

RESEARCH NOTE

Hello, I'm Bishal

What four years of experiments across images and time series taught me about robust learning—and what I'll share here.

July 14, 2026 · 2 minute read




Hello, I'm **Bishal** (ବିଶାଳ, 비살, 丕薩), a computer vision researcher interested in one question: **how can a learning system remain reliable when its world changes?**

My path began with computer science in India, followed by research at the National Atmospheric Research Laboratory and the Indian Institute of Science. I moved to South Korea as a Global Korea Scholarship recipient and began my PhD under Prof. Jaepil Ko at Kumoh National Institute of Technology.


Four years, one recurring question

For four years of my PhD journey, I worked from morning into night, every day—including weekends—without taking a day off. I do not offer that as a productivity formula. It explains the depth behind these posts: reproducing baselines, questioning labels, designing ablations, tracing unstable runs, and restarting when an idea failed.


The work ranged from medical images and steel microstructures to industrial inspection, weather radar, and general time series. It led to research at MIDL, MICCAI, WACV, CVPR, and ICANN—but the most reusable lessons often came from experiments that never entered a paper.



Medical imaging



Materials microscopy



Radar and time series

Across those domains, the pattern stayed remarkably consistent:

DATA

Images + time series



REAL-WORLD FRICTION

Shift + scarcity +
ambiguity



RESEARCH DIRECTION

Context + memory +
feedback

This is why I study **learning under distribution shift**, including memory-augmented models inspired by hippocampal–cortical interaction. Whether the input is a histology image or an evolving signal, I keep returning to the same question: *what should a model remember, and how should that memory change its next decision?*

What readers will get

- **Decisions, not just results:** how I choose baselines, structure ablations, and detect misleading gains.
- **Failures worth keeping:** what broke, why I think it broke, and what changed next.
- **Robust-learning insights:** practical lessons on domain shift, associative memory, liquid neural networks, segmentation, and foundation-model adaptation.
- **Cross-domain judgment:** what transfers between medical, materials, industrial, and temporal data—and what does not.

I will separate measured results from interpretation, show assumptions and negative evidence, and explain evaluation beyond one average score. The aim is simple: help readers design a better next experiment.

Beyond the experiments

I currently serve as **Vice President of the MICCAI Student Board**, following a year as Social Events Officer. I also contribute through conference reviewing—including recognition as an ICML 2026 Gold Reviewer—and community work in Korea. Reaching TOPIK Level 5 lets me share useful posts in both English and Korean.

My [Research & Projects](#), [Academics](#), and [Co-curricular Activities](#) pages hold the complete record. This blog is where I will explain the reasoning behind it.

If your experiment behaves differently from what the paper promised, you are in the right place.