Eric Minwoo Kim

New York, NY (NYU Tandon exchange, Fall 2026) | km7531@nyu.edu | erickim.mk20@gmail.com | ericminwookim.com

Research Interests

Scientific guarantees for modern AI systems in high-risk settings. My core focus is safe autonomy through distribution-free uncertainty quantification, with conformal prediction as my main tool so far, applied to robotics and physical AI, from world models to foundation-model policies. I am also interested in certified robustness, runtime safety, and, increasingly, the interpretability of black-box models.

Education

Korea University — B.Eng. in Electrical Engineering (Intensive major) *Mar. 2022 – Aug. 2027 (Expected)*
GPA: 3.84/4.5 (93.4/100) · **Last two years (62 credits, 2025–2026): 4.07/4.5**

New York University, Tandon School of Engineering — Exchange semester *Sept. 2026 – Dec. 2026*

Research Experience

Undergraduate Researcher, Machines in Motion Laboratory *Sept. 2026 – Present*
New York University, Tandon School of Engineering · Advisor: Prof. Ludovic Righetti

- Studying how physically grounded conformal safety guarantees carry from the perception boundary into a learned world model and the modules downstream of it (planning, RL policies), extending the workshop paper's latent-space result to the robot-learning control pipeline (in progress).
- Studying how far conformal coverage guarantees on world-model rollouts degrade under distribution shift, and how much weighted, online, and conditional recalibration restore them.
- Prepared a literature review of conformal prediction under distribution shift to scope the project.

Undergraduate Student Researcher *Dec. 2024 – Present*
High Performance and Intelligence Computing (HPIC) Lab, Korea University

KEPT: Conformal Keepouts Kept in Planning for Enclosed-Space Navigation in Degraded Sensing
First author · arXiv:2610.37217 (cs.RO), Oct. 2026 · journal submission in preparation

- Proposed a severity-conditioned conformal keepout for navigation in smoke, dust and fog: Mondrian split-conformal margins calibrated offline per severity class and selected at run time per 30° azimuth wedge and LiDAR sweep from a label-free return-fraction estimate. On held-out ground-robot smoke recordings the per-class keepout covered the unsafe set at 0.900 [0.885, 0.913] against a 0.90 target and cut the coverage spread between severity quartiles from 0.042 (one pooled margin) to 0.027 at a smaller keepout share.
- Designed CPMC, a planning layer that enforces one braking-distance barrier condition on the certified keepout at plan time (re-timing of B-spline references) and at every command (a one-step discrete control-barrier projection), with the measured tracking residual as its only margin and the keepout never relaxed; proved per-stratum validity, clearance from coverage, and invariance of the filtered reference between certificate updates.
- Ran a pre-registered, hash-locked closed-loop campaign of 1,344 simulated quadrotor flights (7 arms × 96 cells × 2 calibration regimes) in a foggy tunnel on a FlightBench-based harness (ROS, Docker, headless GPU) with a fog model fitted to recorded smoke sweeps: against the conformal deployment pattern the stack entered the keepout in 8 flights vs. 27 and reached 132 vs. 126 goals on 189 matched cells (entries $p = 0.004$ in the transferred regime; reach +3.2 pp with per-regime CIs including zero), with pose-paired bootstrap intervals and sign-flip tests.
- Authored the full manuscript as first author: problem statement, propositions and proofs, harness construction, all experiments and analysis.

Conformal Safety Certificates in Learned Representations: What Survives and What Must Be Rebuilt
First author · accepted (poster), NeurIPS 2026 Workshop on Robot Learning with World Models

- Designed a three-arm study of how a conformal safety certificate behaves inside a learned latent state (calibrate-then-encode, latent-native calibration, and a metre-valued decoder head), instantiated on real subterranean LiDAR calibration data (160 scenes, 60/50/50 train/calibration/test split).
- Found that statistical coverage transports across deterministic encoders for free (0.89–0.95 realized in every arm) while physical units and region geometry depend on the readout: latent-coordinate sets forbid up to 96% of provably-safe space under frozen DINOv2 patch features, while a small clearance head restores metre-valued margins at native-shell geometry (usable-free-space ratios 0.99/1.00).
- Pre-registered the frozen-DINOv2 second encoder and deviation handling; all results reported with scene-level bootstrap confidence intervals (B=3000, n=50 test scenes).

Attention-Based Qubit Mapping for Scalable Quantum Compilation
Second author · submitted to IEEE Access

- Reimplemented prior DNN-based qubit-mapping methods and analyzed their scalability limits on larger IBM BackendV2-style hardware.
- Designed and evaluated conditional U-Net mapper variants using circuit-as-image inputs, hardware-token attention conditioning, Sinkhorn assignment, and Hungarian decoding.
- Ran ablation experiments against SABRE/NASSC-style baselines, diagnosing routed PST gaps, non-adjacent mappings, mode collapse, and proxy-loss failure modes.
- Iterated on loss/training design, including route/SWAP/non-adjacency penalties, diversity regularization, K-stratified sampling, and active-row assignment.
- Contributed to the final attention/QAP-based optimizer by connecting circuit-hardware attention mechanisms with cost-aware qubit assignment, and to manuscript writing, related-work synthesis, and architecture/pipeline figure design.

Deep Learning for Induction Motor Fault Diagnosis *research assistant · acknowledged in the submitted manuscript*

- Reconstructed prior LSTM-based fault-diagnosis models and implemented them under the project's experimental conditions.
- Implemented and profiled DTW/RGW-based augmentation variants (DTW, RGW, DTW-sD, RGW-sD) for robustness to local timing distortion.
- Ran ablation experiments with ResNet/LSTM models across permutation augmentation, DTW-based augmentation, noisy inputs, FFT-transformed signals, and shifted operating conditions.
- Analyzed sensor-modality and operating-condition effects across vibration/current data, supply-frequency variation, and noisy data distributions; prepared the confusion-matrix figure and several results tables for the manuscript.

Publications and Manuscripts

KEPT: Conformal Keepouts Kept in Planning for Enclosed-Space Navigation in Degraded Sensing — first author; **arXiv:2610.37217** (cs.RO), Oct. 2026; journal submission in preparation. **Role:** method design, theory (propositions and proofs), simulation-harness construction, all experiments and analysis, full manuscript writing.

Conformal Safety Certificates in Learned Representations: What Survives and What Must Be Rebuilt — first author; accepted (poster), NeurIPS 2026 Workshop on Robot Learning with World Models (non-archival). **Role:** study design, all experiments and analysis, full manuscript writing.

AtteQuC: Attention-Based Qubit Mapping for Scalable Quantum Compilation — second author; submitted to IEEE Access, under review. **Role:** baseline reimplementation, U-Net prototyping, ablations, writing, technical figures.

Honors and Scholarships

KU RISE Global TECH Program *2026*
Regional Innovation System & Education (RISE), Ministry of Education / Korea University — tuition and airfare, research semester at NYU Tandon

Korea–U.S. Student Exchange Program in High-Tech Industry, 5th cohort *2026*
Ministry of Trade, Industry and Resources / Korea Institute for Advancement of Technology (KIAT) — national scholarship in designated advanced-industry fields

Leadership and Activities

AKUSSA (Korea University society for US graduate-study preparation) — Assistant board member, 2026–present

One and Zero (OaZ) (Korea University E.E. software society) — Member, Mar. 2022–present; led a team to the Stage 2 National/Regional Round of WorldQuant's International Quant Championship 2025

KU Debate Club (KUDC) — Competitive British Parliamentary debate in English, Sept. 2024–present

Technical Skills

Programming / ML: Python, PyTorch, TensorFlow/Keras, NumPy, Pandas, scikit-learn, C/C++
Robotics / Safe Autonomy: conformal prediction, distribution-free uncertainty quantification, SDF/ESDF motion planning, Fast-Planner, ROS, Gazebo/Unity-Flightmare (FlightBench), Docker, LiDAR/thermal/depth processing, closed-loop evaluation, statistical testing (bootstrap CIs, two-proportion tests)
Modeling & Systems: attention/U-Net/ResNet/LSTM architectures, custom loss design, Sinkhorn/Hungarian assignment, Qiskit and IBM BackendV2 workflows, DTW-based augmentation, SOTA reimplementation, ablation studies, Git, Linux, wandb, LaTeX/Overleaf
Languages: Korean (native), English (fluent)