ISPRS CATCON 9 — Judges' Summary

Project

Title. Remote Sensing Teaching and Practice Platform for Flood Analysis and Student Interaction
One-sentence pitch. A web-based teaching environment built around a real Sentinel-1 SAR flood case — every figure is a real model output, every pipeline step is reproducible, and students close the loop with a guided inquiry form.
Category. Web information package · education-oriented (CATCON Track C).
Public URL. https://main.remote-sensing-platform.pages.dev
Source / reproducibility. Private repository; source available to judges on request.

Team

Role	Name	Affiliation	Country	Email
Principal Investigator · Team Lead	Wei Yuan	IRIDeS, Tohoku University	Japan	wei.yuan@tohoku.ac.jp
Ph.D. Candidate · Model & Data Engineering	Zhongyuan Yang	IRIDeS, Tohoku University	Japan	yang.zhongyuan.t2@dc.tohoku.ac.jp
Ph.D. Candidate · AI Methodology	Weihang Ran	OSCARS Lab, The University of Tokyo	Japan	ran-weihang@g.ecc.u-tokyo.ac.jp

Primary contact: Wei Yuan (`wei.yuan@tohoku.ac.jp`).

Design concept and purpose

Traditional remote sensing teaching shows finished maps and finished equations. This platform teaches the **workflow** — the messy middle where raw satellite archives become decisions. We pick one real disaster (Banda Aceh flood, 2025-11-26), one operational SAR pipeline, and one clearly-explained segmentation model (UNet-RSMamba), and we walk learners through it from L1C GRD to interpreted flood map without hiding the steps.

Functions

1. **SAR physics primer** — VV/VH/RGB imagery paired with backscatter intuition.
2. **End-to-end flood case** — Sentinel-1 GRD to Cross-Scale Mamba (CS-Mamba, 40.55 M params) prediction, side-by-side across four input-clamp configs.
3. **Walkable 6-stage pipeline** (`/pipeline`) — real commands, real artefacts, TL;DR and check-yourself prompts per stage.
4. **Explorable-explanation notebook** (`/paper`) — five interactive widgets replace the essay: tile explorer, clamp playground, architecture walkthrough, scrubbable **37-epoch** training replay (best @ ep 12, test mIoU 79.79 %), and model race. Every teaching point is a control the learner operates.
5. **Live GPU inference** — a ● `LIVE` panel runs a real ~400 ms forward pass against a WSL-hosted FastAPI endpoint and returns the 3-class prediction mask.
6. **Student inquiry loop** — feedback form backed by Cloudflare Workers + D1.
7. **Reproducible artefacts** — every figure is a real output; pipeline reruns with one command.

System requirements

Viewing (judges, learners): any modern browser (Chrome 110+, Firefox 110+, Safari 16+, Edge 110+); responsive 360 px–4K; JavaScript required; no sign-in.
Rebuilding (instructors, reproducers): Node 18+ with Nuxt 3 / Vue 3 / Tailwind; Cloudflare Workers + D1 for the inquiry API; Python 3.10 + PyTorch 2.x + SNAP 9 + rasterio for the modelling pipeline; Cloudflare Pages for static deploy.

Originality

- The segmentation model is **original research**: Cross-Scale Mamba (CS-Mamba), our SAR-specific extension of RSMamba (Chen et al., 2024), submitted to the ISPRS 2026 Congress — the scholarly basis of this CATCON entry.
- Built around **one real disaster** (Banda Aceh, 2025-11-26), not a catalogue of toy datasets.
- Pedagogy is **widget-first**: the `/paper` route carries no prose — just five controls the learner operates.
- Parameter sensitivity is a **first-class lesson**: same model, four input-clamp configs, four visibly different flood maps.
- **Every figure is a real artefact** — no stock images, no staged screenshots.
- The teaching loop closes with an **inquiry form backed by a real database**, not a mailto link.

Target learners

- Undergraduates: first contact with SAR. Suggested path: Home, Primer, Flood Case.
- Graduates: reproducibility and parameter sensitivity. Suggested path: Pipeline page, Training replay, then run `reproduce-no-sar`.
- Instructors: a self-contained case for a one-hour lecture. Suggested path: Flood Case + teaching-flow prompts.