

Joohyung Lee

Research Statement

ML Researcher, AITRICS (Seoul) | PhD applicant (Fall 2026) | M.S. EE, Texas A&M University

Email: joohyunglee.research@gmail.com Scholar: Google Scholar GitHub: github.com/aitrics-chris

Research Focus

Despite rapid progress in foundation models and learned world models, even state-of-the-art systems can violate basic physical and logical regularities, e.g., producing anatomically implausible medical images, undermining trustworthiness in high-stakes domains like healthcare and in high-impact generative AI applications. My research asks what regularities learning systems should respect and how to enforce them at scale without sacrificing downstream utility. Across projects, I formalize domain regularities (symmetries, acquisition mechanisms, temporal dynamics) and translate them into scalable learning mechanisms, from SSL objectives and soft regularizers to architecture-level inductive biases.

Selected contributions

- **Soft Equivariance Regularization (SER) for invariant SSL (ICLR 2026, accepted; first author).** SER decouples invariance and equivariance across layers: train the *final embedding* with an invariant SSL objective, while softly regularizing *intermediate, spatially structured* representations toward equivariance via analytically specified group actions (no labels; no auxiliary networks). On ImageNet-1k MoCo-v3 (ViT-S), SER improves linear-eval Top-1 by +0.84, ImageNet-C/P by +1.11/+1.22, and frozen-backbone COCO detection by +1.7 mAP, with $1.008\times$ training overhead. OpenReview.
- **Acquisition-physics priors in clinical imaging (first author).** In pathology, I used disease-negative slides as anchors to suppress acquisition-induced nuisance variation in multiple-instance learning (ICASSP 2024). In MRI, I introduced a late-fusion 3D design that respects anisotropic acquisition geometry (MICCAI 2022).
- **Temporal regularities for multimodal EHR forecasting (co-first/corresponding author).** I proposed a predictive-coding SSL framework for multimodal clinical trajectories under irregular sampling and missingness (ICLR 2023 TML4H workshop; Best Paper Honorable Mention) and introduced a shared time-embedding with modality-aware attention for cross-modal alignment (MLHC 2023 Spotlight).

PhD research directions

My doctoral goal is to scale *structure-respecting learning* beyond hand-specified groups and geometry, toward learned symmetries and semantic/physical plausibility constraints.

- **Discovering approximate symmetries from spatiotemporal signals.** Learn approximately equivariant transformation families from video and temporal coherence (e.g., motion-induced deformations), and enforce them with SER-style intermediate-layer constraints, leveraging small inter-frame transitions as a local equivariance signal without requiring analytic group actions.
- **Diagnosing and mitigating invariance–equivariance interference via subspace decoupling.** I will study when and why invariant SSL and equivariant objectives compete, tracking effective-rank dynamics, layerwise gradient alignment, and latent-orbit organization under group actions, and develop subspace-structured objectives (e.g., null-space-projected equivariance) that preserve invariant embeddings while encoding transformation-sensitive variation.
- **Semantic/physical constraints for trustworthy generative models.** Extract candidate constraints (anatomy, part consistency, numeracy, logical relations) from pretrained LLMs/VLMs, ground them with lightweight scoring modules, and enforce them with scalable alignment objectives.