Segmentation of SAR Imagery of River Ice in the St. Lawrence River Using Deep Learning

Liubov Poliakova

Supervisor: Andrea K. Scott • lpoliako@uwaterloo.ca

The problem

River ice drives **winter flooding** through ice jams, threatens **hydro-power intakes, bridges and navigation**, and its timing is a sensitive **climate indicator**.

Optical satellites fail exactly when we need them: polar winters are dark and cloudy for weeks.

Synthetic-aperture radar (SAR) sees through. Sentinel-1's C-band radar is an active microwave sensor, it images day or night and is not sensitive to cloud cover or snowfall.

The data

Study area — St. Lawrence River, upstream (SW) of Montréal, QC

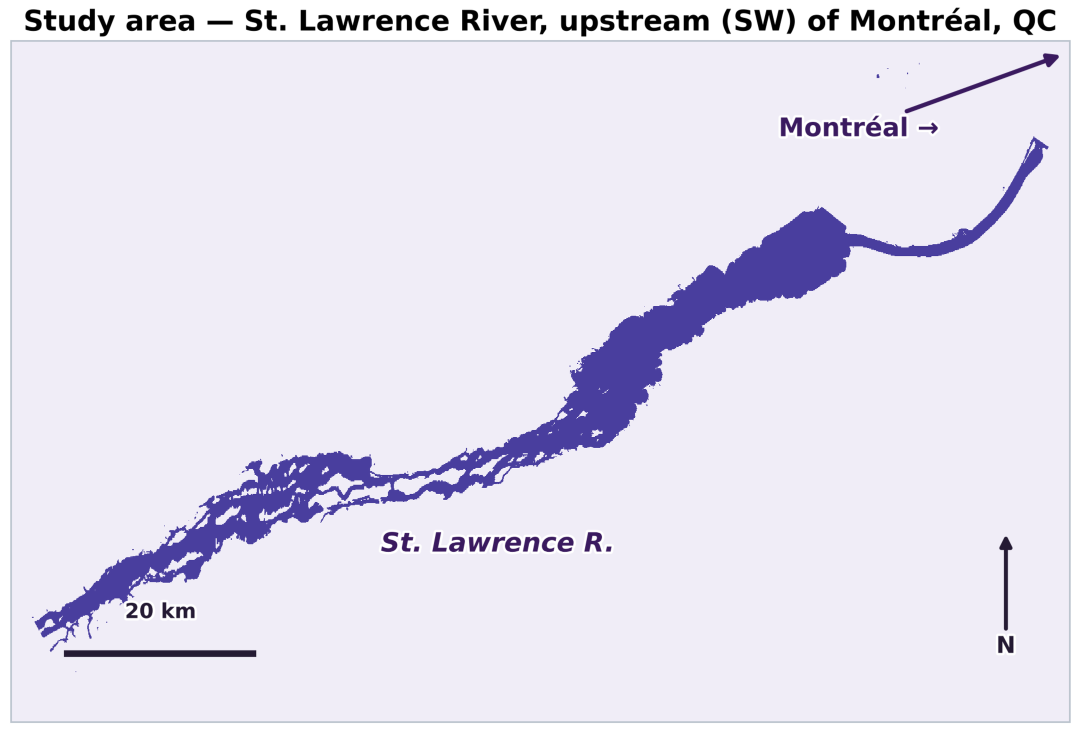

Study area: a ∼110 km reach of the St. Lawrence River just upstream (SW) of Montréal, QC.

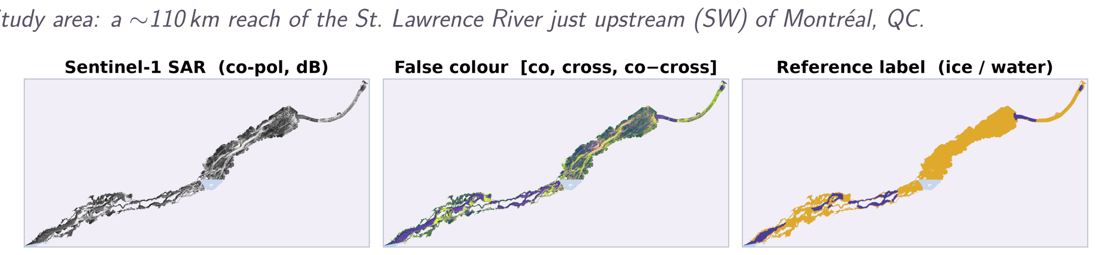

A single Sentinel-1 scene over the St. Lawrence: radar backscatter (left), the three physical input channels as false colour (middle), and the Sentinel-2–derived reference labels used for training (right).

- **60 image–mask pairs**, ten winters (**2017–2026**), St. Lawrence near Montréal, QC (UTM 18N, 10 m).
- **Input:** Sentinel-1 GRD, dual-pol (VV/VH), calibrated to γ^0 backscatter in dB.
- **Labels:** ice / water **hand-labelled** in Segments.ai from cloud-free Sentinel-2 optical imagery, date-matched to each radar pass.
- Everything not confidently ice or water (nodata, fill, ambiguous) is **ignored** in both loss and metrics.

Method

U-Net (ResNet encoder, **ImageNet-pretrained**) predicts a per-pixel ice probability from three physically-motivated SAR channels.

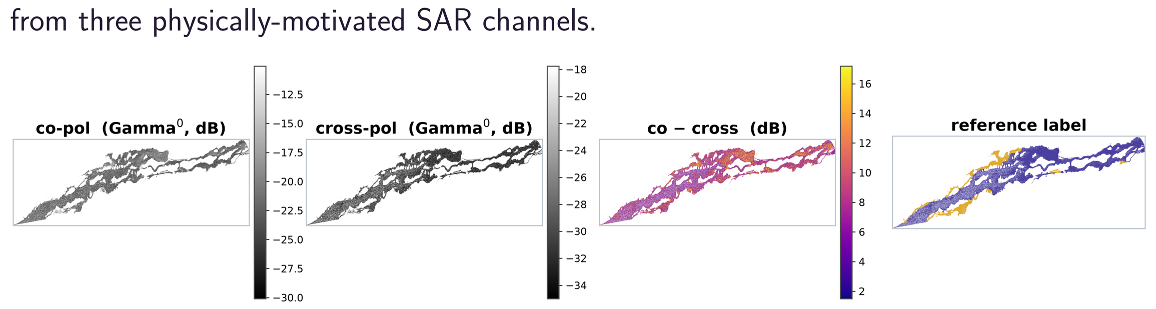

The 3-channel input: co-pol dB, cross-pol dB, and their difference (co−cross).

Sentinel-1 GRD (VV,VH) → [co, cross, co−cross] dB → U-Net ResNet-18 → ice / water mask

- **Loss:** masked **BCE + Dice**.
- **Training:** 256–320 px patches, geometric + radar-speckle augmentation, early-stop on validation IoU.
- **Tuning: Bayesian search (Optuna TPE)** over lr, encoder, loss weights, patch size and augmentation, scored by **leakage-safe K-fold** scene cross-validation.

Leakage-safe splits

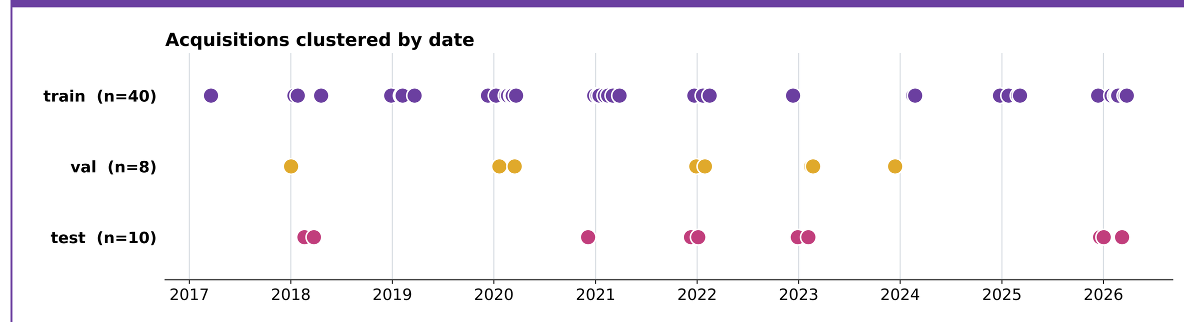

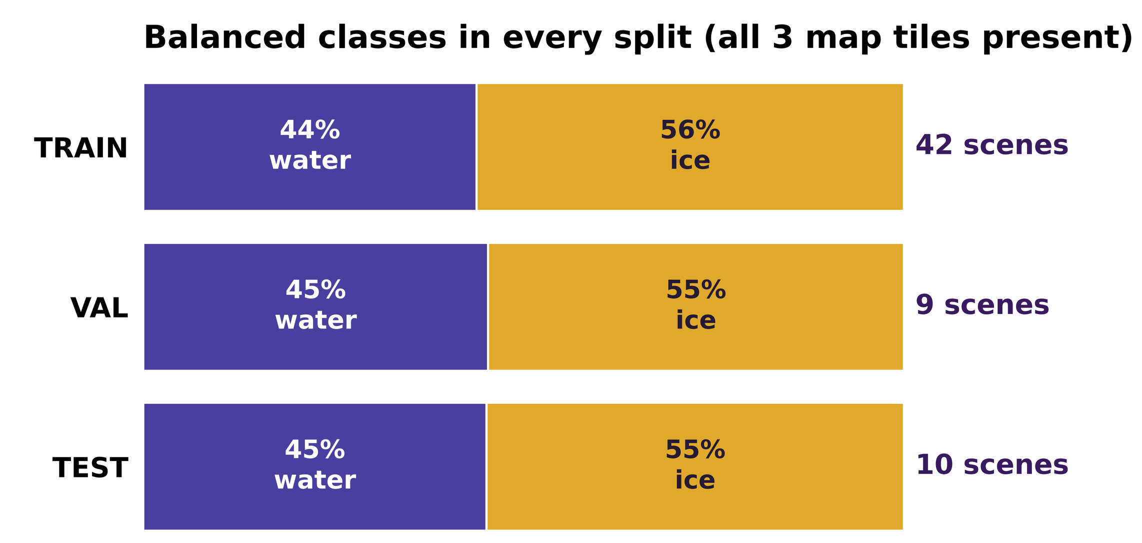

Acquisitions are clustered by date

and each cluster is assigned *whole* to one split, forcing a ≥5-day gap between train, validation and test so **no near-repeat pass leaks** across the boundary. Hard constraints keep every map tile and **both classes (≥12%)** present in each split.

How well does it work?

0.88 test mIoU (ice & water)

0.86 validation mIoU

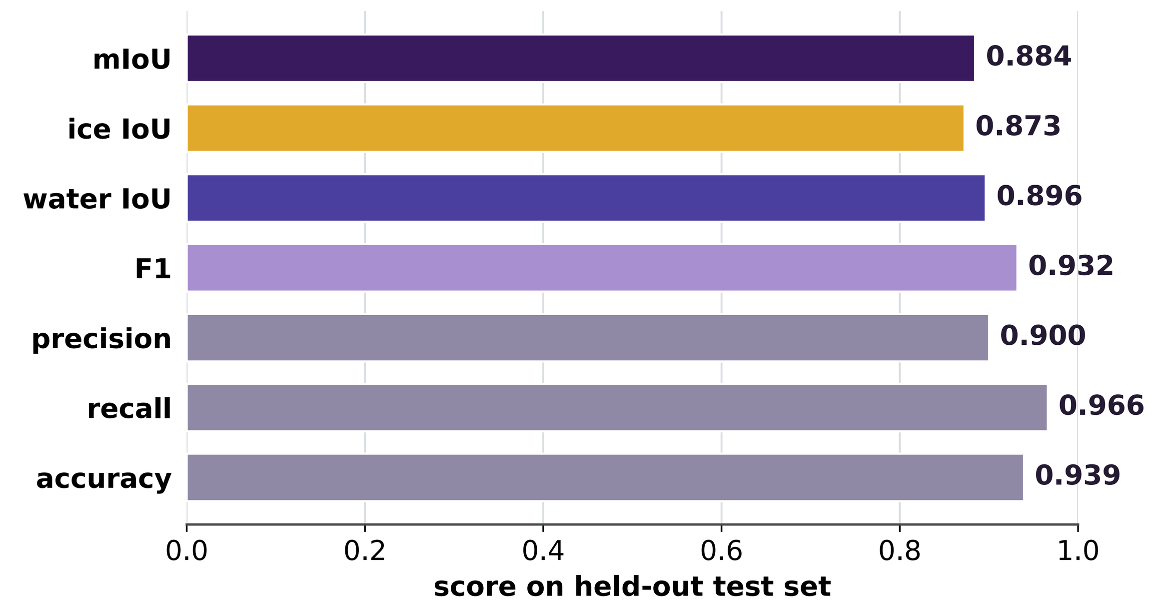

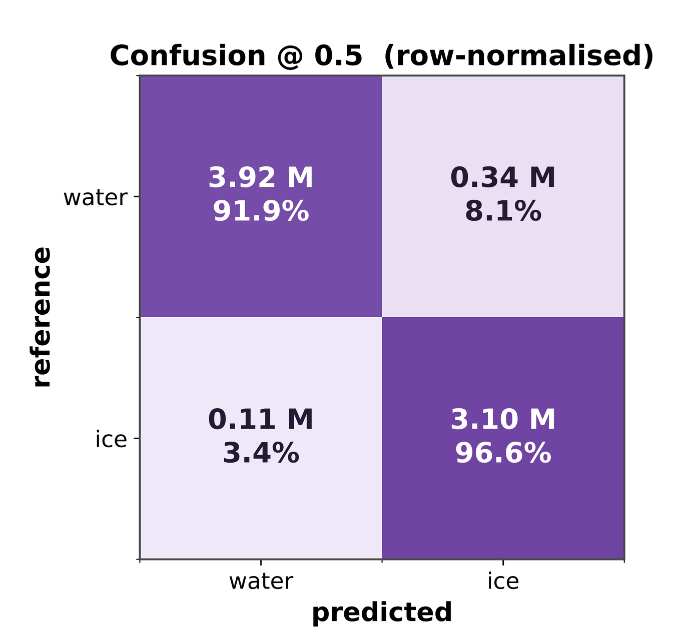

*Headline model on the held-out random-split test (8 substantive scenes, ∼7.5 M labelled px). Balanced **ice 0.87 / water 0.90**; **recall 0.97** / precision 0.90.*

When it works: a full scene

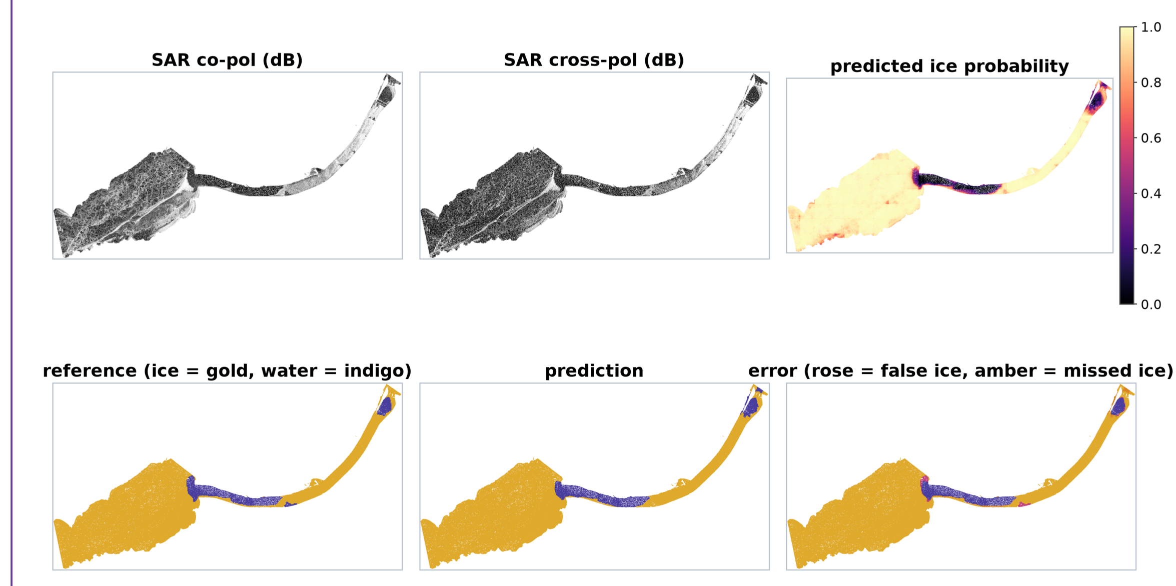

*Held-out scene (3 Feb 2021): a frozen basin threaded by an open-water channel. Top: co/cross-pol radar and predicted ice probability; bottom: reference, prediction, and error map (rose = false ice, amber = missed ice). The model tracks the water channel through the ice from radar texture alone - **scene mIoU 0.89** (ice 0.98 / water 0.81).*

When it fails: wet ice

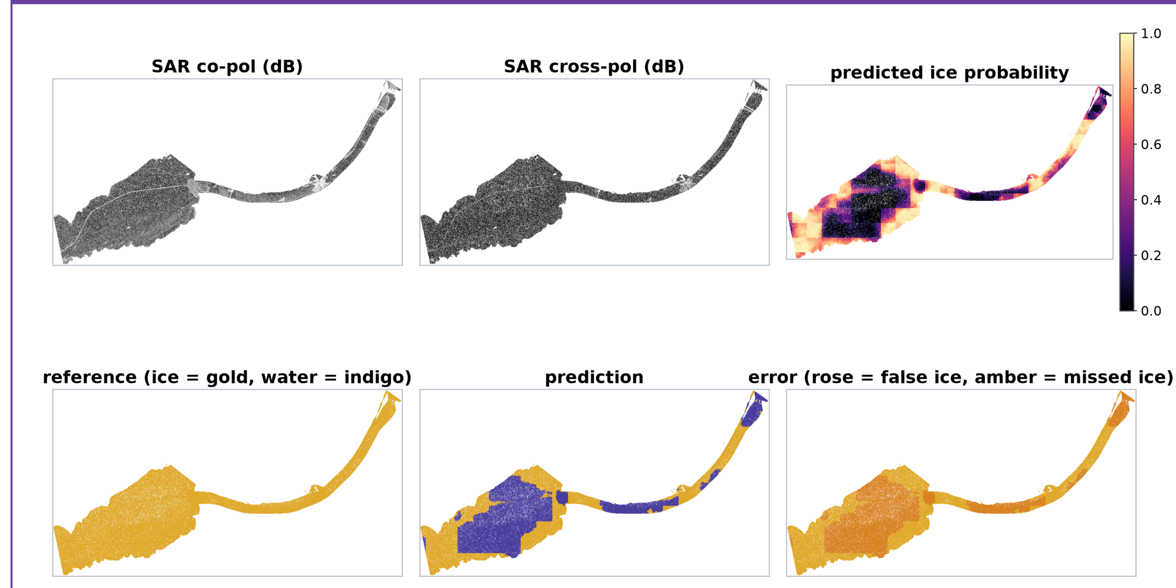

*Spring thaw, 22 Mar 2019. The basin is still fully frozen (reference = all ice), but the ice is **wet**: meltwater and wet snow sit on top. Wet ice is smooth and radar-dark — it back-scatters **like calm open water** — so the model reads large frozen areas as water (amber = missed ice). **Scene mIoU 0.24** (ice IoU 0.48). Separating wet ice from water is a known C-band SAR limitation and our **main failure mode**.*

The hard case: a new winter

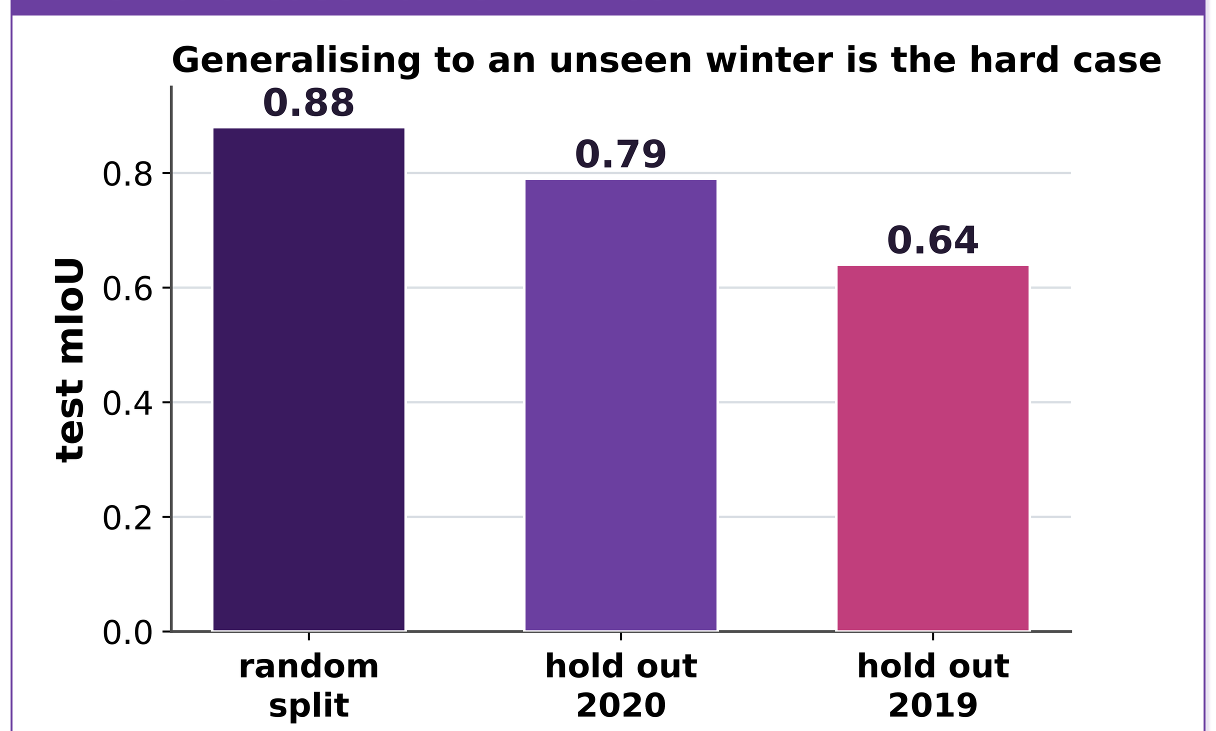

*Holding out an entire winter (leave-one-year-out) is harder than a random split and varies year to year (**0.64–0.79 mIoU**): **2019's wet-ice spring** is the hard case (0.64), while 2020 reaches 0.79. Inter-annual ice regimes differ and a few labelled winters do not yet cover them.*

Does weather data improve ice classification when the surface is wet?

A wet surface on the ice or snow changes the dielectric properties such that the backscatter is significantly attenuated and looks more like water, leading to misclassification. We attempted to improve this by bringing in additional features.

- **Distance-to-shore:** a static 4th channel — each water pixel's metres to the nearest shore, from a 10 m land/water mask (ice forms in consistent places).
- **Weather** via **FiLM**: per-scene temperature, **freezing-degree-days** and snow depth & snowfall (ERA5) — the accumulated cold the radar can't see. A small network turns that weather vector into a per-channel scale γ and shift β that re-tune the U-Net's bottleneck, so cold context reshapes the features without touching the radar pixels.

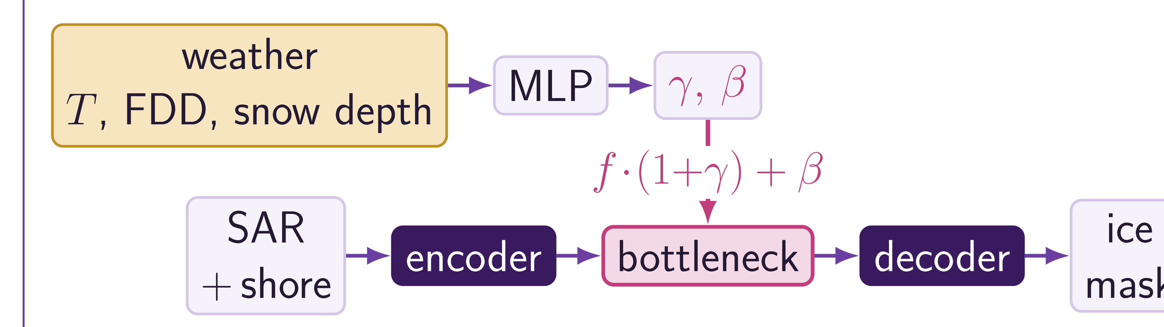

Results

On the wet-ice scene the added context recovers the frozen basin the SAR-only model gives up on: **ice recall 0.48 → 0.95**.

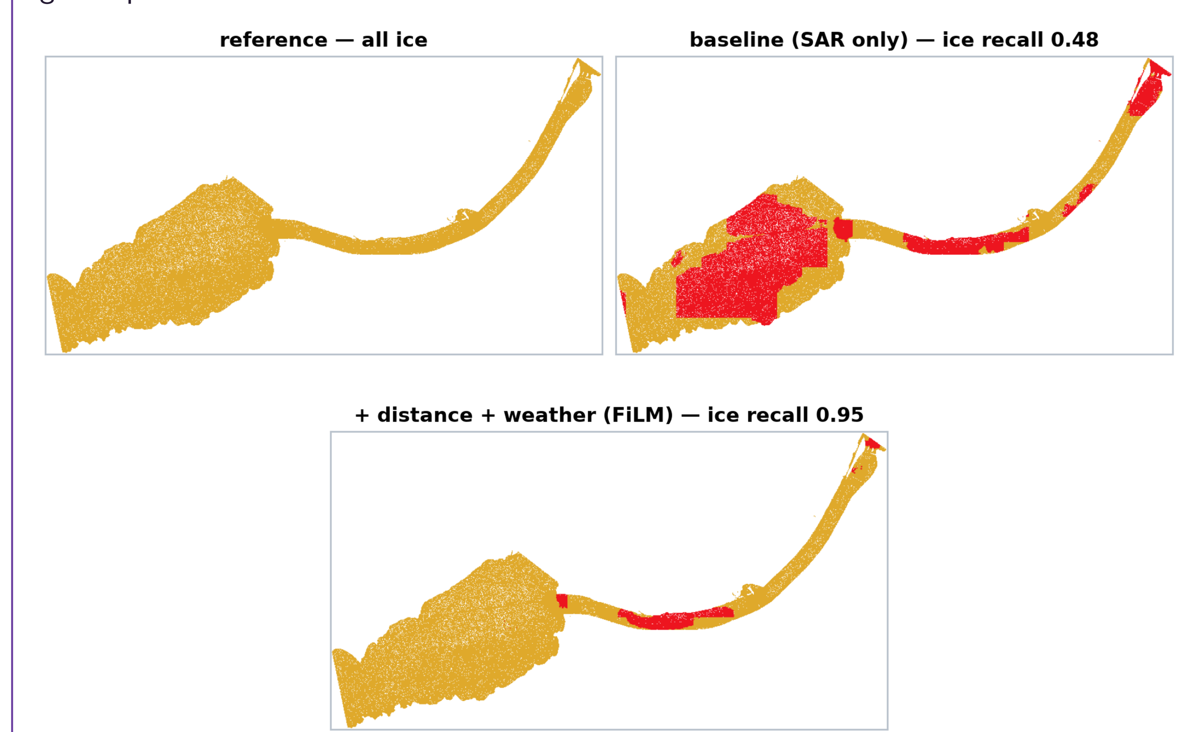

Same wet-ice scene (22 Mar 2019). Red = ice the model missed (read as open water); gold = ice, indigo = water. The baseline loses the whole radar-dark basin; the FiLM model keeps it as ice.

What's next

The model is **data-limited**: 60 labelled scenes across ten winters barely cover the hard regimes (wet ice, unseen winters). The next gains come from *more and harder* data, not a bigger network.

- **More labelled data.** Auto-pair Sentinel-1 passes with **MODIS** (daily) and **Landsat 8/9** optical, not just Sentinel-2, to grow the training set.
- **Richer auxiliary data.** Add physically-grounded context beyond distance-to-shore and weather: operational **ice charts**, river **discharge / stage**, upstream freezing-degree-days, and DEM / flow direction.
- **Multi-sensor fusion.** We plan to fuse Sentinel-1 C-band with **Ka-band SWOT** (KaRIn).

References

[1] Scott, K. A. & Poliakova, L. (2026). *Segmentation of SAR Imagery of the St. Lawrence River Using a U-Net and Preliminary Steps for Northern Lakes.* Technical report, University of Waterloo / National Research Council of Canada.

[2] Gallet, M., Atto, A., Karbou, F. & Trouvé, E. (2024). Wet Snow Detection From Satellite SAR Images by Machine Learning With Physical Snowpack Model Labeling. *IEEE J. Sel. Topics Appl. Earth Obs. Remote Sens.*, 17, 2901–2917.

Funding. This research was undertaken thanks to funding from the **Natural Sciences and Engineering Research Council of Canada (NSERC)** through the CREATE–QCS training program, and from the **National Research Council of Canada (NRC)**.

Data & tools. Copernicus Sentinel-1/-2 (ESA); ESA SNAP; PyTorch, segmentation-models-pytorch, Optuna, MLflow.

Contact: lpoliako@uwaterloo.ca

University of Waterloo • NSERC CREATE–QCS Summer School 2026