

Citation Evidence Report

EB-1A Petition — Original Contributions of Major Significance

8 CFR § 204.5(h)(3)(v) · Criterion 5


Yann LeCun

Chief AI Scientist at Facebook & JT Schwarz Professor at the Courant Institute, New York University

[Google Scholar profile](#)

Generated 2026-05-31 by CiteMap. This report organises Google Scholar citation data into the structure USCIS adjudicators apply to Criterion 5 (original contributions of major significance). 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.

94.5% independent of 4,946 classified citing papers

Citation type	Count
Independent	4,674
Self-citation	28
Co-author	237
Same-institution	7

2,821 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 gradient-based learning for pattern recognition, establishing foundational methods for handwritten digit and document classification that evolved into character-level text analysis.

CLAIM: The researcher’s seminal contribution lies in applying backpropagation to handwritten zip code recognition, a core paper published in 1989 that initiated a sustained line of inquiry into gradient-based learning for document and text recognition.

ORIGINALITY: This work appears to address the challenge of applying neural network training methods to practical recognition tasks. The progression from zip code recognition in 1989 to broader document recognition in 2002 and character-level text classification in 2015 suggests a methodological evolution, extending early backpropagation applications to more complex, high-dimensional data structures over time.

SIGNIFICANCE: The impact of this research is evidenced by substantial citation counts, with the core paper cited over 20,000 times and the 2002 follow-up cited over 86,000 times. Furthermore, analysis of nearly 5,000 citing papers reveals that 98.1% originate from independent researchers, indicating widespread adoption and influence across the broader scientific community rather than self-citation or institutional clustering.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 1,695 · 168 flagged influential by Semantic Scholar

CORE PAPER

[Backpropagation Applied to Handwritten Zip Code Recognition](#)

1989 · Neural Computation · 20,530 citations (GS)

Field-normalised: 12,010 Semantic Scholar citations place it in the top 1% of Computer Science papers from 1989 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Designing Network Design Spaces (2020)	Facebook AI Research	—	—
2	EVA: Exploring the Limits of Masked Visual Representation Learning at Scale (2023)	Beijing Academy of Artificial Intelligence, Beijing Institute of Technology, Huazhong University of Science and Technology	China	—
3	InternImage: Exploring Large-Scale Vision Foundation Models with Deformable Convolutions (2023)	Nanjing University, SenseTime, SenseTime Research	China, Hong Kong	Background
4	Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour (2017)	Facebook AI Research	United States	Background
5	Rethinking Atrous Convolution for Semantic Image Segmentation (2017)	Google Inc.	United States	Methodology
6	Activation functions: Comparison of trends in practice and research for deep learning (2018)	University of Strathclyde	United Kingdom	Methodology
7	Mastering Diverse Domains through World Models (2023)	Google DeepMind, University of Toronto	Canada, United Kingdom, United States	Background
8	A Survey of Deep Active Learning (2021)	Carnegie Mellon University, Monash University, National	Australia, China, India	Methodology

No.	Citing paper	Citing institution(s)	Country	S2
		Institute of Technology Kurukshetra		
9	Transformers in Vision: A Survey (2022)	Inception Institute of Artificial Intelligence, MBZ University of Artificial Intelligence, Monash University	Australia, United Arab Emirates, United States	Methodology
10	Natural Language Understanding and Inference with MLLM in Visual Question Answering: A Survey (2025)	Sun Yat-sen University, Tsinghua University, Worcester Polytechnic Institute	China, United States	—
11	Visual Attention Network (2023)	Fitten Tech, Nankai University, Tsinghua University	China	—
12	Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review (2017)	University of South Africa	South Africa	Methodology
13	A Survey of the Usages of Deep Learning for Natural Language Processing (2020)	University of Colorado Colorado Springs	United States	Background
14	Compute Trends Across Three Eras of Machine Learning (2022)	Complutense University of Madrid, Epoch, Epoch AI	Germany, Spain, United States	—
15	A Survey of Convolutional Neural Networks: Analysis, Applications, and Prospects (2022)	Griffith University, Hohai University, The Hong Kong University of Science and Technology	Australia, China, Hong Kong	Methodology
16	Introduction to Machine Learning, Neural Networks, and Deep Learning (2020)	Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, OHSU	United States	—
17	Deep learning in computational mechanics: a review (2024)	Bauhaus-Universität Weimar	Germany	—
18	Deep Learning for Brain MRI Segmentation: State of the Art and Future Directions (2017)	Mayo Clinic, Stanford University School of Medicine	United States	Background
19	A Survey of the Recent Architectures of Deep Convolutional Neural Networks (2020)	Pakistan Institute of Engineering and Applied Sciences	Pakistan	—
20	Artificial intelligence in the creative industries: a review (2021)	University of Bristol	United Kingdom	Background
21	Artificial intelligence to deep learning: machine intelligence approach for drug discovery (2021)	Delhi Technological University, Delhi Technological University (Formerly DCE)	India	Background
22	Survey on deep learning with class imbalance (2019)	Florida Atlantic University	United States	—
23	Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning (2017)	Google Inc.	—	—
24	Deep Learning for Computer Vision: A Brief Review (2018)	National Technical University of Athens, Technological Educational Institute of Athens	Greece	—

No.	Citing paper	Citing institution(s)	Country	S2
25	MLP-Mixer: An all-MLP Architecture for Vision (2021)	Google Research	Switzerland	Background
26	Diffusion for World Modeling: Visual Details Matter in Atari (2024)	Microsoft Research, University of Edinburgh, University of Geneva	Switzerland, United Kingdom, United States	—
27	Building machines that learn and think like people (2017)	Harvard University, Massachusetts Institute of Technology, New York University	United States	Background
28	A Comprehensive Review of Deep Learning: Architectures, Recent Advances, and Applications (2024)	University of Johannesburg	South Africa	—
29	Deep learning in optical metrology: a review (2022)	Nanjing University of Science and Technology, Nanyang Technological University, Queen Mary University of London	China, Singapore, United Kingdom	Methodology
30	Neuromorphic computing at scale (2025)	Google DeepMind, Indian Institute of Science, Intel Labs	China, Germany, India	—

Showing the 30 most-cited of 776 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).

Citing-text excerpts — how the field used this work

METHODOLOGY Rethinking Atrous Convolution for Semantic Image Segmentation

“For the task of semantic segmentation [18, 53, 12, 82], we consider two challenges in applying Deep Convolutional Neural Networks (DCNNs) [42].”

METHODOLOGY Activation functions: Comparison of trends in practice and research for deep learning

“used for recognition tasks had few layers in their entire architecture, with LeNet5, having just five layers [6].”

METHODOLOGY A Survey of Deep Active Learning

“1a, we present a standard deep learning model example: convolutional neural network (CNN) [91, 130].”

METHODOLOGY Transformers in Vision: A Survey

“LeViT [72] (name inspired from LeNet [111]) applies a four-layered CNN block (with 3×3 convolutions) at the beginning with progressively increasing channels (3,32,64,128,256).”

METHODOLOGY Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review

“The purpose of the pooling layers is to reduce the spatial resolution of the feature maps and thus achieve spatial invariance to input distortions and translations (LeCun et al., 1989a, 1989b; LeCun et al., 1998, 2015; Ranzato et al., 2007).”

FOLLOW-UP WORK

[Gradient-based learning applied to document recognition](#)

2002 · Proceedings of the IEEE 86 (11), 2278-2324, 2002 · 86,349 citations (GS)

Field-normalised: 60,115 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2002 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	iCaRL: Incremental Classifier and Representation Learning (2017)	Institute of Science and Technology Austria, IST Austria	Austria	Methodology

No.	Citing paper	Citing institution(s)	Country	S2
2	VideoMAE V2: Scaling Video Masked Autoencoders with Dual Masking (2023)	Nanjing University, Shanghai AI Lab, Shanghai Artificial Intelligence Laboratory	China	Background
3	InternVL: Scaling up Vision Foundation Models and Aligning for Generic Visual-Linguistic Tasks (2024)	Nanjing University, SenseTime, Shanghai AI Laboratory	China, Hong Kong	—
4	Adding Conditional Control to Text-to-Image Diffusion Models (2023)	Stanford University	United States	Background
5	Perceiver: General Perception with Iterative Attention (2021)	DeepMind	United Kingdom	Background
6	Neural Architecture Search with Reinforcement Learning (2017)	Google	United States	Background
7	Improved Regularization of Convolutional Neural Networks with Cutout	University of Guelph	Canada	Methodology
8	Hyper-Parameter Optimization: A Review of Algorithms and Applications (2020)	Inspur, Inspur Electronic Information Industry Co., Ltd	China	Methodology
9	FedBN: Federated Learning on Non-IID Features via Local Batch Normalization (2021)	Monash University, The Chinese University of Hong Kong	Australia, China	Methodology
10	Advancing Transformer Architecture in Long-Context Large Language Models: A Comprehensive Survey (2023)	—	—	Methodology
11	Vision Mamba: Efficient Visual Representation Learning with Bidirectional State Space Model (2024)	Beijing Academy of Artificial Intelligence, Horizon Robotics, Huazhong University of Science and Technology	China	—
12	A Survey on Deep Learning: Algorithms, Techniques, and Applications (2018)	—	—	Background
13	A Survey of Deep Active Learning (2021)	Carnegie Mellon University, Monash University, National Institute of Technology Kurukshetra	Australia, China, India	Background
14	Understanding World or Predicting Future? A Comprehensive Survey of World Models (2025)	Tsinghua University	China	—
15	A Comprehensive Survey of Small Language Models in the Era of Large Language Models: Techniques, Enhancements, Applications, Collaboration with LLMs, and Trustworthiness (2024)	Amazon, Industry Research, LinkedIn	United States	—
16	nnFormer: Volumetric Medical Image Segmentation via a 3D Transformer (2023)	The Chinese University of Hong Kong (Shenzhen), The University of Hong Kong, Xiamen University	China, Hong Kong	Background
17	Vision-Language Models for Vision Tasks: A Survey (2024)	Nanyang Technological University, Shenzhen Institutes	China, Singapore	—

No.	Citing paper	Citing institution(s)	Country	S2
		of Advanced Technology, Chinese Academy of Sciences		
18	Advancements in Generative AI: A Comprehensive Review of GANs, GPT, Autoencoders, Diffusion Model, and Transformers (2024)	Bowie State University, Morgan State University, University of the District of Columbia	United States	Methodology
19	A Comprehensive Review of Convolutional Neural Networks for Defect Detection in Industrial Applications (2024)	University of Huddersfield	United Kingdom	Methodology
20	Visual Attention Network (2023)	Fitten Tech, Nankai University, Tsinghua University	China	—
21	Deep Learning for Health Informatics (2016)	Imperial College London	United Kingdom	Background
22	Towards Evaluating the Robustness of Neural Networks (2016)	University of California, Irvine Medical Center	United States	—
23	Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review (2017)	University of South Africa	South Africa	Methodology
24	Efficient Processing of Deep Neural Networks: A Tutorial and Survey (2017)	Massachusetts Institute of Technology	United States	Methodology
25	A Comprehensive Survey on Graph Neural Networks (2021)	Monash University, University of Illinois Chicago, University of Technology Sydney	Australia, United States	—
26	A Survey of the Usages of Deep Learning for Natural Language Processing (2020)	University of Colorado Colorado Springs	United States	Background
27	Self-Supervised Visual Feature Learning With Deep Neural Networks: A Survey (2020)	City University of New York, The Graduate Center, The City University of New York	United States	Background
28	A Survey on Performance Metrics for Object-Detection Algorithms (2020)	Federal University of Rio de Janeiro	Brazil	—
29	Image Segmentation Using Deep Learning: A Survey (2021)	Australian National University, Snapchat, University of California, Irvine Medical Center	Australia, Canada, Spain	Background
30	Compute Trends Across Three Eras of Machine Learning (2022)	Complutense University of Madrid, Epoch, Epoch AI	Germany, Spain, United States	—

Showing the 30 most-cited of 903 independent citing papers.

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

Citing-text excerpts — how the field used this work

METHODOLOGY iCaRL: Incremental Classifier and Representation Learning

“ons (the current training data). In particular, this is also how iCaRL learns about the existence of new classes. Architecture. Under the hood, iCaRL makes use of a convolutional neural network (CNN) [14]1. We interpret the network as a trainable feature extractor, XIRd, followed by a single classification layer with as many sigmoid output nodes as classes observed so far [3]. All feature vectors ar”

METHODOLOGY Improved Regularization of Convolutional Neural Networks with Cutout

“Much of these improvements can be attributed to the use of convolutional neural networks (CNNs) [9], which are capable of learning complex hierarchical feature representations of images.”

METHODOLOGY Hyper-Parameter Optimization: A Review of Algorithms and Applications

“For documentation recognition (LeCun et al., 1998), several transformations for enhanced model accuracy and robustness are mainly applied.”

METHODOLOGY FedBN: Federated Learning on Non-IID Features via Local Batch Normalization

“Specifically, we use the digits classification task and treat the two unseen datasets – Morpho-global and Morpho-local from Morpho-MNIST (Castro et al., 2019) as the two new clients.”

METHODOLOGY Advancing Transformer Architecture in Long-Context Large Language Models: A Comprehensive Survey

“Inspired by convolutional neural networks (CNNs) [94, 89], another approach is to use sliding-window techniques, as demonstrated in Longformer [11].”

FOLLOW-UP WORK

Character-level Convolutional Networks for Text Classification

2015 · Advances in Neural Information Processing Systems 28 (NIPS 2015) · 9,245 citations (GS)

Field-normalised: 6,940 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2015 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	A Comprehensive Survey on Pretrained Foundation Models: A History from BERT to ChatGPT (2025)	Beihang University, Duke University, Hangzhou Dianzi University	Australia, China, Singapore	—
2	Unleashing the potential of prompt engineering for large language models (2025)	Beijing Normal-Hong Kong Baptist University, Beijing Normal University, BNU-HKBU United International College	China, Singapore, United States	—
3	DepGraph: Towards Any Structural Pruning (2023)	Huawei Technologies Ltd., National University of Singapore, Zhejiang University	China, Singapore	—
4	Continual Learning of Large Language Models: A Comprehensive Survey (2025)	Google DeepMind, Google Inc, Rutgers The State University of New Jersey	—	—
5	Simple and Effective Masked Diffusion Language Models (2024)	Cornell Tech	United States	Methodology
6	Can Large Language Models Be an Alternative to Human Evaluations? (2023)	National Taiwan University	Taiwan	Methodology
7	Sentiment Analysis in the Era of Large Language Models: A Reality Check (2024)	DAMO Academy, Alibaba Group, University of Illinois at Chicago	Singapore, United States	Methodology
8	Finetuned Language Models Are Zero-Shot Learners (2021)	Google Research	United States	—
9	Aya Model: An Instruction Finetuned Open-Access Multilingual Language Model (2024)	Brown University, Carnegie Mellon University, Cohere	Canada, United States	—
10	From Coordinate Matching to Structural Alignment: Rethinking Prototype Alignment in Heterogeneous Federated Learning	Beihang University, University at Buffalo	China, United States	—
11	Lightweight Stylistic Consistency Profiling: Robust Detection of LLM-Generated Textual Content for Multimedia Moderation	Institute of Automation, Chinese Academy of Sciences, Shanghai Jiao Tong University, Zhejiang University	China	—

No.	Citing paper	Citing institution(s)	Country	S2
12	Safety Anchor: Defending Harmful Fine-tuning via Geometric Bottlenecks	Nanjing University of Posts and Telecommunications	China	—
13	Budgeted Attention Allocation: Cost-Conditioned Compute Control for Efficient Transformers	Independent Researcher	United States	—
14	When2Speak: A Dataset for Temporal Participation and Turn-Taking in Multi-Party Conversations for Large Language Models	Duke University	United States	—
15	Revisiting the Role of Label Smoothing in Enhanced Text Sentiment Classification	Shanghai International Studies University	China	Methodology
16	DRIFT: Drift-Resilient Invariant-Feature Transformer for DGA Detection	Sookmyung Women's University	South Korea	—

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

Citing-text excerpts — how the field used this work

METHODOLOGY Simple and Effective Masked Diffusion Language Models

“...and evaluating likelihoods on the validation splits of 7 datasets: Penn Tree Bank (PTB; Marcus et al. [28]), Wikitext [29], One Billion Word Language Model Benchmark (LM1B; Chelba et al. [5]), Lambada [31], AG News [51], and Scientific Papers (Pubmed and Arxiv subsets; Cohan et al. [7]).”

METHODOLOGY Can Large Language Models Be an Alternative to Human Evaluations?

“We use the adversarial samples generated against a bert-base-uncased text classifier trained on AG-News, using three different adversarial attacks: Textfooler, PWWS, and BAE.”

METHODOLOGY Sentiment Analysis in the Era of Large Language Models: A Reality Check

“However, the Yelp-5 dataset offers a more fine-grained sentiment classification by introducing three additional sentiment classes: very positive, very negative, and neutral.”

METHODOLOGY Revisiting the Role of Label Smoothing in Enhanced Text Sentiment Classification

“Moving to the TextCNN architecture, LS models exhibit superior performance on three-category classification datasets compared to binary classification datasets.”

Contribution 2

Claim — Contribution 2

The researcher pioneered foundational convolutional network architectures for multimodal data and subsequently advanced next-generation AI through neuro-inspired computational frameworks.

CLAIM: The researcher established a seminal contribution in deep learning by developing convolutional networks applicable to images, speech, and time series, as evidenced by the highly cited 1998 paper. This work serves as the cornerstone for subsequent research into neuro-inspired artificial intelligence.

ORIGINALITY: The titles suggest a trajectory from establishing fundamental neural network structures for diverse data types to exploring advanced, next-generation AI paradigms. The follow-up papers indicate a shift toward catalyzing a ‘neuroai revolution,’ implying an effort to bridge biological inspiration with computational efficiency in modern AI systems.

SIGNIFICANCE: The core paper has accumulated nearly 10,000 citations, demonstrating widespread adoption. Crucially, 98.1% of citing papers originate from independent researchers, confirming that this work has significantly influenced the broader scientific community beyond the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 752 · 27 flagged influential by Semantic Scholar

CORE PAPER

Convolutional networks for images, speech, and time series

1998 · The handbook of brain theory and neural networks, 1998 · 9,915 citations (GS)

Field-normalised: 6,117 Semantic Scholar citations place it in the top 1% of Computer Science papers from 1998 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	TransNeXt: Robust Foveal Visual Perception for Vision Transformers (2023)	Independent Researcher	—	—
2	U-Mamba: Enhancing Long-range Dependency for Biomedical Image Segmentation (2024)	University Health Network	Canada	—
3	From System 1 to System 2: A Survey of Reasoning Large Language Models (2025)	AiShiWeiLai AI Research, Chinese Academy of Sciences, City University of Hong Kong and the Hong Kong University of Science and Technology (Guangzhou)	China, United Arab Emirates, United Kingdom	—
4	A Survey on Deep Learning: Algorithms, Techniques, and Applications (2018)	—	—	Methodology
5	A Comprehensive Survey on Graph Neural Networks (2021)	Monash University, University of Illinois Chicago, University of Technology Sydney	Australia, United States	Background
6	A Survey of the Usages of Deep Learning for Natural Language Processing (2020)	University of Colorado Colorado Springs	United States	Background
7	Self-normalizing neural networks (2017)	Johannes Kepler University Linz	Austria	—
8	How to Build the Virtual Cell with Artificial Intelligence: Priorities and Opportunities (2024)	Agilent Technologies, Allen Institute for Cell Science, Arc Institute	Canada, Germany, Sweden	—
9	An Introduction to Deep Reinforcement Learning (2018)	Google, McGill University, McGill University and Mila - Quebec AI Institute	Canada	—
10	Recent advances and applications of deep learning methods in materials science (2022)	Carnegie Mellon University, Columbia University, Lawrence Berkeley National Laboratory	United States	Background
11	Deep learning models for plant disease detection and diagnosis (2018)	—	—	—
12	A high-bias, low-variance introduction to Machine Learning for physicists (2019)	Boston University, The Graduate Center, City University of New York, University of California, Irvine Medical Center	United States	—
13	The Evolution of Distributed Systems for Graph Neural Networks and their Origin in Graph Processing and Deep Learning: A Survey (2023)	Technical University of Munich, University of Bayreuth, University of Toronto	Canada, Germany	—

No.	Citing paper	Citing institution(s)	Country	S2
14	A Comprehensive Survey on Community Detection With Deep Learning (2022)	Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Macquarie University, Tianjin University	Australia, China, United States	Background
15	An on-chip photonic deep neural network for image classification (2022)	University of Pennsylvania	—	—
16	Traffic Flow Prediction via Spatial Temporal Graph Neural Network (2020)	—	—	—
17	Deep learning for chest X-ray analysis: A survey (2021)	Radboud University Medical Center	Netherlands	—
18	Review of Deep Learning Algorithms and Architectures (2019)	University of Bridgeport	United States	Background
19	Deep Convolutional Neural Network Based Regression Approach for Estimation of Remaining Useful Life (2016)	Institute for Infocomm Research	Singapore	Methodology
20	Federated learning on non-IID data: A survey (2021)	East China University of Science and Technology, Queen's University Belfast	China, United Kingdom	—
21	A survey of deep neural network architectures and their applications (2017)	Brunel University London, King Abdulaziz University, Xiamen University	China, Saudi Arabia, United Kingdom	—
22	Deep Learning in Mobile and Wireless Networking: A Survey (2019)	Imperial College London, Institute for Computing Systems Architecture (ICSA)	United Kingdom	Background
23	Concepts of Artificial Intelligence for Computer-Assisted Drug Discovery (2019)	ETH Zurich, Sichuan University	China, Switzerland	—
24	Deep networks for system identification: A survey (2024)	Linköping University, University of Padova, University of Washington	Italy, Sweden, United States	Methodology
25	Object-Centric Learning with Slot Attention (2020)	ETH Zurich, Google	Switzerland, United States	—
26	Deep Learning: Methods and Applications (2014)	Microsoft, Microsoft Research	United States	—
27	Convolutional Neural Networks for Speech Recognition (2014)	Microsoft Research, University of Toronto, York University	Canada, United States	Background
28	A Survey of Multimodal Retrieval-Augmented Generation (2025)	Huawei Cloud BU	—	—
29	Neural Network Methods for Natural Language Processing by Yoav Goldberg (2018)	Tsinghua University	China	—
30	Toward an Integration of Deep Learning and Neuroscience (2016)	Google DeepMind, Massachusetts Institute of Technology, Northwestern University	United Kingdom, United States	Background

Showing the 30 most-cited of 746 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).

Citing-text excerpts — how the field used this work

METHODOLOGY Deep Convolutional Neural Network Based Regression Approach for Estimation of Remaining Useful Life

“In this paper we treat RUL estimation problem as multivariate time series regression and solve it by adapting one particular deep learning model, namely Convolutional Neural Network (CNN) adapted from deep learning model for image classification [1,12,13], which is the first attempt to leverage deep learning to estimate RUL in prognostics.”

METHODOLOGY Deep networks for system identification: A survey

“Neural networks have been around for more than 70 years [209] and multilayer networks at least since 1970s, notable examples being the Necognitron used to determine characters [116] and the convolutional networks [175].”

FOLLOW-UP WORK

Catalyzing next-generation artificial intelligence through neuroai

2023 · Nature communications 14 (1), 1597, 2023 · 407 citations (GS)

Field-normalised: 277 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	Emerging NeoHebbian Dynamics in Forward-Forward Learning: Implications	TECNALIA, University of Deusto	Spain	—
2	Driver Behavior Modeling with Subjective Risk-Driven Inverse Reinforcement Learning	Wuhan University of Technology	China	—
3	Neurodatascience: Past, Present, and Future	University of California, Irvine Medical Center	United States	—
4	Cortex-Canvas: An Interactive Web Interface for Executing and Evaluating Models of Category-Selective Regions in Human Visual Cortex	Amrita Vishwa Vidyapeetham, Georgia Institute of Technology	India, United States	—
5	Mechanistic Foundations of Goal-Directed Control	University of the Basque Country	Spain	—

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

Toward next-generation artificial intelligence: Catalyzing the neuroai revolution

2022 · arXiv preprint arXiv:2210.08340, 2022 · 94 citations (GS)

Field-normalised: 54 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	Rethinking the Simulation vs. Rendering Dichotomy: No Free Lunch in Spatial World Modelling	Carnegie Mellon University, Johns Hopkins University, University of Michigan	United States	—

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

Contribution 3

Claim – Contribution 3

The researcher pioneered the application of multilayer graph transformer networks to reading checks, establishing a foundational framework later expanded into broader geometric deep learning paradigms.

The researcher's contribution centers on the seminal 1997 paper 'Reading checks with multilayer graph transformer networks,' which introduced a novel approach to processing structured data. This core work serves as the intellectual foundation for subsequent research, including the highly cited 2015 Nature article on deep learning and the 2017 IEEE Signal Processing Magazine paper on geometric deep learning, suggesting a sustained effort to generalize neural network architectures beyond Euclidean domains.

This line of work appears to address the challenge of applying deep learning techniques to non-Euclidean data structures. By moving from specific graph transformer applications in 1997 to broader theoretical frameworks in 2015 and 2017, the researcher demonstrates an original trajectory in extending deep learning capabilities to complex, geometric data types that traditional methods struggled to handle effectively.

The significance of this contribution is evidenced by substantial citation metrics and broad independent adoption. The core paper has accumulated 140 citations, while the follow-up works have garnered over 115,000 and 5,400 citations respectively. Notably, 98.1% of the 4,968 classified citations originate from independent researchers, indicating that this body of work has been widely recognized and utilized by the global scientific community outside the researcher's immediate circle.

INDEPENDENT CITATIONS FOR THIS CONTRIBUTION: 1,019 · 37 flagged influential by Semantic Scholar

CORE PAPER

[Reading checks with multilayer graph transformer networks](#)

1997 · Acoustics, Speech, and Signal Processing, 1997. ICASSP 1997. IEEE ..., 1997 · 140 citations (GS)

No.	Citing paper	Citing institution(s)	Country	S2
1	Deep Learning	Grenoble Alpes University	France	—

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

FOLLOW-UP WORK

[Deep learning](#)

2015 · Nature · 115,491 citations (GS)

Field-normalised: 74,422 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2015 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Rewrite the Stars (2024)	Microsoft, Northeastern University	United States	—
2	Adding Conditional Control to Text-to-Image Diffusion Models (2023)	Stanford University	United States	—
3	Perceiver: General Perception with Iterative Attention	DeepMind	United Kingdom	—
4	What do we need to build explainable AI systems for the medical domain?	Medical University Graz, The University of Manchester, Universität Hamburg	Austria, Cyprus, Germany	Background

No.	Citing paper	Citing institution(s)	Country	S2
5	Activation functions: Comparison of trends in practice and research for deep learning	University of Strathclyde	United Kingdom	Methodology
6	On the Opportunities and Risks of Foundation Models (2021)	Stanford Institute for Human-Centered Artificial Intelligence, Stanford University	United States	—
7	Large Language Models for Robotics: A Survey	University of Illinois Chicago	United States	Background
8	From System 1 to System 2: A Survey of Reasoning Large Language Models	AiShiWeiLai AI Research, Chinese Academy of Sciences, City University of Hong Kong and the Hong Kong University of Science and Technology (Guangzhou)	China, United Arab Emirates, United Kingdom	—
9	Fake news detection on social media: A data mining perspective	Arizona State University, Charles River Analytics, Michigan State University	United States	—
10	A Survey on Deep Learning: Algorithms, Techniques, and Applications	—	—	—
11	A Survey of Deep Active Learning	Carnegie Mellon University, Monash University, National Institute of Technology Kurukshetra	Australia, China, India	—
12	Transformers in Vision: A Survey (2022)	Inception Institute of Artificial Intelligence, MBZ University of Artificial Intelligence, Monash University	Australia, United Arab Emirates, United States	Background
13	Recent Advances in Natural Language Processing via Large Pre-trained Language Models: A Survey (2023)	Amazon AWS AI Labs, Harvard University, Synoptic Engineering	Spain, United States	—
14	A Survey on Evaluation of Large Language Models (2024)	Carnegie Mellon University, Hong Kong University of Science and Technology, Institute of Automation, Chinese Academy of Sciences	China, Hong Kong, United States	Background
15	Deep Multimodal Data Fusion (2024)	The University of Alabama at Birmingham	United States	Background
16	Object Detection Using Deep Learning, CNNs and Vision Transformers: A Review	Ibn Zohr University, University Ibn Zohr	—	Background
17	Advancements in Generative AI: A Comprehensive Review of GANs, GPT, Autoencoders, Diffusion Model, and Transformers (2024)	Bowie State University, Morgan State University, University of the District of Columbia	United States	Background
18	Deep Learning for Health Informatics	Imperial College London	United Kingdom	Background
19	Deep convolutional neural network for inverse problems in imaging	Dassault Aviation, École Polytechnique Fédérale de Lausanne, École polytech-	France, Switzerland	Background

No.	Citing paper	Citing institution(s)	Country	S2
		nique fédérale de Lausanne (EPFL)		
20	Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review	University of South Africa	South Africa	—
21	Efficient Processing of Deep Neural Networks: A Tutorial and Survey	Massachusetts Institute of Technology	United States	Background
22	Object Detection with Deep Learning: A Review	Hefei University of Technology, University of Louisiana at Lafayette	China, United States	Methodology
23	Deep Multi-Modal Object Detection and Semantic Segmentation for Autonomous Driving: Datasets, Methods, and Challenges	Robert Bosch GmbH, Ulm University, University of Stuttgart	Germany	Background
24	A Survey of the Usages of Deep Learning for Natural Language Processing	University of Colorado Colorado Springs	United States	Background
25	AI in Medical Imaging Informatics: Current Challenges and Future Directions	AstraZeneca, Boston Healthcare System, Emory University	Cyprus, Greece, New Zealand	—
26	A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises	Case Western Reserve University, Case Western Reserve University and Louis Stokes Cleveland Veterans Administration Medical Center, Emory University and Georgia Institute of Technology	China, Germany, Israel	Background
27	Self-Supervised Learning: Generative or Contrastive	Beijing Institute of Technology, Renmin University of China, Tsinghua University	China	—
28	Domain Adaptation for Medical Image Analysis: A Survey	University of North Carolina at Chapel Hill	United States	—
29	Deep learning for electroencephalogram (EEG) classification tasks: a review	University of Houston	United States	—
30	Introduction to Machine Learning, Neural Networks, and Deep Learning	Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, OHSU	United States	—

Showing the 30 most-cited of 1,011 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).

Citing-text excerpts — how the field used this work

METHODOLOGY Activation functions: Comparison of trends in practice and research for deep learning

“The numerous deep architectures include deep feedforward NNs, CNN, long short term memory, RNN, DBN, and deep generative models like deep Boltzmann machines, etc [1], [4].”

METHODOLOGY Object Detection with Deep Learning: A Review

“CNN is the most representative model of deep learning [26].”

FOLLOW-UP WORK

Geometric Deep Learning: Going beyond Euclidean data

2017 · IEEE Signal Processing Magazine · 5,458 citations (GS)

Field-normalised: 3,771 Semantic Scholar citations place it in the top 1% of Computer Science papers from 2017 indexed by Semantic Scholar, by citation count.

No.	Citing paper	Citing institution(s)	Country	S2
1	Open Graph Benchmark: Datasets for Machine Learning on Graphs (2020)	Harvard University, Microsoft, Stanford University	Germany, United States	Background
2	A review of large language models and autonomous agents in chemistry (2024)	—	—	—
3	Graph of Thoughts: Solving Elaborate Problems with Large Language Models (2023)	Cedar, ETH Zurich, Warsaw University of Technology	Poland, Switzerland	—
4	Machine learning for functional protein design (2024)	Harvard Medical School, Seismic Therapeutic	United States	—
5	How Attentive are Graph Attention Networks? (2021)	Carnegie Mellon University, Technion	Israel, United States	Background
6	Consistent Geometric Deep Learning via Hilbert Bundles and Cellular Sheaves	Harvard University, Northeastern University, Sakana AI	United States	—
7	SpheronizaTor: Spherical Voxelization for Interpretable Protein Microenvironment Modeling	University of Florida	United States	—

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

D. Citing-Institution Prestige & Geography

Top citing institutions

Institution	Country	World ranking	Citing papers
Tsinghua University	PR China	SCImago #8 · THE 12 · QS =17	151
University of California, Irvine Medical Center	United States	—	149
Stanford University	United States	SCImago #18 · THE =5 · QS 3	136
Massachusetts Institute of Technology	United States	SCImago #41 · THE 2 · QS 1	100
Carnegie Mellon University	United States	SCImago #266 · THE 24 · QS 52	93
Nanyang Technological University	Singapore	SCImago #137	76
New York University	United States	SCImago #116 · THE =31 · QS 55	72
University of Oxford	United Kingdom	SCImago #26 · THE 1 · QS 4	71
Microsoft Research	United States	—	68
ETH Zurich	Switzerland	THE 11 · QS 7	65
University of Toronto	Canada	SCImago #39 · THE 21 · QS 29	64
National University of Singapore	Singapore	SCImago #59 · THE 17 · QS 8	62
Chinese Academy of Sciences	China	SCImago #2	62

Institution	Country	World ranking	Citing papers
UC Berkeley	United States	—	61
Shanghai Jiao Tong University	China	SCImago #10 · THE 40 · QS =47	61

Geographic distribution of citing authors


Country	Citing papers
United States	1,877
China	1,149
United Kingdom	478
Germany	318
Canada	276
Australia	237
Singapore	180
France	175
Switzerland	157
India	145
Italy	124
South Korea	124

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.

E. Citation Growth Over Time

Distinct citing papers by publication year. Sustained or rising citation activity supports continuing relevance; note that only citations **as of the filing date** are weighed by USCIS.

1989 | 6
1991 | 8
1992 | 8
1993 | 8
1994 | 8
1995 | 3
1996 | 8
1997 | 7
1998 | 12
1999 | 10
2000 | 13
2001 | 3



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	Backpropagation Applied to Handwritten Zip Code Recognition	1,695	8 CFR 204.5(h)(3)(v) – Criterion 5
Contribution 2	Convolutional networks for images, speech, and time series	752	8 CFR 204.5(h)(3)(v) – Criterion 5
Contribution 3	Reading checks with multilayer graph transformer networks	1,019	8 CFR 204.5(h)(3)(v) – Criterion 5