

MICHELE MERLER

Principal Research Scientist, IBM Research AI

New York, NY, USA • +1 (646) 510-1702 • michele.merler@gmail.com

www.michelemerler.com • [Google Scholar](https://scholar.google.com/citations?user=...) • [LinkedIn](https://www.linkedin.com/in/michelemerler) • [X](https://x.com/michelemerler)

SUMMARY

AI researcher with 14+ years of experience at IBM Research spanning foundation models, evaluation, and multimodal machine learning. Core contributor to the Granite Code Models (open foundation models for code) and lead of large-scale empirical evaluations of LLMs, including an award-recognized study of bugs introduced by LLMs in code translation. Long track record in multimodal AI across vision, speech, audio, and language: Emmy Award-winning multimodal sports highlight curation deployed at major tennis and golf tournaments, interpretable speech biomarkers for ALS, and vision-language transfer for low-resource languages. 50 peer-reviewed publications and preprints, 20 U.S. patents and applications; Area Chair for WACV 2027, CVPR 2026, ECCV 2024, ACM Multimedia 2016-17. Senior Member of the IEEE.

PROFESSIONAL EXPERIENCE

IBM T.J. Watson Research Center — IBM Research AI

Yorktown Heights, NY

Principal Research Scientist

June 2025 – present

- Lead research on training, steering, and evaluation of LLMs and agents for code: constrained reasoning injection for code agents via nullspace editing (CRANE, 2026), and multi-task capability acquisition via data mixing vs. model merging (LLM4Code @ ICSE 2026).
- Contributed to the creation of ScarfBench, a benchmark for evaluating LLM-driven cross-framework application migration in enterprise Java (2026).
- Co-led one of the first empirical studies of AI coding assistant usage, effects, and requirements in the enterprise (LLM4Code @ ICSE 2026).
- Mentor PhD research interns on code LLM specialization, model merging, and agent steering.

Senior Research Scientist

June 2022 – June 2025

- **Granite Code Models:** core contributor to IBM's family of open foundation models for code — data preparation and curation, post-training of checkpoints for specialized enterprise Java use cases, and evaluation.
- **LLM evaluation:** contributed to “Lost in Translation” (ICSE 2024), a systematic study and taxonomy of bugs introduced by LLMs while translating code — Pat Goldberg Best Paper Honorable Mention (2025).
- **AI for digital medicine:** developed interpretable attention-based deep learning models for assessing dysarthria in ALS from remotely collected speech (npj Digital Medicine 2025; Best Paper Award, IEEE ICDH 2024).
- **Model efficiency:** task-agnostic knowledge distillation and neural architecture search for compressing transformer language models (EMNLP 2023).
- **Application modernization:** contrastive entity standardization (CoSiNES, ACL Matching Workshop 2023) and contributions to the Konveyor open-source community — IBM Outstanding Technical Achievement (2023).

Research Staff Member

May 2012 – June 2022

- **Multimodal sports highlight curation:** co-invented automatic highlight extraction from live sports broadcasts fusing audio, visual, and text excitement features; deployed at premier tennis and golf events including Wimbledon, the US Open, and the Masters — Technology & Engineering Emmy Award (2023), IEEE Transactions on Multimedia 2018.
- **Neural architecture search:** large-scale NAS for vision and language models, including architecture transferability analysis (AAAI 2021) and polyharmonic-spline search.
- **Face and attribute analysis:** Diversity in Faces dataset and deep learning facial attribute estimation; multimodal attribute inference from social media profiles (ICME 2015).
- **Fine-grained food recognition:** semantics-aware visual food recognition and the Snap, Eat, RepEat dietary-logging engine (ACM Multimedia 2016).
- **Medical imaging:** multimodal (visual + text) medical image modality recognition — 1st place, ImageCLEF Medical Modality Classification, 2012 and 2013 (MICCAI 2014).
- **Video understanding:** semantic model vectors for complex event recognition in large-scale video collections (IEEE Transactions on Multimedia 2012); transformer-based topic modeling and visualization.

Columbia University — High Level Vision Lab

Sept 2007 – May 2012

Graduate Research Assistant — Advisor: Prof. John Kender

- Analysis, indexing, and enhancement of unstructured presentation videos based on textual, graphical, and facial cues.

IBM T.J. Watson Research Center — Multimedia and Vision Group

Summers 2008–2010

Research Intern — Manager: Apostol Natsev; Mentors: Rong Yan, Gang Hua, Lexing Xie

- Complex video event detection, imbalanced boosting for ranking large-scale collections (CVPR 2009), and local feature-based representations for face recognition in multimedia collections.

California Institute for Telecommunications and Information Technology (Calit2), UC San Diego

Summer 2006

Research Intern — Advisor: Prof. Serge Belongie

- Built the GroZi multimedia grocery-recognition dataset and benchmarked state-of-the-art object detection and recognition algorithms on it (CVPR SLAM Workshop 2007).

Create-Net International Research Center, Trento, Italy

Summer 2004

Researcher — Advisor: Prof. Francesco De Natale

- Real-time human detection, visual tracking, and human-computer interaction algorithms.

EDUCATION

Ph.D. in Computer Science — Columbia University, New York

2013

Advisor: Prof. John Kender

M.S. in Computer Science — Columbia University, New York

2008

M.Eng. in Telecommunications Engineering — University of Trento, Italy

2007

Summa cum Laude

B.Eng. in Telecommunications Engineering — University of Trento, Italy

2004

Summa cum Laude

RESEARCH INTERESTS

Current focus: large language models for code (data curation, pre- and post-training, model merging, evaluation); LLM evaluation and benchmarking; multimodal foundation models; interpretable deep learning for digital medicine (ALS detection and monitoring from speech).

Broader interests: computer vision, NLP, multimedia indexing and retrieval, neural architecture search, video understanding and highlight extraction, large-scale semantic image/video retrieval.

Tools: Python, PyTorch, C/C++, OpenCV.

HONORS & AWARDS

- Pat Goldberg Best Paper Honorable Mention — “Lost in Translation: A Study of Bugs Introduced by Large Language Models while Translating Code” (2025)
- Best Paper Award, IEEE ICDH — “Harnessing Remote Speech Tasks for Early ALS Biomarker Identification” (2024)
- Technology & Engineering Emmy Award — AI/ML Curation of Sports Highlights (2023)
- Outstanding Reviewer, CVPR (2021)
- Best Digital Development, Yahoo Sports Tech Awards — Wimbledon Cognitive Highlights (2018)
- Best Reviewer Award, ICMR (2015)
- 1st Place, ImageCLEF Medical Image Modality Classification (2012, 2013)
- Yahoo! Key Scientific Challenge Award (2009)
- VideoOlympics “People's Choice Award” (group), IMARS Multimedia Retrieval System, ACM CIVR (2008)
- IBM Awards: Outstanding Technical Achievement (Konveyor Open-Source Community, 2023; Multimedia Group, 2012), Corporate Award (AI Video Enrichment and Editing, 2019), Research Division Awards (Watson Visual Recognition, 2018; Multimedia Semantic Modeling, 2013)
- Calit2 Summer Undergraduate Research Scholarship (2006)
- University of Trento International Cooperation and Mobility Program: University of California Scholarship (2005–2006)

PUBLICATIONS

Peer-reviewed publications and preprints. Full list also available on [Google Scholar](#).

- [50] Mingzhi Zhu, **Michele Merler**, Raju Pavuluri, Stacy Patterson. "CRANE: Constrained Reasoning Injection for Code Agents via Nullspace Editing". *arXiv preprint arXiv:2605.14084*, 2026.
- [49] Anudeep Pavuluri, B. McGinn, A. Saxena, G. Safta, S. Tamilselvam, Raju Pavuluri, **Michele Merler**, et al.. "ScarfBench: A Benchmark for Cross-Framework Application Migration in Enterprise Java". *arXiv preprint arXiv:2605.06754*, 2026.
- [48] Mingzhi Zhu, Boris Sobolev, Rahul Krishna, Raju Pavuluri, Stacy Patterson, **Michele Merler**. "Multi-task Code LLMs: Data Mix or Model Merge?". *LLM4Code Workshop at ICSE*, 2026.
- [47] Maja Vukovic, Rangeet Pan, Tin Kam Ho, Rahul Krishna, Raju Pavuluri, **Michele Merler**. "Usage, Effects and Requirements for AI Coding Assistants in the Enterprise: An Empirical Study". *LLM4Code Workshop at ICSE*, 2026.
- [46] **Michele Merler**, Carla Agurto, Julian Peller, Esteban Roitberg, Alan Taitz, Marcos A. Trevisan, Indu Navar, James D. Berry, Ernest Fraenkel, Lyle W. Ostrow, Guillermo A. Cecchi, Raquel Norel. "Clinical assessment and interpretation of dysarthria in ALS using attention based deep learning AI models". *npj Digital Medicine*, 2025.
- [45] Indu Navar Bingham, Raquel Norel, Esteban G. Roitberg, Julián Peller, Marcos A. Trevisan, Carla Agurto, **Michele Merler**, Diego E. Shalom, Felipe Aguirre, Iair Embon, Alan Taitz, Donna Harris, Amy Wright, Katie Seaver, Stacey Sullivan, Jordan R. Green, Lyle W. Ostrow, Ernest Fraenkel, James D. Berry. "Listener effort measures clinically meaningful change of dysarthria in amyotrophic lateral sclerosis". *Brain Communications*, 2025.
- [44] Raquel Norel, **Michele Merler**, Pavitra Modi. "Toward Continuous Neurocognitive Monitoring: Integrating Speech AI with Relational Graph Transformers for Rare Neurological Diseases". *arXiv preprint arXiv:2512.04938*, 2025.
- [43] Carla Agurto, **Michele Merler**, Esteban Roitberg, Alan Taitz, Marcos A. Trevisan, Diego E. Shalom, Julian Peller, Lyle W. Ostrow, Indu Navar, Ernest Fraenkel, James Berry, Guillermo A. Cecchi, Raquel Norel. "Harnessing Remote Speech Tasks for Early ALS Biomarker Identification". *IEEE International Conference on Digital Health (ICDH)*, 2024.
- [42] Carla Agurto, **Michele Merler**, Jenna Reinen, Pritish Parida, Guillermo Cecchi, Jeffrey L. Rogers, Navitas and Envision Studies Physician Author Group, Boston Scientific Research Scientists Consortium. "Exploring Chronic Pain Experiences: Leveraging Text and Audio Analysis to Infer Well-Being Metrics". *IEEE International Conference on Digital Health (ICDH)*, 2024.
- [41] Rikke Gade, **Michele Merler**, Graham Thomas, Thomas B. Moeslund. "The (Computer) Vision of Sports: Recent Trends in Research and Commercial Systems for Sport Analytics". *Chapter 14 in Computer Vision: Challenges, Trends, and Opportunities*, CRC Press, 2024.
- [40] Mayank Mishra, Matt Stallone, Gaoyuan Zhang, Yikang Shen, Aditya Prasad, Adriana Meza Soria, **Michele Merler**, Parameswaran Selvam, Saptha Surendran, Shivdeep Singh, Manish Sethi, Xuan-Hong Dang, Pengyuan Li, Kun-Lung Wu, Syed Zawad, Andrew Coleman, Matthew White, Mark Lewis, Raju Pavuluri, Yan Koyfman, Boris Lublinsky, Maximilien de Bayser, Ibrahim Abdelaziz, Kinjal Basu, Mayank Agarwal, Yi Zhou, Chris Johnson, Aanchal Goyal, Hima Patel, Yousaf Shah, Petros Zeros, Heiko Ludwig, Asim Munawar, Maxwell Crouse, Pavan Kapanipathi, Shweta Salaria, Bob Calio, Sophia Wen, Seetharami Seelam, Brian Belgodere, Carlos Fonseca, Amith Singhee, Nirmal Desai, David D. Cox, Ruchir Puri, Rameswar Panda. "Granite Code Models: A Family of Open Foundation Models for Code Intelligence". *arXiv preprint arXiv:2405.04324*, 2024.
- [39] Rangeet Pan, Ali Reza Ibrahimzada, Rahul Krishna, Divya Sankar, Lambert Pougues Wassi, **Michele Merler**, Boris Sobolev, Raju Pavuluri, Saurabh Sinha, Reyhaneh Jabbarvand. "Lost in translation: A study of bugs introduced by large language models while translating code". *International Conference on Software Engineering (ICSE)*, 2024.
- [38] Rangeet Pan, Rahul Krishna, **Michele Merler**, Raju Pavuluri, Saurabh Sinha, Maja Vukovic. "What would be an Ideal Code LLM? An Industrial Perspective". *2030 Software Engineering Workshop at FSE*, 2024.
- [37] Masayasu Muraoka, Bishwaranjan Bhattacharjee, **Michele Merler**, Graeme Blackwood, Yulong Li, Yang Zhao. "Cross-Lingual Transfer of Large Language Model by Visually-Derived Supervision Toward Low-Resource Languages". *ACM Multimedia*, 2023.
- [36] Takuma Udagawa, Aashka Trivedi, **Michele Merler**, Bishwaranjan Bhattacharjee. "A Comparative Analysis of Task-Agnostic Distillation Methods for Compressing Transformer Language Models". *EMNLP*, 2023.
- [35] Jiaqing Yuan, **Michele Merler**, Mihir Choudhury, Raju Pavuluri, Munindar P. Singh, Maja Vukovic. "CoSiNES: Contrastive Siamese Network for Entity Standardization". *Matching Workshop at ACL*, 2023.
- [34] Aashka Trivedi, Takuma Udagawa, **Michele Merler**, Rameswar Panda, Yousef El-Kurdi, Bishwaranjan Bhattacharjee. "Neural Architecture Search for Effective Teacher-Student Knowledge Transfer in Language Models". *arXiv preprint arXiv:2303.09639*, 2023.

- [33] Rameswar Panda, **Michele Merler**, Mayoore Jaiswal, Hui Wu, Kandan Ramakrishnan, Ulrich Finkler, Chun-Fu Chen, Minsik Cho, David Kung, Rogerio Feris, Bishwaranjan Bhattacharjee. "NASTransfer: Analyzing Architecture Transferability in Large Scale Neural Architecture Search". *35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [32] Ulrich Finkler, **Michele Merler**, Rameswar Panda, Mayoore S. Jaiswal, Hui Wu, Kandan Ramakrishnan, Chun-Fu Chen, Minsik Cho, David Kung, Rogerio Feris, Bishwaranjan Bhattacharjee. "Large Scale Neural Architecture Search with Polyharmonic Splines". *AAAI Workshop on Meta-Learning for Computer Vision (MeL4CV)*, 2021.
- [31] **Michele Merler**, Cicero N. Santos, Mauro Martino, Alfio M. Gliozzo, John R. Smith. "Covering the news with (AI) Style". *arXiv preprint arXiv:2002.02369*, 2020.
- [30] **Michele Merler**, Nalini Ratha, Rogerio S. Feris, John R. Smith. "Diversity in Faces". *arXiv preprint arXiv:1901.10436*, 2019.
- [29] **Michele Merler**, Dhiraj Joshi, Quoc-Bao Nguyen, Stephen Hammer, John Kent, Jinjun Xiong, Minh N. Do, John R. Smith, Rogerio S. Feris. "Automatic Curation of Sport Highlights using Multimodal Excitement Features". *IEEE Transactions on Multimedia*, 2018.
- [28] **Michele Merler**, Dhiraj Joshi, Quoc-Bao Nguyen, Stephen Hammer, John Kent, John R. Smith, Rogerio S. Feris. "Automatic Curation of Golf Highlights using Multimodal Excitement Features". *3rd Workshop on Computer Vision in Sports at CVPR*, 2017.
- [27] Dhiraj Joshi, **Michele Merler**, Quoc-Bao Nguyen, Stephen Hammer, John Kent, John R. Smith, Rogerio S. Feris. "IBM High-Five: Highlights From Intelligent Video Engine". *ACM Multimedia*, 2017.
- [26] Xiaolong Wang, Guodong Guo, **Michele Merler**, Noel C.F. Codella, M.V. Rohith, John R. Smith, Chandra Kambhamettu. "Leveraging multiple cues for recognizing family photos". *Journal of Image and Vision Computing*, 2017.
- [25] **Michele Merler**, Hui Wu, Rosario Uceda-Sosa, Quoc-Bao Nguyen, John R. Smith. "Snap, Eat, RepEat: a food recognition engine for dietary logging". *ACM Multimedia*, 2016.
- [24] Hui Wu, **Michele Merler**, Rosario Uceda-Sosa, John R. Smith. "Learning to make better mistakes: Semantics-aware visual food recognition". *ACM Multimedia*, 2016.
- [23] Quan Zou, Wei Liu, **Michele Merler**, Rongrong Ji. "Advanced learning for large-scale heterogeneous computing". *Neurocomputing Journal, Issue 217*, 2016.
- [22] **Michele Merler**, Liangliang Cao, John R. Smith. "You are what you tweet... pic! Gender prediction based on semantic analysis of social media images". *IEEE International Conference on Multimedia and Expo (ICME)*, 2015.
- [21] Junjie Cai, **Michele Merler**, Sharath Pankanti, Qi Tian. "Heterogeneous semantic level features fusion for action recognition". *IEEE International Conference on Multimedia Retrieval (ICMR)*, 2015.
- [20] Mani Abedini, Noel C.F. Codella, Jonathan H. Connell, Rahil Garnavi, **Michele Merler**, Sharath Pankanti, John R. Smith, Tanveer Syeda-Mahmood. "A generalized framework for medical image classification and recognition". *IBM Journal of Research and Development*, 2015.
- [19] John R. Smith, Liangliang Cao, Noel C.F. Codella, Matthew L. Hill, **Michele Merler**, Quoc-Bao Nguyen, E. Pring, Rosario A. Uceda-Sosa. "Massive-scale learning of image and video semantic concepts". *IBM Journal of Research and Development, Volume 59, Issue 2/3*, 2015.
- [18] Felix X. Yu, Liangliang Cao, **Michele Merler**, Noel Codella, Tao Chen, John R. Smith, Shih-Fu Chang. "Modeling attributes from category-attribute proportions". *ACM Multimedia*, 2014.
- [17] Noel Codella, Jonathan Connell, Sharath Pankanti, **Michele Merler**, John R. Smith. "Automated medical image modality recognition by fusion of visual and text information". *MICCAI*, 2014.
- [16] Noel Codella, Gang Hua, Liangliang Cao, **Michele Merler**, Leiguang Gong, Matt Hill, John R. Smith. "Large-scale video event classification using dynamic temporal pyramid matching of visual semantics". *ICIP*, 2013.
- [15] Mani Abedini, Liangliang Cao, Noel Codella, Jonathan H. Connell, Rahil Garnavi, Amir Geva, **Michele Merler**, Quoc-Bao Nguyen, Sharathchandra U. Pankanti, John R. Smith, Xingzhi Sun, Asaf Tzadok. "IBM research at ImageCLEF 2013 medical tasks". *AMIA ImageCLEF Workshop*, 2013.
- [14] Lisa Brown, Liangliang Cao, Shih-Fu Chang, Yu Cheng, Alok Choudhary, Noel Codella, Courtenay Cotton, Dan Ellis, Quanfu Fan, Rogerio Feris, Leiguang Gong, Matthew Hill, Gang Hua, John Kender, **Michele Merler**, Yadong Mu, Sharath Pankanti, John R. Smith, F.X. Yu. "IBM Research and Columbia University TRECVID-2013 Multimedia Event Detection (MED), Multimedia Event Recounting (MER), Surveillance Event Detection (SED), and Semantic Indexing (SIN) Systems". *NIST TRECVID Workshop*, 2013.
- [13] **Michele Merler**, Bert Huang, Lexing Xie, Gang Hua, Apostol Natsev. "Semantic Model Vectors for Complex Video Event Recognition". *IEEE Transactions on Multimedia*, 2012.

- [12] F. Yu, L. Cao, S.F. Chang, N. Codella, C. Cotton, D. Ellis, L. Gong, M. Hill, G. Hua, J. Kender, **M Merler**, Y. Mu, J. Smith. "IBM Research and Columbia University TRECVID-2012 Multimedia Event Detection (MED), Multimedia Event Recounting (MER), and Semantic Indexing (SIN) Systems". *NIST TRECVID Workshop*, 2012.
- [11] Liangliang Cao, Yuan-Chi Chang, Noel Codella, **Michele Merler**, Quoc-Bao Nguyen, John R. Smith. "Multimedia Analytics: Modality Classification and Case-Based Retrieval tasks of ImageCLEF 2012". *ImageCLEF Workshop*, 2012.
- [10] **Michele Merler**, John Kender. "Selecting the Best Faces to Index Presentation Videos". *ACM Multimedia*, 2011.
- [9] **Michele Merler**. "Analysis, indexing and visualization of presentation videos". *ACM Multimedia*, 2011.
- [8] Matthew Hill, Gang Hua, Apostol Natsev, John R. Smith, Lexing Xie, Bert Huang, **Michele Merler**, Hua Ouyang, Mingyuan Zhou. "IBM Research TRECVID-2010 Video Copy Detection and Multimedia Event Detection System". *NIST TRECVID Workshop*, 2010.
- [7] Apostol Natsev, Matthew Hill, John R. Smith, Lexing Xie, Rong Yan, Shenghua Bao, Dong Wang, **Michele Merler**, Yi Zhang. "IBM Research TRECVID-2009 video retrieval system". *NIST TRECVID Workshop*, 2009.
- [6] **Michele Merler**, John Kender. "Semantic Keyword Extraction via Adaptive Text Binarization of Unstructured Unsourced Video". *IEEE Conference on Image Processing (ICIP)*, 2009.
- [5] Rong Yan, Marc-Olivier Fleury, **Michele Merler**, Apostol Natsev, John R. Smith. "Large-Scale Multimedia Semantic Concept Modeling using Robust Subspace Bagging and Map-Reduce". *ACM Multimedia Workshop LS-MMRM*, 2009.
- [4] **Michele Merler**, Rong Yan, John R. Smith. "Imbalanced RankBoost for Efficiently Ranking Large-Scale Image/Video Collections". *CVPR*, 2009.
- [3] Jane Chang, **Michele Merler**, Paul Natsev, John R. Smith. "IBM Content Based Copy Detection System for TRECVID 2009". *NIST TRECVID Workshop*, 2009.
- [2] Apostol Natsev, Wei Jiang, **Michele Merler**, John Smith, Jelena Tesic, Lexing Xie, Rong Yan. "IBM Research TRECVID-2008 video retrieval system". *NIST TRECVID Workshop*, 2008.
- [1] **Michele Merler**, Carolina Galleguillos, Serge Belongie. "Recognizing Groceries in situ Using in vitro Training Data". *SLAM Workshop at CVPR*, 2007.

PATENTS & APPLICATIONS

- US Patent 12,561,950* — AI system and method for automatic analog gauge reading, 2026.
- US Patent App. 18/647,841* — Differential handwriting skill score based motor skill evaluation, 2025.
- US Patent App. 18/616,498* — Enabling a machine learning model to run predictions on domains where training data is limited by performing knowledge distillation from features, 2025.
- US Patent App. 18/614,861* — Automated recommendation tool to improve intelligibility in speech dysarthria, 2025.
- US Patent App. 18/540,228* — Neural network-based context-aware code translation and optimization, 2025.
- US Patent 12,475,316* — Mono-lingual language models using parallel data, 2025.
- US Patent 12,468,931* — Configuring a neural network using smoothing splines, 2025.
- US Patent 12,346,665* — Neural architecture search of language models using knowledge distillation, 2025.
- US Patent 12,249,041* — Oblique Image Rectification, 2025.
- US Patent App. 18/160,301* — Entity standardization for application modernization, 2024.
- US Patent App. 18/071,911* — Adaptable and explainable application modernization disposition, 2024.
- US Patent App. 17/977,880* — Attribute-based calibration for machine learning, 2024.
- US Patent App. 17/935,198* — Knowledge expansion for improving machine learning, 2024.
- US Patent 11,941,038* — Transparent and controllable topic modeling, 2024.
- US Patent 11,170,270* — Automatic generation of content using multimedia, 2021.
- US Patent 10,595,101* — Auto-curation and personalization of sports highlights, 2020.
- US Patent 10,282,677* — Individual and user group attributes discovery and comparison from social media visual content, 2019.
- US Patent 9,928,448* — Image classification utilizing semantic relationships in a classification hierarchy, 2018.
- US Patent 9,684,852* — Systems and methods for inferring gender by fusion of multimodal content, 2017.
- US Patent 9,177,229* — Kalman filter approach to augment object tracking, 2015.

PROFESSIONAL ACTIVITY

Associate Editor: IEEE Transactions on Multimedia, 2021–2023

Area Chair: WACV 2027; CVPR 2026; ECCV 2024; ICIAP 2019 (Brave New Ideas); ACM Multimedia 2016, 2017 (Multimedia Search and Recommendation)

Program Chair: International Workshop on Computer Vision in Sports (CVsports) @ CVPR, 2020–2026; Workshop on Fair, Data Efficient and Trusted Computer Vision (FA.DE.TR.CV) @ CVPR, 2020–2025; Workshops on Bias Estimation in Face Analytics (BEFA) @ ECCV 2018 and CVPR 2019

Guest Editor: Neurocomputing, Special Issue on Advanced Learning for Large-Scale Heterogeneous Computing, 2016

Award Committee: Best Paper Award Committee, ACM Multimedia 2017; NSF IIS Division Grant Reviewing Panelist, 2015

Organization: Demo Chair, MMM 2019; Registration Chair, VCIP 2017; Local Organization & Web Chair, ICMR 2016; Co-organizer, Greater New York Area Multimedia and Vision Meeting, 2012–2014

Reviewer (since 2008): CVPR, ECCV, ICCV, NeurIPS, ICML, AAAI, ACM Multimedia, ICMR, FG, ICPR, ICME; IEEE T-PAMI, IEEE T-MM, IEEE T-IP, IEEE T-CSVT, CVIU, ACM TOMM, Neurocomputing, J-BHI, and others

Memberships: Senior Member of the IEEE; IEEE Computer Society (CS); IEEE CS Technical Community on Pattern Analysis and Machine Intelligence (TCPAMI); IEEE Circuits and Systems Society (CAS); New York Academy of Sciences (NYAS)

MENTORSHIP

- **Mingzhi Zhu (PhD, RPI)** — model merging and reasoning injection for code LLMs (2025-2026).
- **Jiaqing Yuan (PhD, NCSU)** — contrastive Siamese networks for entity standardization (2022).
- **Aashka Trivedi (MS, NYU)** — NAS-guided language model distillation (2021).
- **Paul Pritz (PhD, Imperial College London)** — zero-shot learning for IoT device classification from network traffic (2021).
- **Xiaolong Wang (PhD, University of Delaware)** — face analysis in consumer photos (2015).
- **Junjie Cai (PhD, UT San Antonio)** — heterogeneous feature fusion for action recognition in videos (2014).

Thesis committees: Adrià Arbués Sanguësa (Universitat Pompeu Fabra, 2021); Chun-Yu (Claire) Tsai (Columbia University, 2017).

TEACHING

Adjunct Instructor, Columbia University — School of Professional Studies

Fall 2026 (upcoming)

- Applied Generative AI.

Instructor, Columbia University — Department of Computer Science

2009–2011

- COMS 1003 — Introduction to Computer Science and Programming in C (full course, 3 credits), Fall 2010 and Spring 2011.
- COMS 3101 — Programming Languages: MATLAB (short course, 1 credit), Fall 2009 and Spring 2010.

LANGUAGES & OTHER INTERESTS

Languages: Italian (native), English (fluent), Spanish (intermediate).

Other interests: Soccer, AI art, graphic novels, philosophy, light painting.