Seokeon Choi

Staff Research Engineer, Qualcomm AI Research

San Diego, CA | +1 (619) 953-8503 | bismex@gmail.com | seokchoi@qti.qualcomm.com

Website | Google Scholar: 3.4k+ citations | GitHub: 2k+ stars | LinkedIn | DBLP

Work Authorization: U.S. L-1B status; EB-2 NIW I-140 approved (Feb. 2023); I-485 filed (Apr. 2026); EAD expected in Q3 2026.

Summary

AI research engineer with a Ph.D. in Electrical Engineering and 10+ years of research experience across generative AI, multimodal learning, on-device learning, and computer vision. Currently building agentic AI systems at Qualcomm AI Research in the U.S.; previously worked in Qualcomm Korea's On-device Learning (ODL) team on efficient fine-tuning, diffusion personalization, PEFT/LoRA, quantization-aware workflows, and deployment-oriented evaluation. Published 20+ peer-reviewed papers across CVPR, ICCV, ECCV, NeurIPS, TPAMI, and TIFS, with multiple U.S. patent applications in generative AI and on-device learning.

Core Expertise

Generative AI	Diffusion personalization, LoRA/PEFT, training-free personalization, multi-human generation, synthetic data
Agentic AI	Hybrid on-device/cloud agents, agentic RAG, user memory/profile lifecycle, preference modeling
Efficient AI	On-device learning, low-precision training, quantization workflows, SLM reasoning efficiency, consistency models
Multimodal AI	LLMs, LVMs, LMMs, text-to-image systems, personalized generation, human-centric visual understanding
Computer Vision	Domain generalization/adaptation, person re-identification, detection/tracking, pose, segmentation
Tools	Python, PyTorch, TensorFlow, MATLAB, C/C++, model training/conversion/evaluation/deployment

Research Experience

Qualcomm AI Research **Nov. 2025 - Present | San Diego, CA**

Staff Research Engineer, Agentic AI Core System Architecture Team

- Developing hybrid agentic AI frameworks that combine on-device and cloud execution for data collection, memory extraction, preference/profile modeling, and agentic RAG over unstructured data.
- Improving reasoning performance and efficiency of small language models through test-time adaptation, reinforcement learning, on-device learning, and automated prompt optimization.
- Designing cost-aware benchmarking and evaluation frameworks for hybrid AI systems under resource constraints.

Qualcomm AI Research **Nov. 2023 - Nov. 2025 | Seoul, South Korea**

Staff Research Engineer, On-device Learning (ODL) Team

- Led efficient fine-tuning and on-device learning research for LLMs, LVMs, and LMMs, including PEFT, low-precision training, zeroth-order optimization, checkpointing, and diffusion guidance.
- Applied efficient fine-tuning to diffusion models and LLMs for deployment on Qualcomm chips, including SD1.5, SSD-1B, SANA 1.6B, Llama 3, and Qwen 2.5.
- Introduced consistency-model workflows and techniques to transfer personalized LoRA adapters to consistency models for efficient on-device inference.
- Published demos/results for single/multi-human generation and email personalization; collaborated on quantization, conversion, evaluation, and deployment workflows.
- Output: 7 conference papers across CVPR, NeurIPS, ICCV, and ICCVW; 6 U.S. patent applications.

Qualcomm AI Research **Sep. 2021 - Nov. 2023 | Seoul, South Korea**

Senior Research Engineer, On-device Learning (ODL) Team

- Developed domain generalization and adaptation algorithms for on-device learning; integrated research output into Qualcomm toolkit workflows.
- Proposed source-free object detection for surveillance applications and delivered the method to software teams for SDK integration.
- Output: 4 conference papers across CVPR, ECCV, and CVPRW; 3 U.S. patent applications.

Google **Jun. 2021 - Sep. 2021 | Seoul, South Korea**

Research Intern, TensorFlow Model Optimization Team

- Contributed low-bit quantization code to the TensorFlow Model Optimization Toolkit and worked within Google's code submission, documentation, and code-review workflows.
- Conducted survey, benchmarking, and analysis work for low-bit quantization methods and deployment-oriented model optimization.

Carnegie Mellon University, Language Technologies Institute (LTI) **Jan. 2020 - Jul. 2020 | Pittsburgh, PA**

Visiting Researcher, advised by Prof. Alex Hauptmann

- Conducted research on robust long-term object tracking and large-scale multi-camera vehicle re-identification for intelligent surveillance systems.
- Participated in the AI City Challenge and developed an algorithm from coursework/project work that led to an ECCVW 2020 long-term object tracking publication with Prof. Hauptmann.
- Completed six courses through the IITP-supported CMU program and worked on an AI chatbot project in addition to research.

KAIST, Computational Intelligence Lab **Mar. 2015 - Jun. 2021 | Daejeon, South Korea**

Research Assistant; Ph.D. and M.S. researcher

- Led and contributed to multiple research/industry projects across person re-identification, domain adaptation/generalization, object detection/tracking, pose estimation, gait recognition, and 3D reconstruction.
- Served as lab leader for multiple years and contributed to papers co-authored with Prof. Changick Kim across CVPR, ICCV, WACV, BMVC, TIFS, and TPAMI.

Education

Ph.D., Electrical Engineering, KAIST (2017 - 2021). Thesis: Person Re-Identification to Bridge Domain Gaps. Advisor: Prof. Changick Kim. GPA: 4.06/4.30.

M.S., Electrical Engineering, KAIST (2015 - 2017). Thesis: Robust Model-based Gait Recognition via Candidate Selection and Pose-aware Decision Fusion. GPA: 4.03/4.30.

B.S., Electronic and Electrical Engineering, Sungkyunkwan University (2009 - 2015). Summa Cum Laude. GPA: 4.41/4.50.

Research Interests

- **LLM and Agentic AI:** Hybrid on-device/cloud agentic systems, memory extraction, user profile lifecycle, preference modeling, and agentic RAG; Small-language-model reasoning efficiency through test-time adaptation, reinforcement learning, on-device learning, and automated prompt optimization
- **Generative AI and Multimodality:** Personalization / efficient fine-tuning for diffusion and foundation models: [C25], [C23], [C22], [C21], [C20], [P03]; Multi-human generation, benchmarks, and controllable human synthesis: [C26], [C24], [P02]; Synthetic data and representation generation: [C18], [C17], [C11], [C06]
- **Generalizability and Transferability:** Domain generalization: [C19], [C18], [C14]; Domain adaptation and test-time adaptation: [C16], [C06], [P01]; Cross-modal learning: [C11]; Transfer, contrastive, and few-shot learning: [C17], [C13]
- **Human Understanding:** Person re-identification: [C14], [C11], [P01]; Pose estimation: [C12]; Action detection: [J02]; Gait recognition / biometrics: [J01]
- **Machine Perception:** Object detection: [C06]; Object tracking: [C09], [C07], [C05]; Object segmentation: [C10]; 3D reconstruction / multi-view stereo: [C15], [C03]; Challenge reports: [C08], [C04]; Low-level vision, image enhancement, and template matching: [C02], [C01]

Full Publication List

International Journals

[J02] Sumin Lee, Hyunjun Eun, Jinyoung Moon, **Seokeon Choi**, Yoonhyung Kim, Chanho Jung, and Changick Kim, "Learning to Discriminate Information for Online Action Detection: Analysis and Application," *IEEE TPAMI, 2022*. [Link]

[J01] Seokeon Choi, Jonghee Kim, Wonjun Kim, and Changick Kim, "Skeleton-based Gait Recognition via Robust Frame-level Matching," *IEEE TIFS, 2019*. [Link]

International Conferences

[C26] Shubhankar Borse, Phuc Pham, Farzad Farhadzadeh, **Seokeon Choi**, Phong Ha Nguyen, Anh Tuan Tran, Sungrack Yun, Munawar Hayat, and Fatih Porikli, "Ar2Can: An Architect and an Artist Leveraging a Canvas for Multi-Human Generation," *CVPR 2026*. [Link]

[C25] Sunghyun Park*, Jeongho Kim*, Hyoungwoo Park, Debasmit Das, Sungrack Yun, Munawar Hayat, Jaegul Choo, Fatih Porikli, and **Seokeon Choi**, "Memory-Efficient Fine-Tuning Diffusion Transformer via Dynamic Patch Sampling and Block Skipping," *CVPR 2026*. [Link]

[C24] Shubhankar Borse, **Seokeon Choi**, Sunghyun Park, Jeongho Kim, Shreya Kadambi, Risheek Garrepalli, Sungrack Yun, Munawar Hayat, and Fatih Porikli, "MultiHuman-Testbench: Benchmarking Image Generation for Multiple Humans," *NeurIPS 2025*. [Link]

[C23] Seokeon Choi*, Sunghyun Park*, Hyoungwoo Park, Jeongho Kim, and Sungrack Yun, "Memory-Efficient Personalization of Text-to-Image Diffusion Models via Selective Optimization Strategies," *ICCVW 2025, oral presentation, best paper award*. [Link]

[C22] Sunghyun Park*, **Seokeon Choi***, Sungrack Yun, and Fatih Porikli, "Steering Guidance for Personalized Text-to-Image Diffusion Models," *ICCV 2025*. [Link]

[C21] Debasmit Das, Hyoungwoo Park, Munawar Hayat, **Seokeon Choi**, Sungrack Yun, and Fatih Porikli, "ConsNoTrainLoRA: Data-driven Weight Initialization of Low-rank Adapters using Constraints," *ICCV 2025*. [Link]

[C20] Wonguk Cho, **Seokeon Choi**, Debasmit Das, Matthias Reisser, Taesup Kim, Sungrack Yun, and Fatih Porikli, "Hollowed Net for On-Device Personalization of Text-to-Image Diffusion Models," *NeurIPS 2024*. [Link]

[C19] Seunghan Yang, **Seokeon Choi**, Hyunsin Park, Sungha Choi, Simyung Chang, and Sungrack Yun, "Feature Diversification and Adaptation for Federated Domain Generalization," *ECCV 2024*. [Link]

[C18] Seokeon Choi, Debasmit Das, Sungha Choi, Seunghan Yang, Hyunsin Park, and Sungrack Yun, "Progressive Random Convolutions for Single Domain Generalization," *CVPR 2023*. [Link]

[C17] Taekyung Kim, Debasmit Das, **Seokeon Choi**, Minki Jeong, Seunghan Yang, Sungrack Yun, and Changick Kim, "Neural Transformation Network to Generate Diverse Views for Contrastive Learning," *CVPRW 2023*. [Link]

[C16] Sungha Choi, Seunghan Yang, **Seokeon Choi**, and Sungrack Yun, "Improving Test-Time Adaptation via Shift-agnostic Weight Regularization and Nearest Source Prototypes," *ECCV 2022*. [Link]

[C15] Taekyung Kim, Jaehoon Choi, **Seokeon Choi**, Dongki Jung, and Changick Kim, "Just a Few Points are All You Need for Multi-view Stereo: A Novel Semi-supervised Learning Method for Multi-view Stereo," *ICCV 2021*. [Link]

[C14] Seokeon Choi, Taekyung Kim, Minki Jeong, Hyoungseob Park, and Changick Kim, "Meta Batch-Instance Normalization for Generalizable Person Re-Identification," *CVPR 2021*. [Link]

[C13] Minki Jeong, **Seokeon Choi**, and Changick Kim, "Few-shot Open-set Recognition by Transformation Consistency," *CVPR 2021*. [Link]

[C12] Sangbum Choi, **Seokeon Choi**, and Changick Kim, "MobileHumanPose: Toward Real-time 3D Human Pose Estimation in Mobile Devices," *CVPRW 2021*. [Link]

[C11] Seokeon Choi, Sumin Lee, Youngeun Kim, Taekyung Kim, and Changick Kim, "Hi-CMD: Hierarchical Cross-Modality Disentanglement for Visible-Infrared Person Re-Identification," *CVPR 2020*. [Link]

[C10] Youngeun Kim, **Seokeon Choi**, Hankyeol Lee, Taekyung Kim, and Changick Kim, "RPM-Net: Robust Pixel-Level Matching Networks for Self-Supervised Video Object Segmentation," *WACV 2020*. [Link]

[C09] Seokeon Choi, Junhyun Lee, Yunsung Lee, and Alex Hauptmann, "Robust Long-Term Object Tracking via Improved Discriminative Model Prediction," *ECCVW 2020*. [Link]

[C08] Matej Kristan, et al. (including **Seokeon Choi**), "The eighth visual object tracking VOT2020 challenge results," *ECCVW 2020*. [Link]

[C07] Hankyeol Lee, **Seokeon Choi**, Youngeun Kim, and Changick Kim, "Bilinear Siamese Networks with Background Suppression for Visual Object Tracking," *BMVC 2019, spotlight*. [Link]

[C06] Taekyung Kim, Minki Jeong, Seunghyeon Kim, **Seokeon Choi**, and Changick Kim, "Diversify and Match: A Domain Adaptive Representation Learning Paradigm for Object Detection," *CVPR 2019*. [Link]

[C05] Hankyeol Lee*, **Seokeon Choi***, and Changick Kim, "A Memory Model based on the Siamese Network for Long-term Tracking," *ECCVW 2018*. [Link]

[C04] Matej Kristan, et al. (including **Seokeon Choi**), "The sixth visual object tracking VOT2018 challenge results," *ECCVW 2018*. [Link]

[C03] Hankyeol Lee, **Seokeon Choi**, and Changick Kim, "A Robust Rectification Algorithm for Non-Calibrated Multi-view Images," *3DSA 2018*. [Link]

[C02] Jonghee Kim, Jinsu Kim, **Seokeon Choi**, Muhammad Abul Hasan, and Changick Kim, "Robust Template Matching Using Scale-Adaptive Deep Convolutional Features," *APSIPA ASC 2017*. [Link]

[C01] Seokeon Choi, Jonghee Kim, and Changick Kim, "A Selective Upscaling Method via Super-Resolution Validity Measure," *ICCE-Asia 2016*. [Link]

Preprints

[P03] Seokeon Choi*, Sunghyun Park*, Hyoungwoo Park, Jeongho Kim, and Sungrack Yun, "Mix-Opt: Mixed Optimization for Memory-Efficient Personalization of Text-to-Image Diffusion Models," *Preprint, 2025*.

[P02] Jeongho Kim, Sunghyun Park, Hyoungwoo Park, Sungrack Yun, Jaegul Choo, and **Seokeon Choi**, "From Wardrobe to Canvas: Wardrobe Polyptych LoRA for Part-level Controllable Human Image Generation," *Preprint, 2025*. [Link]

[P01] Youngeun Kim, **Seokeon Choi**, Taekyung Kim, Sumin Lee, and Changick Kim, "Learning to Align Multi-Camera Domains using Part-Aware Clustering for Unsupervised Video Person Re-Identification," *Preprint, 2019*. [Link]

Domestic Papers

[D07] Kibum Yun, **Seokeon Choi**, and Changick Kim, "Performance Analysis of Discriminative Correlation Filter-based Object Tracking Methods for Small Objects," *IEIE Fall Conference, 2020*. [Link]

[D06] Seokeon Choi, Kyujin Shim, JunHaeng Lee, Taeyoung Lee, and Changick Kim, "A Comprehensive Study of Long-term Object Tracking," *The Magazine of the IEIE, 2020*. [Link]

[D05] Seokeon Choi, Kyujin Shim, JunHaeng Lee, Taeyoung Lee, and Changick Kim, "A Survey of Recent Trends in Short-term Object Tracking," *The Magazine of the IEIE, 2020*. [Link]

[D04] Kyujin Shim, **Seokeon Choi**, JunHaeng Lee, Taeyoung Lee, and Changick Kim, "Trends in Object Classification," *The Magazine of the IEIE, 2020*. [Link]

[D03] Seokeon Choi, Jonghee Kim, and Changick Kim, "Introduction to Gait Recognition," *The Magazine of the IEIE, 2019*. [Link]

[D02] Hankyeol Lee, **Seokeon Choi**, and Changick Kim, "Performance Evaluation of Local Descriptor Algorithms for Narrow Baseline Stereo," *IEIE Fall Conference, 2017*. [Link]

[D01] Seokeon Choi, Jonghee Kim, and Changick Kim, "A Region-based Texture Synthesizability Prediction Method for High-Resolution Image Generation," *Korea Multimedia Society Fall Conference, 2015*. [Link]

Theses

[T02] Seokeon Choi, "Person re-identification to bridge domain gaps," *Doctoral dissertation, KAIST, 2021*. [Link]

[T01] Seokeon Choi, "Robust Model-based Gait Recognition via Candidate Selection and Pose-aware Decision Fusion," *Master's thesis, KAIST, 2017*. [Link]

U.S. Patents

[U07] Seokeon Choi, Sunghyun Park, and Sungrack Yun, "Spatiotemporal Attention in Generative Machine Learning Models," *U.S. Patent Application No. 19/019,882*. [Link]

[U06] Seokeon Choi, Sunghyun Park, and Sungrack Yun, "Personalized Output Generation in Generative Artificial Intelligence Models," *U.S. Patent Application No. 18/959,042*. [Link]

[U05] Wonguk Cho, Matthias Reisser, Debasmit Das, **Seokeon Choi**, Sungrack Yun, and Fatih Porikli, "Finetuning One or More Neural Networks," *U.S. Patent Application No. 18/663,903*. [Link]

[U04] Seokeon Choi, Juntae Lee, and Jaewon Choi, "Privacy-Aware Multi-Modal Generative Autoreply," *U.S. Patent Application No. 18/454,456*. [Link]

[U03] Seunghan Yang, **Seokeon Choi**, Hyunsin Park, Sungha Choi, and Sungrack Yun, "Client-agnostic learning and zero-shot adaptation for federated domain generalization," *U.S. Patent Application No. 18/238,998*. [Link]

[U02] Seokeon Choi, Sungha Choi, Seunghan Yang, Hyunsin Park, Debasmit Das, and Sungrack Yun, "Semantic-aware random style aggregation for single domain generalization," *U.S. Patent Application No. 18/157,723*. [Link]

[U01] Sungha Choi, Seunghan Yang, **Seokeon Choi**, and Sungrack Yun, "Test-time adaptation with unlabeled online data," *U.S. Patent Application No. 18/086,586*. [Link]

Industry Collaborations

- **Hanwha R&D Center**, Team Leader, Small Object Tracking (Sep. 2020 - Aug. 2021). small-object/IR tracking, military applications, DCF-based methods; related ECCVW 2020 long-term tracking paper.
- **LIG Nex1**, Team Member, Image Enhancement and Object Tracking (Mar. 2020 - Aug. 2020). fast-moving object tracking, IR-based tracking, embedded systems, super-resolution, deblurring, and video stabilization.
- **Hanwha Systems**, Team Leader, AI-Based Object Tracking and Recognition (Mar. 2018 - Feb. 2020). small-object/IR tracking and IR image generation; related BMVC 2019 spotlight and ECCVW 2018 long-term tracking papers.
- **ETRI**, Team Leader, Image Rectification for Super Multi-View Display (Sep. 2017 - Feb. 2018). super multi-view display, image rectification, feature extraction/matching, and robust estimation.
- **SAIT**, Team Leader, IR Image Quality Assessment (Mar. 2017 - Aug. 2017). image quality assessment, face recognition, score-level fusion, and adaptive learning.
- **Samsung DMC / Samsung Research**, Team Member, Super-Resolution and High-Resolution Creation (Mar. 2015 - Feb. 2017). video super-resolution, metadata-guided texture transfer, texture classification, and perceptual high-resolution creation.

Awards, Honors, and Scholarships

- Best Paper Award, ICCVW 2025 LIMIT Workshop, 2025.
- Top Reviewer, NeurIPS 2025.
- CVPR 2021 Doctoral Consortium Presenter, mentor: Tong Xiao from Meta.
- KAIST Chul-Hi Han Augustino Scholarship Foundation academic scholarship, 2021.
- Samsung Humantech Paper Award, Silver Prize, top 2%, 2021.
- Samsung Humantech Paper Award, Finalist, top 8%, 2020.
- Carnegie Mellon University visit funded by Korean government, 2020.
- CVPR 2020 AI City Challenge, 16th prize, vehicle re-identification track.
- ECCV 2020 VOT2020-LT Challenge, 5th prize.
- ECCV 2018 VOT2018-LT Challenge, 3rd prize.
- Samsung Advanced Institute of Technology Industry-Academic Scholarship, 2017-2021.
- Summa Cum Laude, Sungkyunkwan University, 2015.
- Best Work/Thesis Award, Sungkyunkwan University, 2014.
- Dean's List, Sungkyunkwan University, 2012 and 2013.
- National Science and Technology Scholarship, 2012-2014.
- Academic scholarship, Sungkyunkwan University, 2009 and 2012.

Professional Activities

- **Conference Reviewer:** NeurIPS (2025, 2026); CVPR (2023, 2024, 2025, 2026); ICCV (2023, 2025); ECCV (2024, 2026); ACM MM (2025); WACV (2025).
- **Journal Reviewer:** IJCV (2023, 2024); IEEE TNNLS (2022, 2023); IEEE TIP (2023); IEEE TIFS (2022); IEEE TMM (2020-2023); IEEE TCSVT (2022); IEEE TGRS (2021); IEEE SPL (2021, 2022); IEEE IoT Journal (2021); IET Computer Vision (2022); Neurocomputing (2020).
- **Teaching Assistant:** Programming Structure for Electrical Engineering; Image Engineering; Image Understanding; Electronics Design Lab; Signals and Systems.
- **Invited Talks:** Solving the Domain Shift Problem in Deep Learning (KAIST, 2023); Person Re-Identification to Bridge Domain Gaps (ETRI, 2020); Introduction to Person Re-Identification (SAIT, 2020); Introduction to Long-Term Object Tracking (SAIT, 2020).

Skills

Languages: Korean (native), English (fluent). **Programming:** Python, PyTorch, MATLAB, TensorFlow, C/C++.

References available upon request. Last updated: July 2026.