feature Mixture for Batch-Level Personalization in Heterogeneous Federated Learning	Nankai University, Nanyang Technological University, The University of British Columbia	Canada, China, Singapore	—

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

FOLLOW-UP WORK

[Revisiting Data-Free Knowledge Distillation with Poisoned Teachers](#)

2023 · International Conference on Machine Learning, 2023 · 21 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Hydra-fl: Hybrid knowledge distillation for robust and accurate federated learning	University of Massachusetts, Amherst	United States	—
2	Unlearning backdoor attacks for llms with weak-to-strong knowledge distillation	MIT, Nanyang Technological University, Northwest Normal University	China, Singapore, United States	—
3	Evading data provenance in deep neural networks	Carnegie Mellon University, Shanghai Jiao Tong University, Southeast University	China, United States	—
4	Teacher as a lenient expert: Teacher-agnostic data-free knowledge distillation	Inha University	South Korea	—
5	Fusing pruned and backdoored models: Optimal transport-based data-free backdoor mitigation	The Chinese University of Hong Kong, Shenzhen, The Hong Kong University of Science and Technology (Guangzhou)	China	—
6	How to backdoor the knowledge distillation	Carnegie Mellon University Africa, The Pennsylvania State University	Rwanda, United States	—
7	Unlocking tuning-free few-shot adaptability in visual foundation models by recycling pre-tuned loras	Nanyang Technological University, Sun Yat-sen University, Tsinghua University	China, Singapore	—
8	Revisiting the Auxiliary Data in Backdoor Purification	The Chinese University of Hong Kong	China	—
9	Exploited Intelligence: AI, Cybercrime, and the Global South's Digital Struggle	German Institute for Global and Area Studies (GIGA)	Philippines	—

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

FOLLOW-UP WORK

Safe and Robust Watermark Injection with a Single OoD Image

2023 · ICLR 2024, 2023 · 5 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Bridging today and the future of humanity: Ai safety in 2024 and beyond	University of California, Irvine	United States	—
2	Watermarking techniques for large language models: A survey	Jinan University, University of Illinois Chicago	China, United States	—
3	Reliable model watermarking: Defending against theft without compromising on evasion	Shanghai Jiao Tong University, Southeast University	China	—

Independent citing papers only; self- and co-author citations excluded. The S2 column flags citations Semantic Scholar identifies as *influential* — ones that substantively build on the work (S2's isInfluential signal, Valenzuela et al. 2015) — the “built on / relied upon” pattern the AAO credits. Counsel should quote the citing text for the strongest of these.

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Nanyang Technological University	Singapore	SCImago #137	51
Zhejiang University	China	SCImago #6 · THE 39 · QS 49	48
Tsinghua University	China	SCImago #8 · THE 12 · QS =17	42
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	39
University of Science and Technology of China	China	SCImago #77 · THE 51 · QS =132	26
Beihang University	China	SCImago #160 · THE 251–300 · QS =388	24
Michigan State University	United States	SCImago #436 · THE =105 · QS 161	22
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	22
Peking University	China	SCImago #11 · THE 13 · QS 14	22
Huazhong University of Science and Technology	China	SCImago #25 · THE =176 · QS 319	22
Sun Yat-sen University	China	SCImago #40 · THE 201–250 · QS =276	21
The Hong Kong University of Science and Technology	Hong Kong	SCImago #483 · THE =58 · QS 44	20
University of Southern California	United States	SCImago #192 · THE =73 · QS 146	20
Northeastern University	United States	QS 384	19
Chinese Academy of Sciences	China	SCImago #2	18

Geographic distribution of citing authors

Country	Citing papers
China	456
United States	397
Singapore	84
United Kingdom	68
Australia	63
Hong Kong	42
South Korea	41
Germany	39
Canada	37
Japan	27
Switzerland	23
France	15

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	Data-free knowledge distillation for heterogeneous federated learning	349	Dhanasar – Prong 2 (well-positioned)
Contribution 2	DP-OPT: Make Large Language Model Your Privacy-Preserving Prompt Engineer	117	Dhanasar – Prong 2 (well-positioned)
Contribution 3	Resilient and communication efficient learning for heterogeneous federated systems	31	Dhanasar – Prong 2 (well-positioned)