HydroSphereAI

March 4th, 2025

Andre R. Erler (aerler@aquanty.com)
Steven K. Frey (sfrey@aquanty.com)
Brayden McNeill (bmcneill@aquanty.com)

About Aquanty Inc.

HydroGeoSphere™

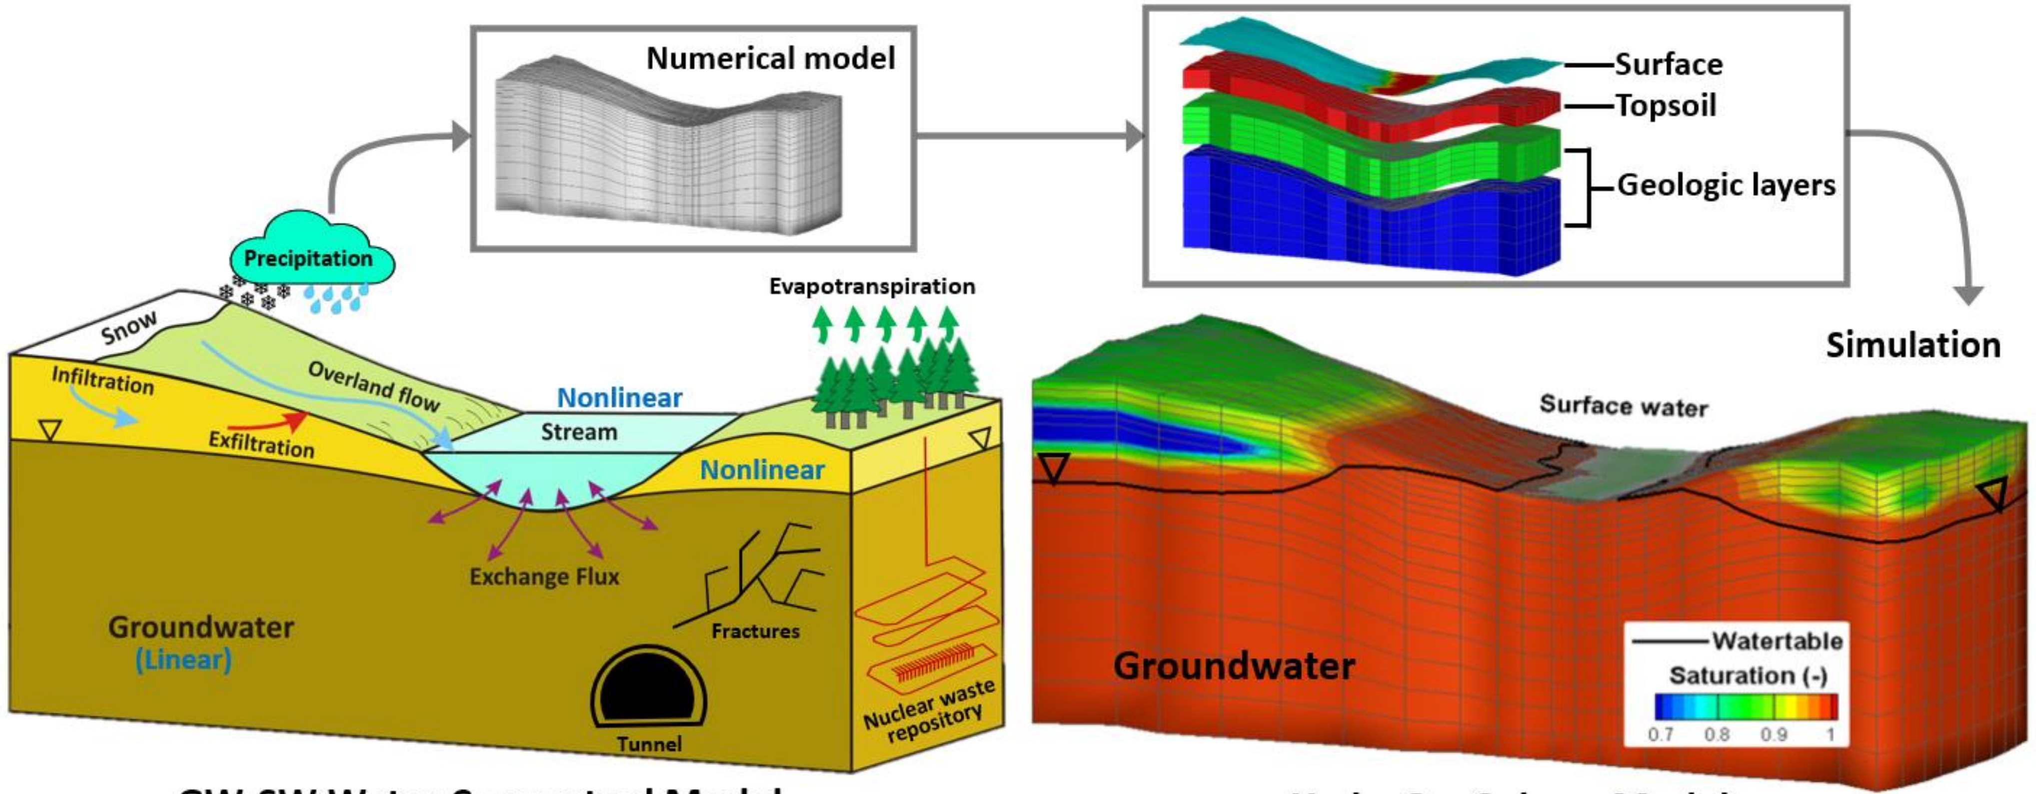

GW-SW Water Conceptual Model

HydroGeoSphere Model

HydroClimateSight™

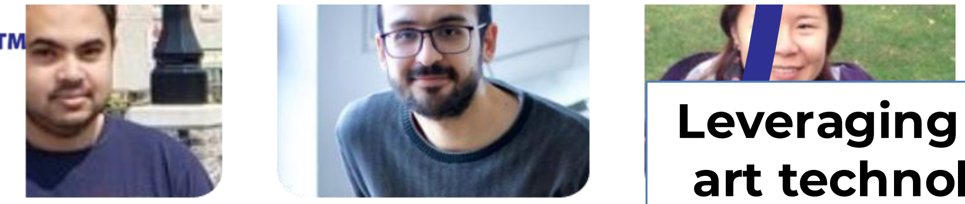

Leveraging state-of-the-art technology for real-time hydrologic forecasting

HydroSphereAI™

HGS REAL TIME™

CANADA ONE WATER

Modelling On Demand™

Aquanty Products and Services

HydroGeoSphere

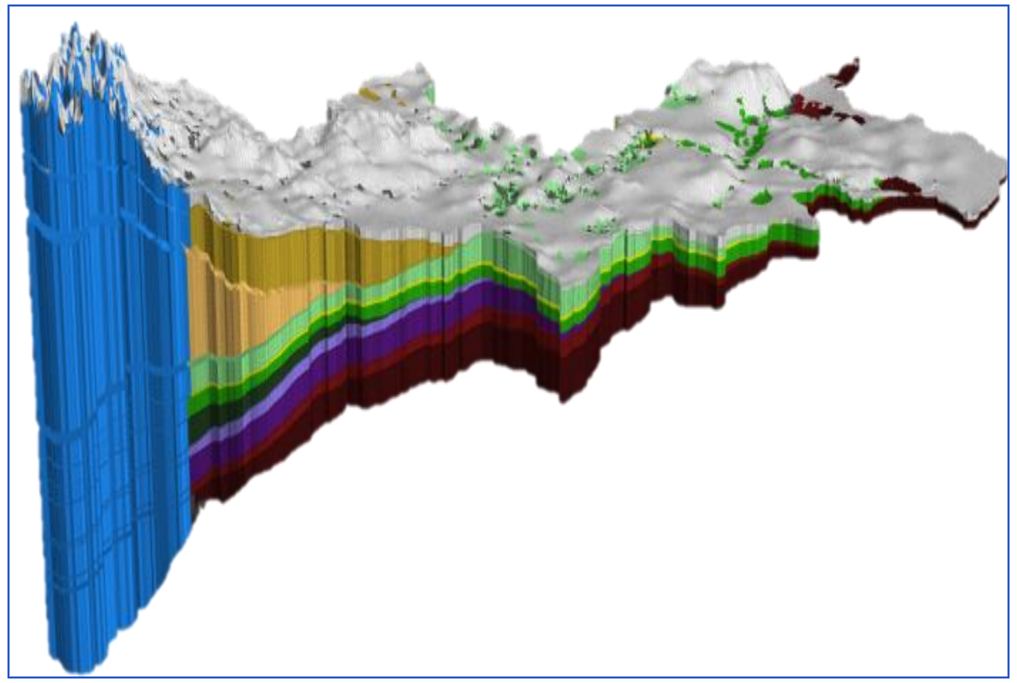

- Fully integrated, hydrological model
- Entirely physics-based
- Used worldwide and across Canada

HydroClimateSight

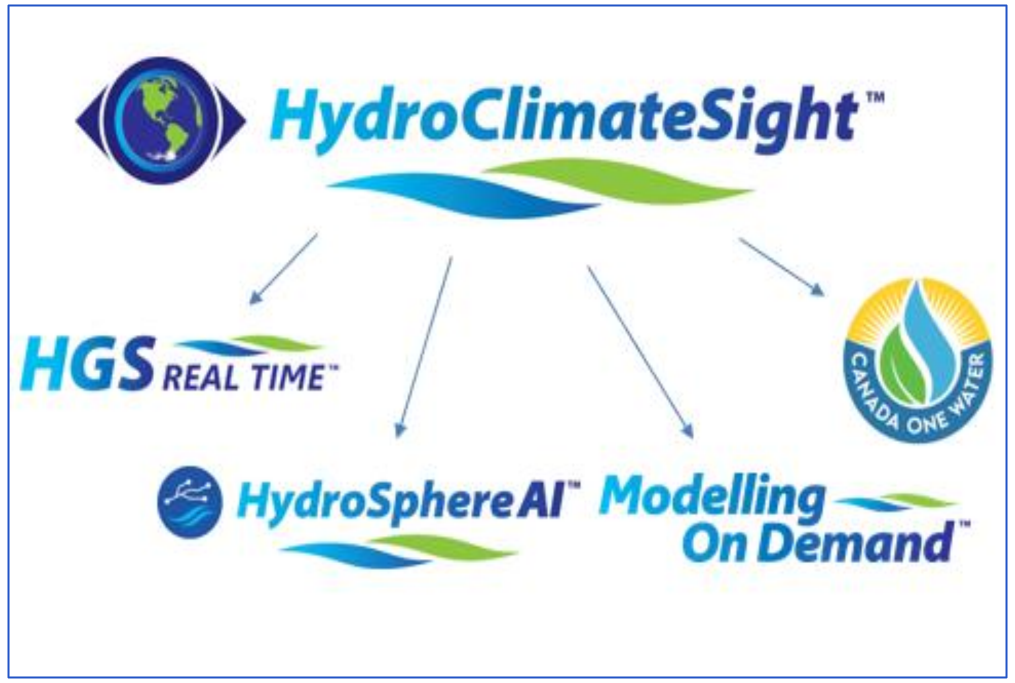

- Web-based SaaS platform since 2018
- Operational forecasting services
- Including HydroSphereAI

Canada1Water

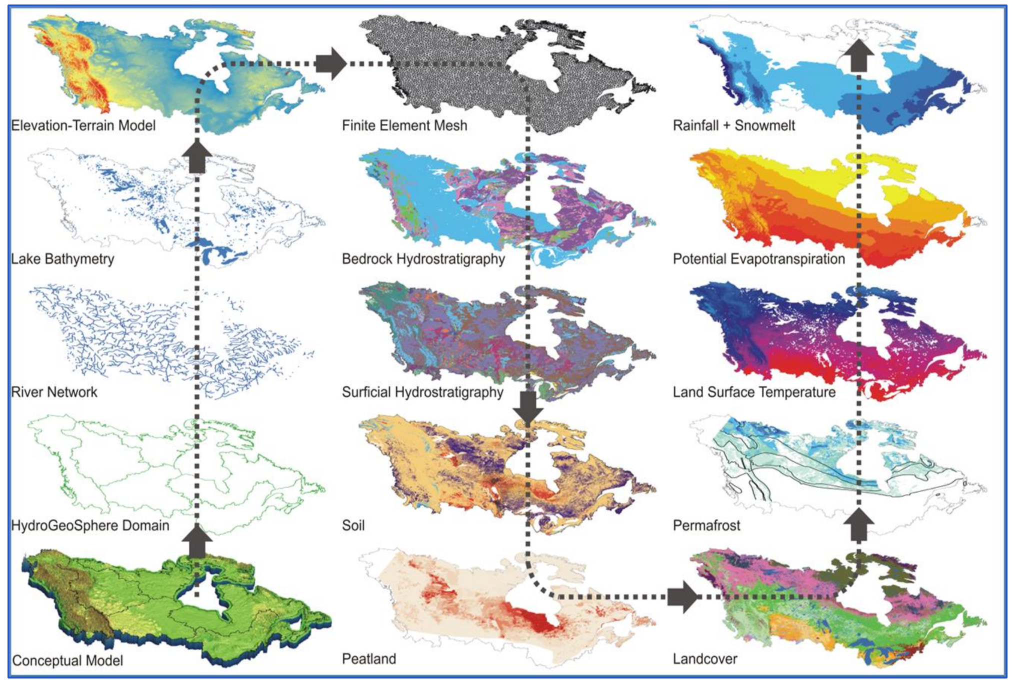

- Comprehensive modelling framework at continental scale
- Federally funded
- High resolution datasets available today

HydroSphereAI: Product Overview

Machine Learning-based hydrologic forecasts across Canada and the US

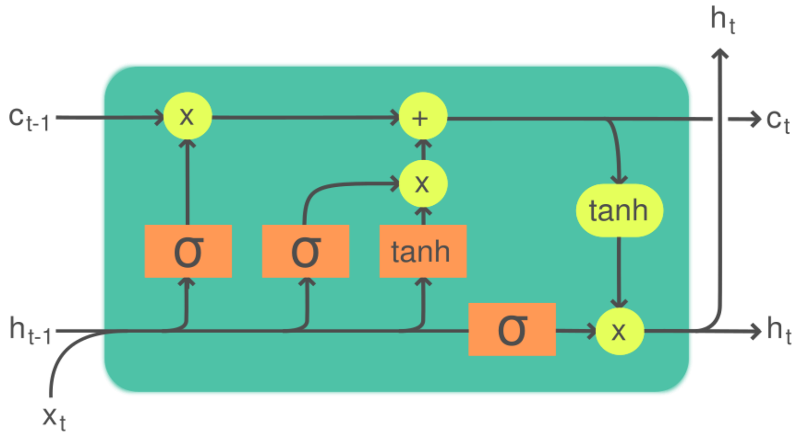

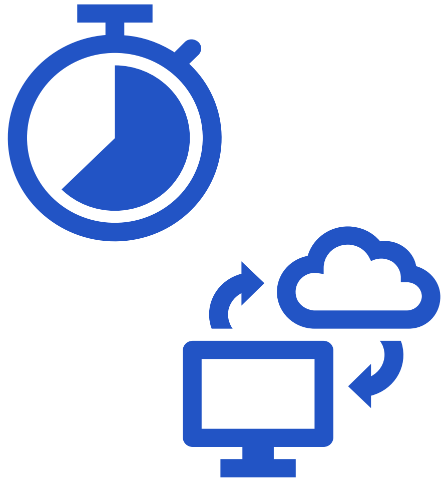

Continental coverage

- ~2k stations active
- Easily scalable to US and other regions

Cutting edge algorithms

- LSTMs
- Fine tuning
- Custom end-points

Forecast integration

- GEPS, HRDPS, SnoDAS, CaPA…
- Daily, seasonal, hourly

Operational platform

- Web app
- APIs
- Robust & Automated

Some History, sort of…

- Training ANNs or Random Forests as Rainfall-Runoff models is actually an old idea and has been tried before

- However, trained on a single watershed, ML models do not perform better than existing Rainfall-runoff models

- They either over-fit the training data or have poor skill

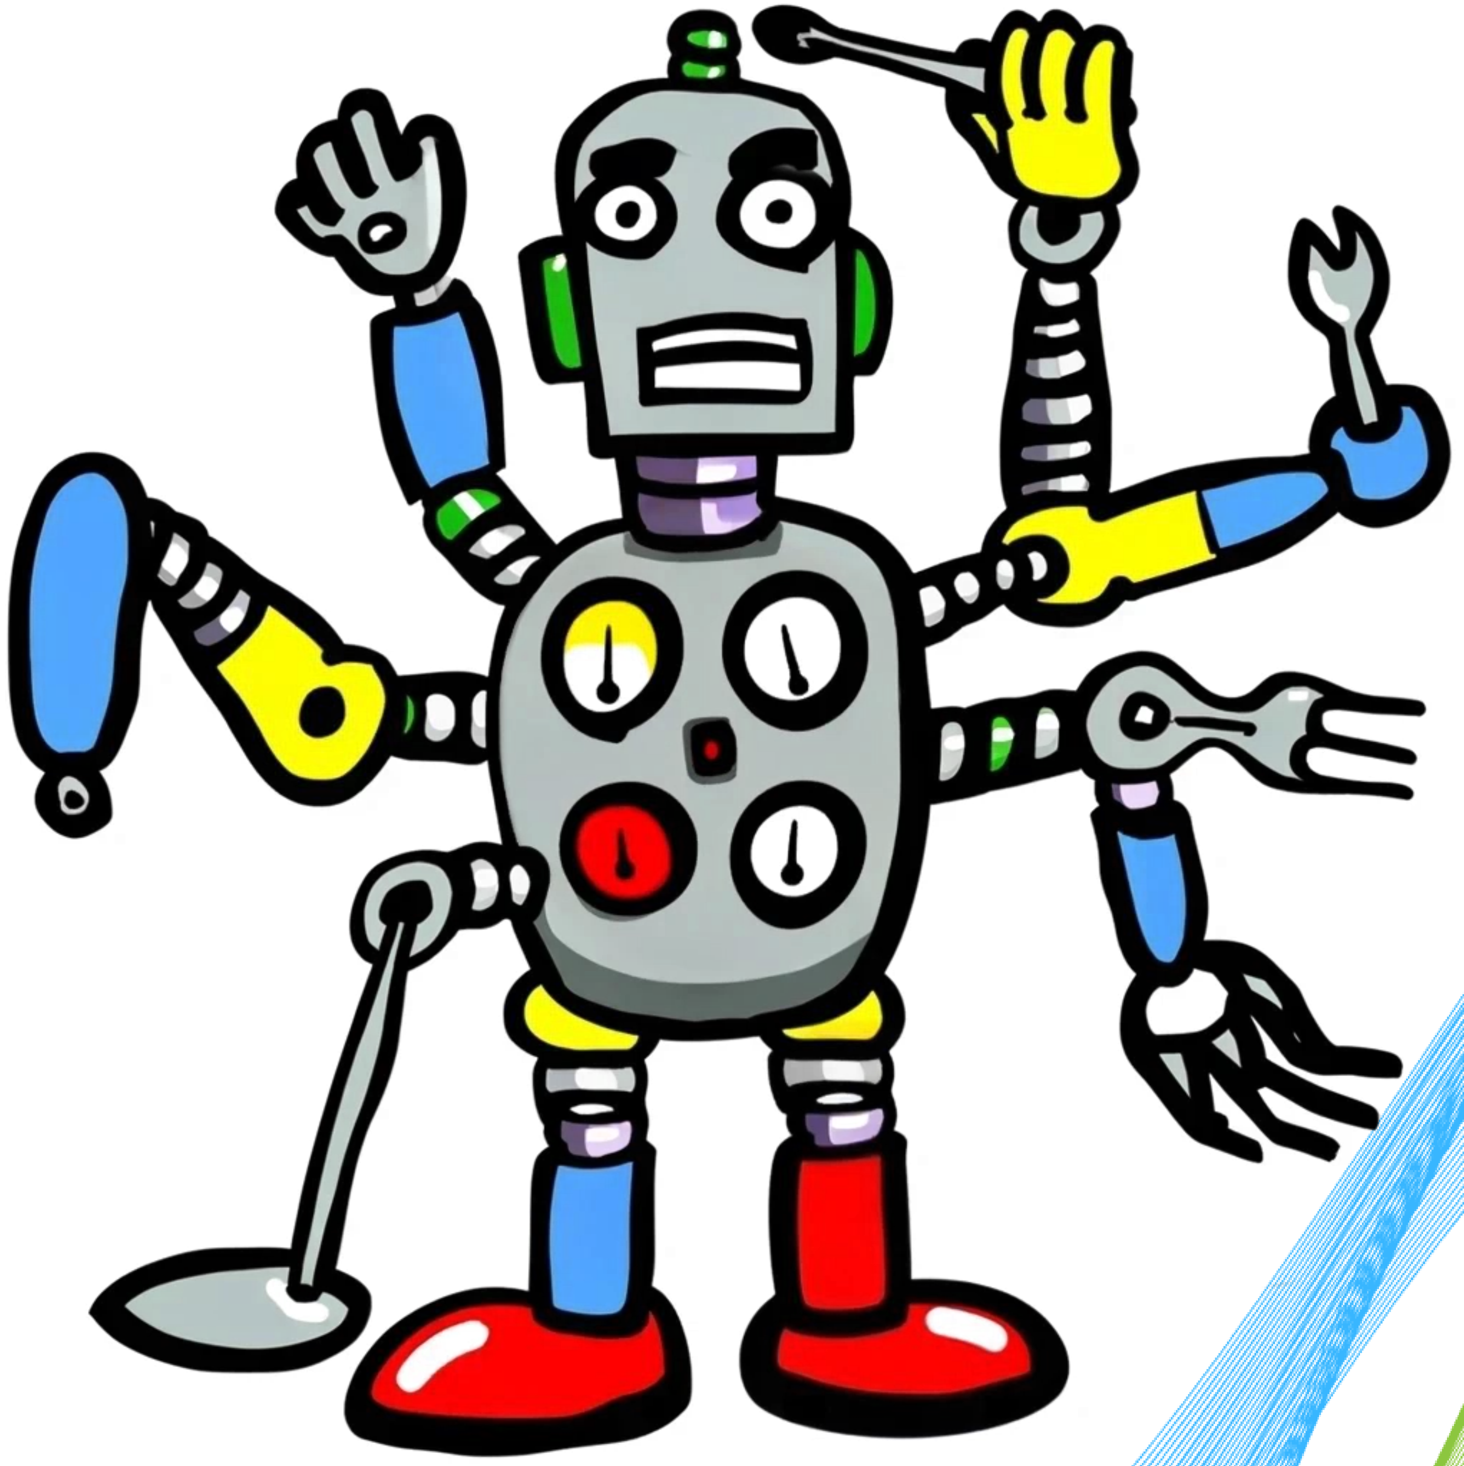

LSTMs and Large Datasets

- Based Kratzert et al. (2018) and follow-on work (> 1000 citations)

Two key advances:

- Use of LSTMs, utilizing memory cell
- Trained on a very large set of watersheds, using geographic attributes as predictors

LSTM Schematic

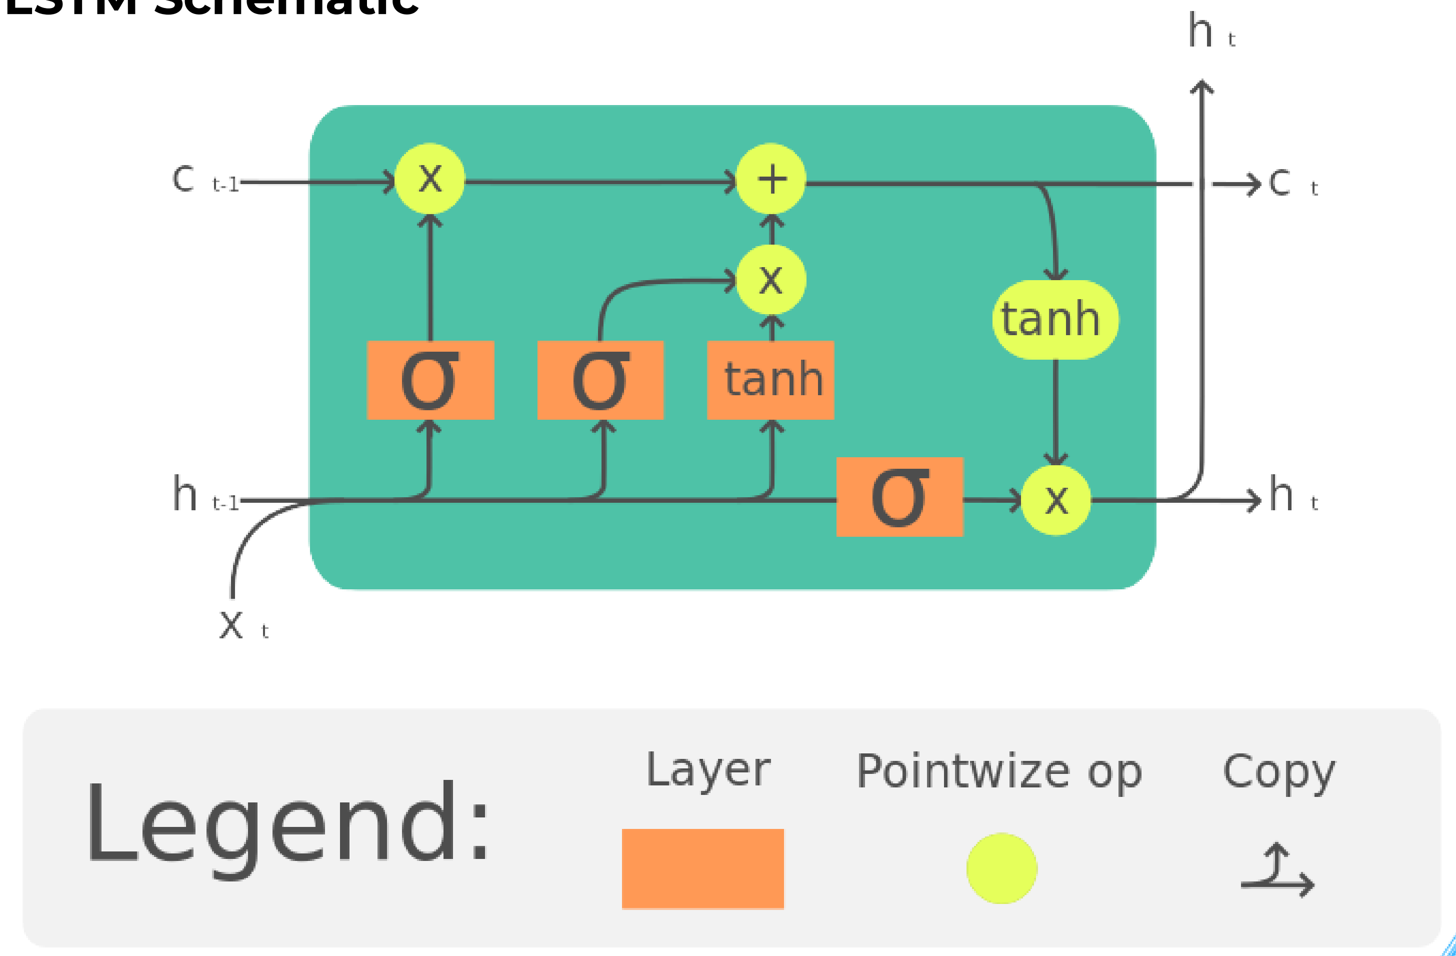

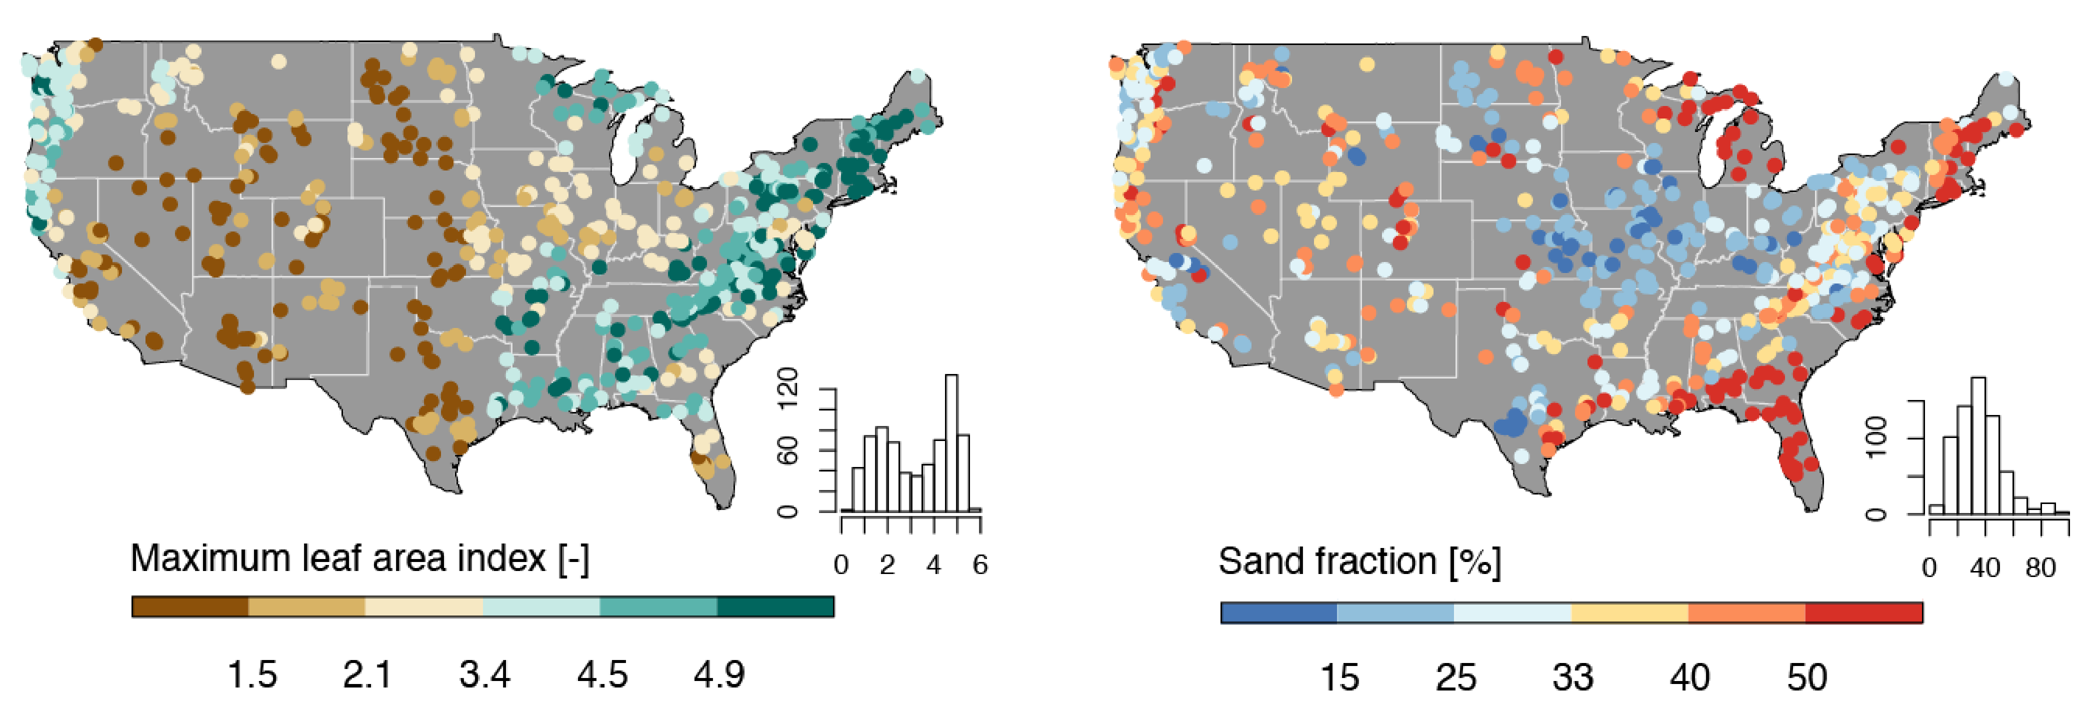

Excerpt of watershed features from the CAMELS dataset (US)

- Academic focus on historical evaluation and interpretability

Why LSTMs?

- LSTMs have an internal memory cell, which can represent state variables
- They learn to track the physical state of the watershed, quite similar to physical models
  - e.g. snowpack or groundwater table
  - This means we also need a long spin-up, like physical models (~1 year)

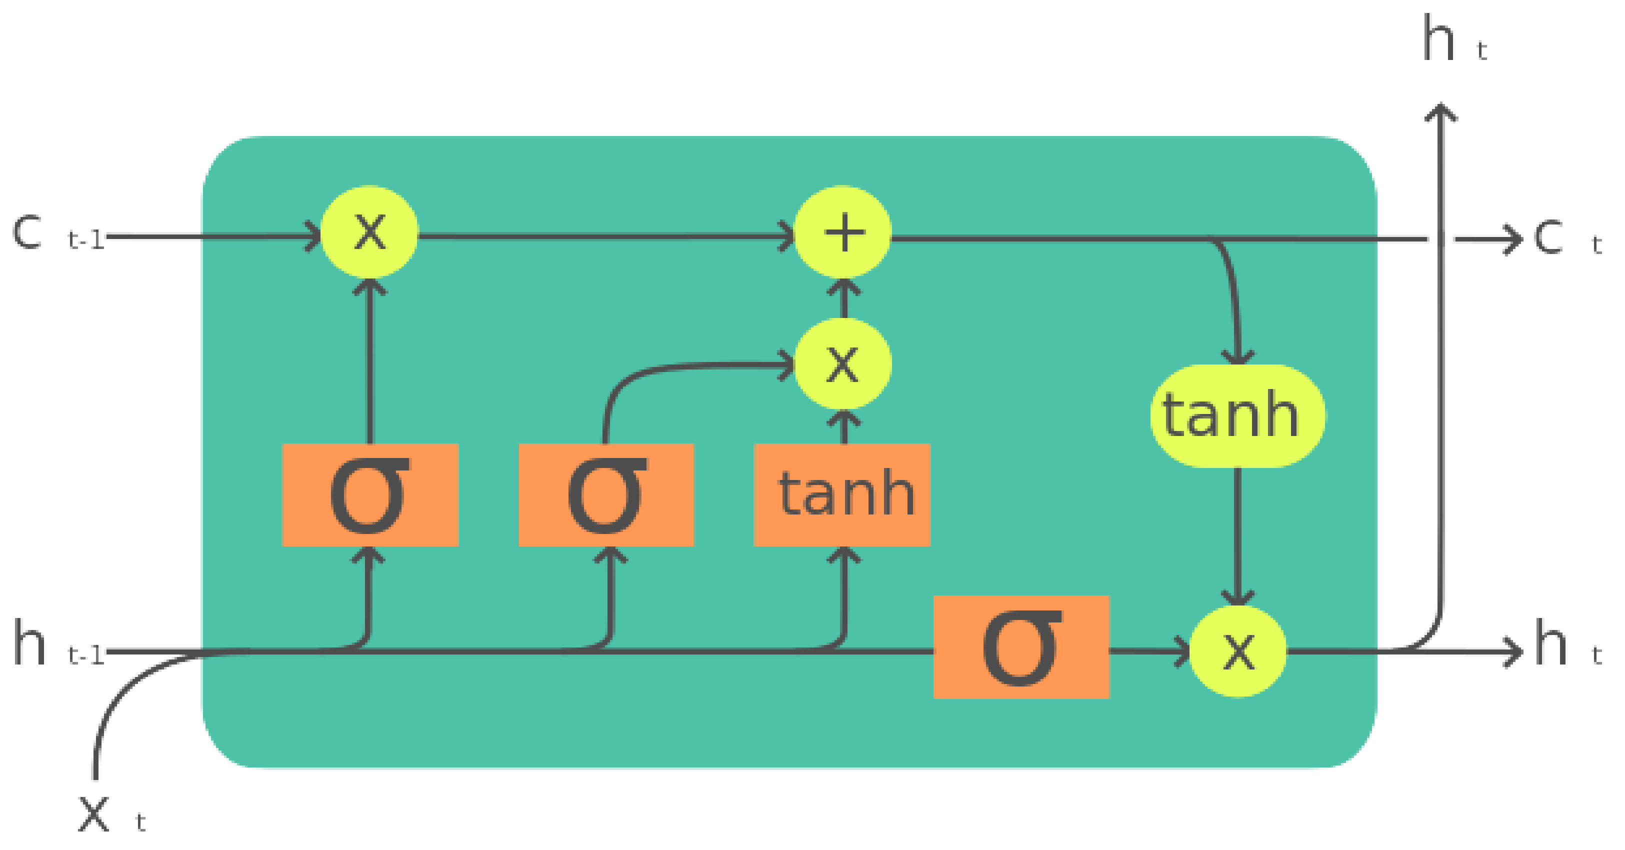

General architecture of an Long- Short-Term Memory (LSTM) network; "C" is the memory cell state. (Guillaume Chevalier - CC BY-SA 4.0)

N.B.: Originally LSTMs come out of natural language processing, but are frequently used for timeseries prediction

The Training Data

- CAMELS and HYSETS provide geographic attributes and meteorology for many catchments in North America
    - In Canada mainly based on HYSETS and ERA5

- We added **our own custom data**, based on:
    - CaPA & HRDPS
    - AgERA5
    - HyDat
    - **[Canada1Water]** geographic data

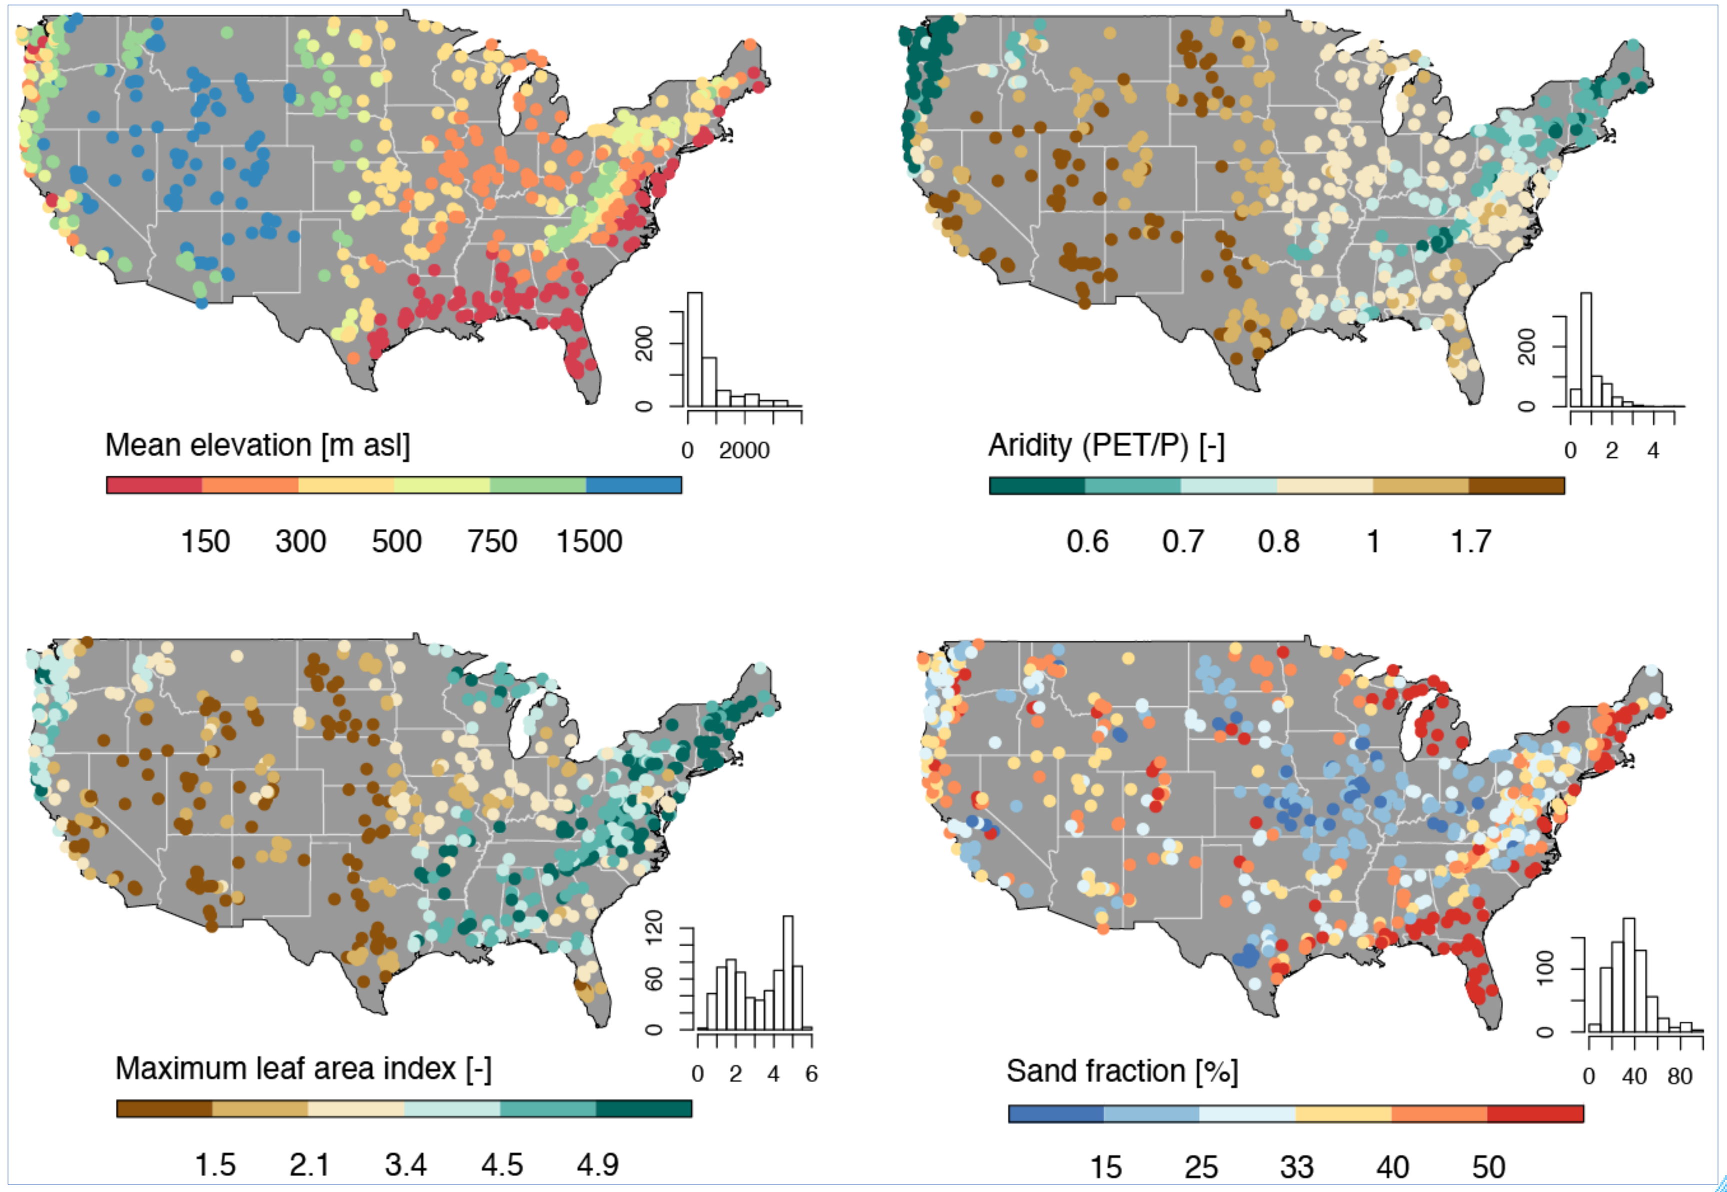

https://ncar.github.io/hydrology/datasets/CAMELS_attributes

- Note that all inputs for the ML models are catchment-averaged: **essentially an HRU approach**

Model Inputs & Outputs

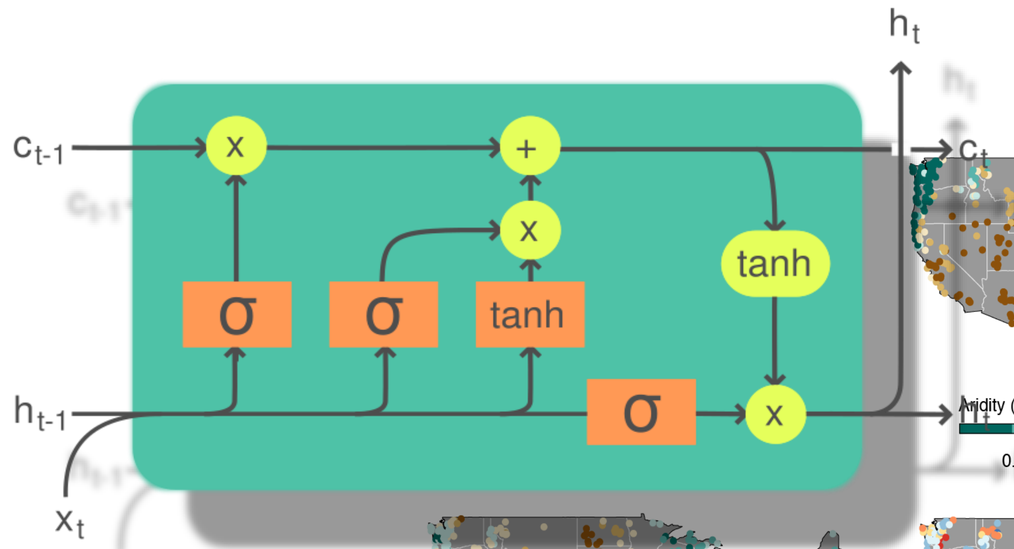

LSTM

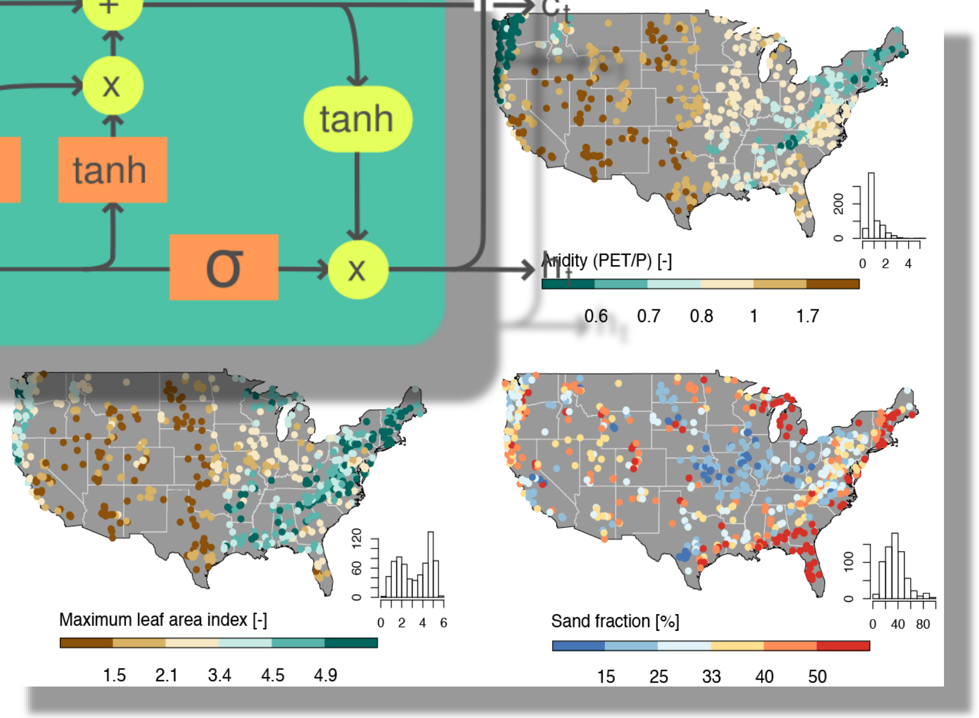

Watershed features

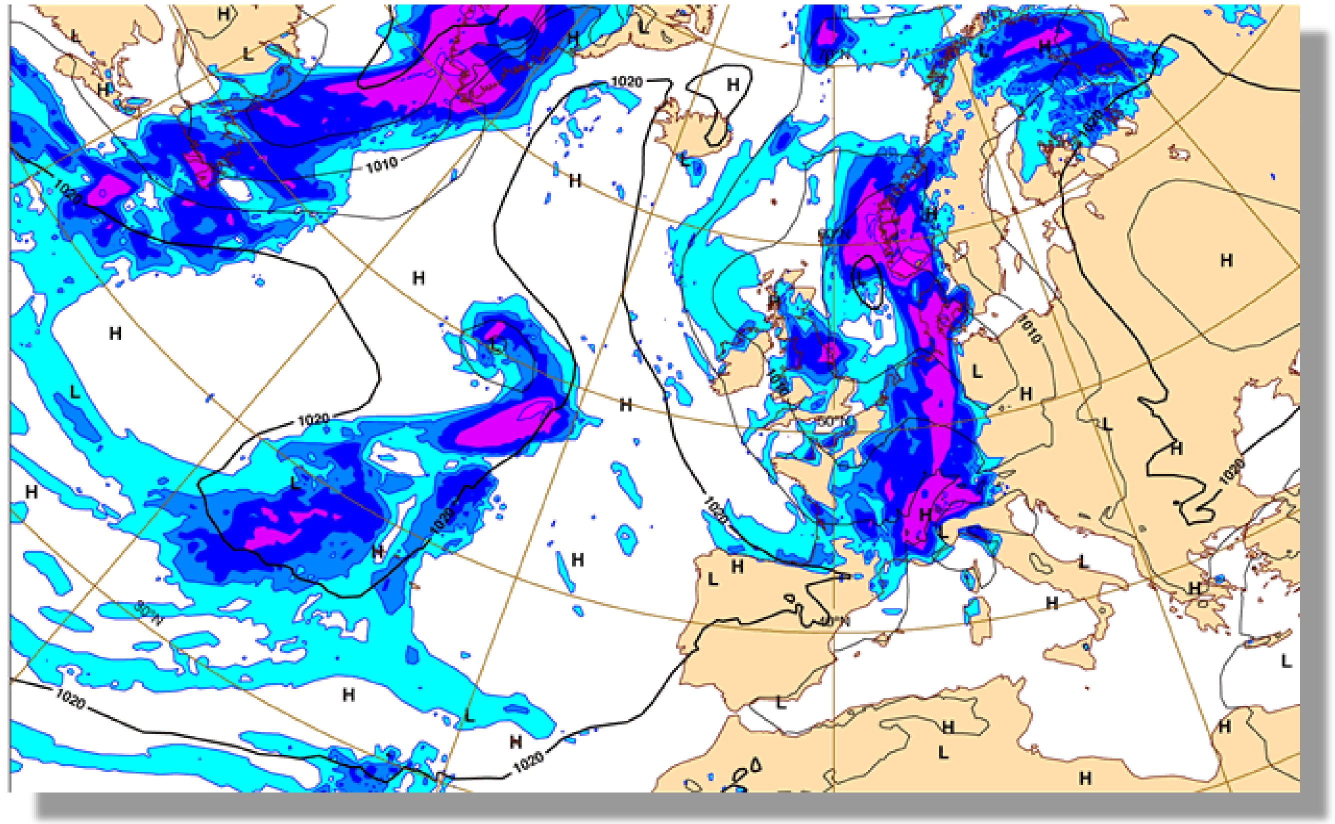

Numerical Weather Prediction (or AI-based)

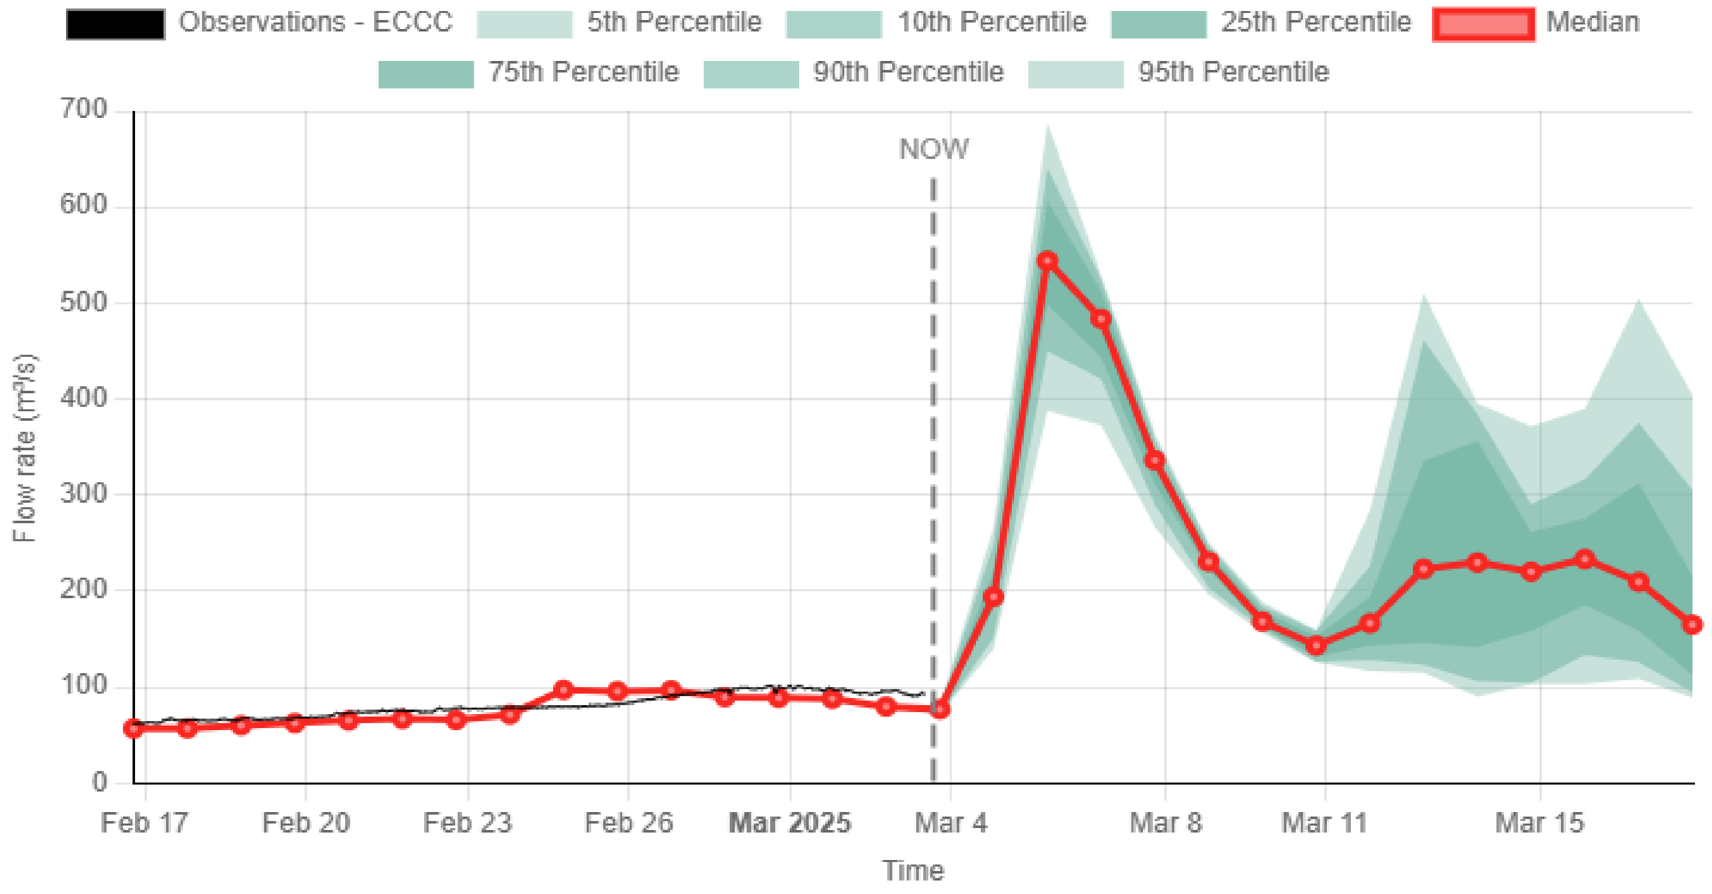

Daily forecasts at Water Survey of Canada stations

Inputs: geographic data and daily weather data, averaged over catchment/HRU

Outputs: daily streamflow

ML Model Skill (historical evaluation)

ML models are trained using our <u>custom dataset</u>:

- Weather data mainly based on AgERA5
- Fine-tuning with CaPA further improves skill (not shown here)
- Static geographic attributes are derived from **Canada1Water**
- HYDAT gauge data from WSC

- Models achieve excellent skill in **gauged basins** (if catchment was in training dataset – blue line)
- Skill deteriorates in **ungauged basins** (i.e. unknown), but is still quite good! (green dashed line)

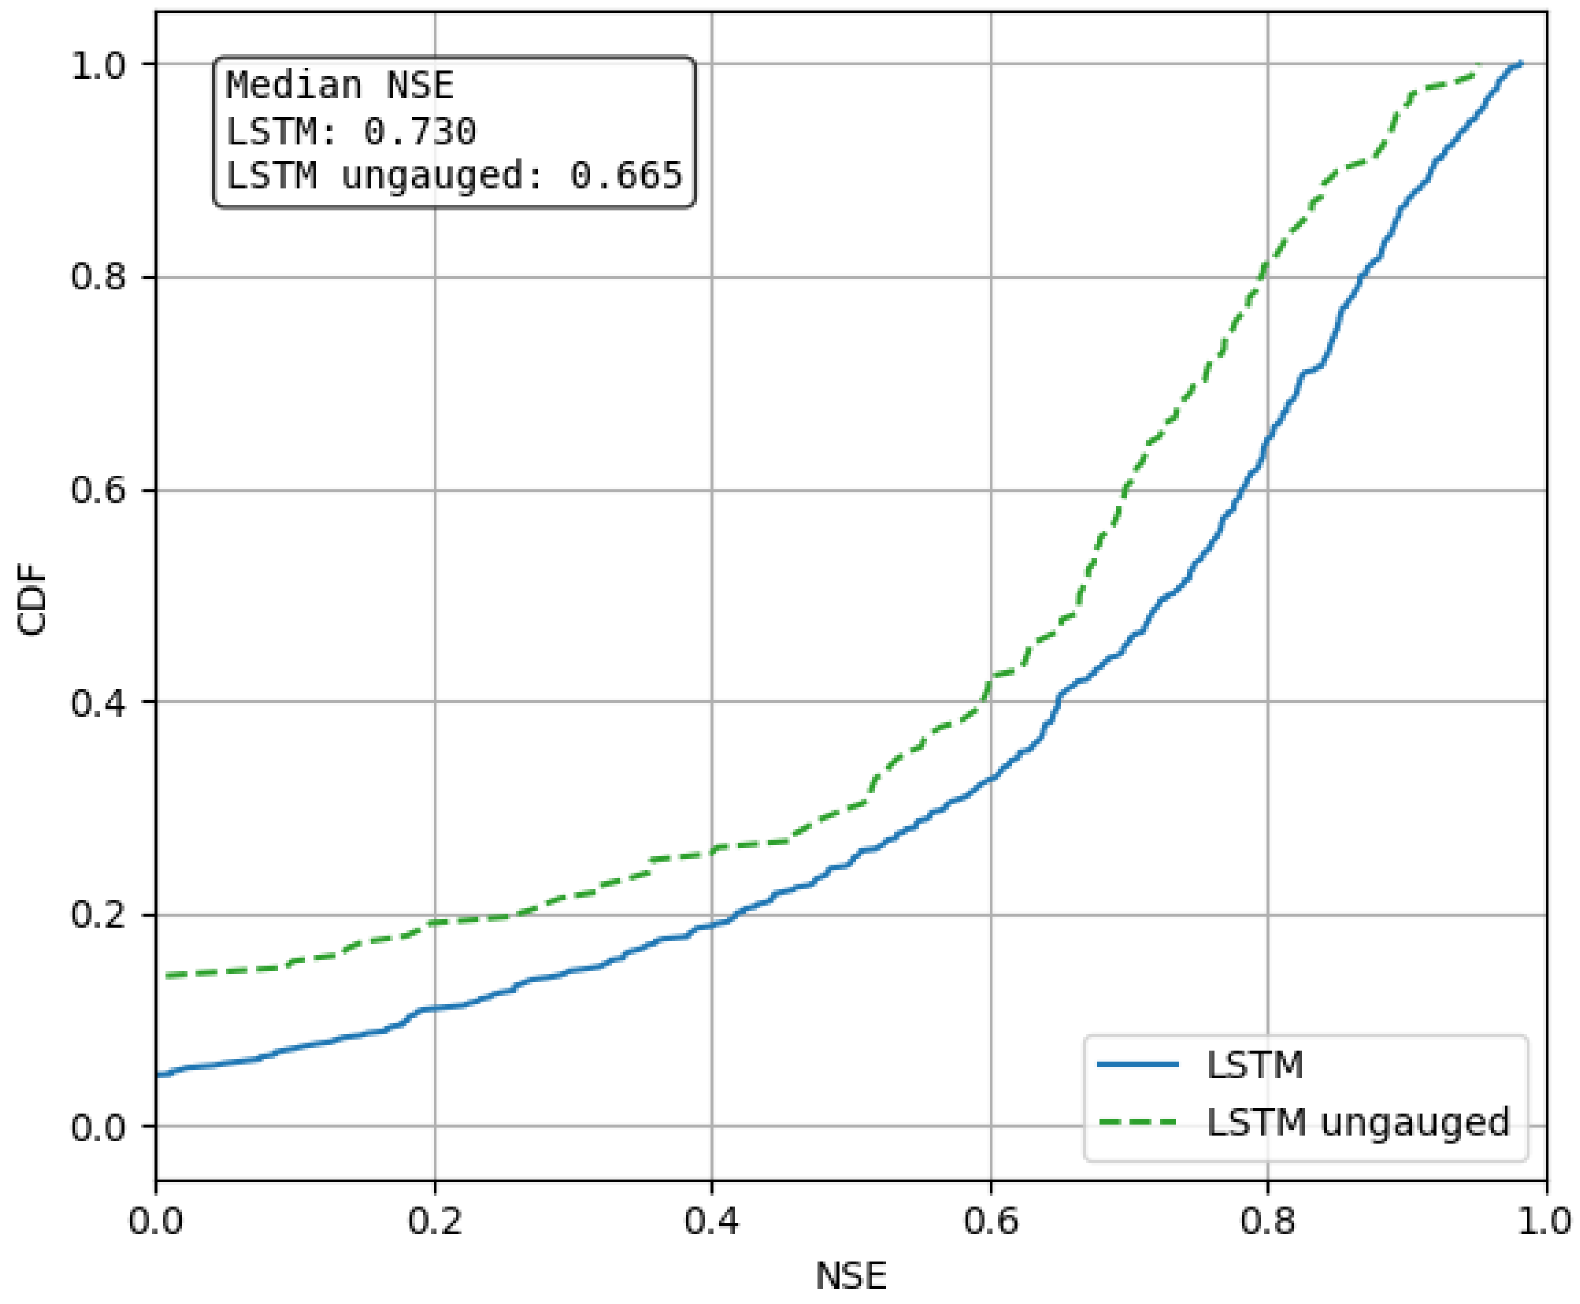

Historical simulation skill of a 10-member ensemble of LSTMs evaluated across Canada over a train/test split in time (LSTM) and over catchments (LSTM ungauged).

Gauge Data Integration (Assimilation)

ML models are trained using our <u>custom dataset</u>:

- Weather data mainly based on AgERA5
- Fine-tuning with CaPA further improves skill (not shown here)
- Static geographic attributes are derived from [Canada1Water]
- HYDAT gauge data from WSC

QLSTM: Data Integration

- QLSTMs use past streamflow observations as additional predictor
- Significantly increases skill, if streamflow observations are available

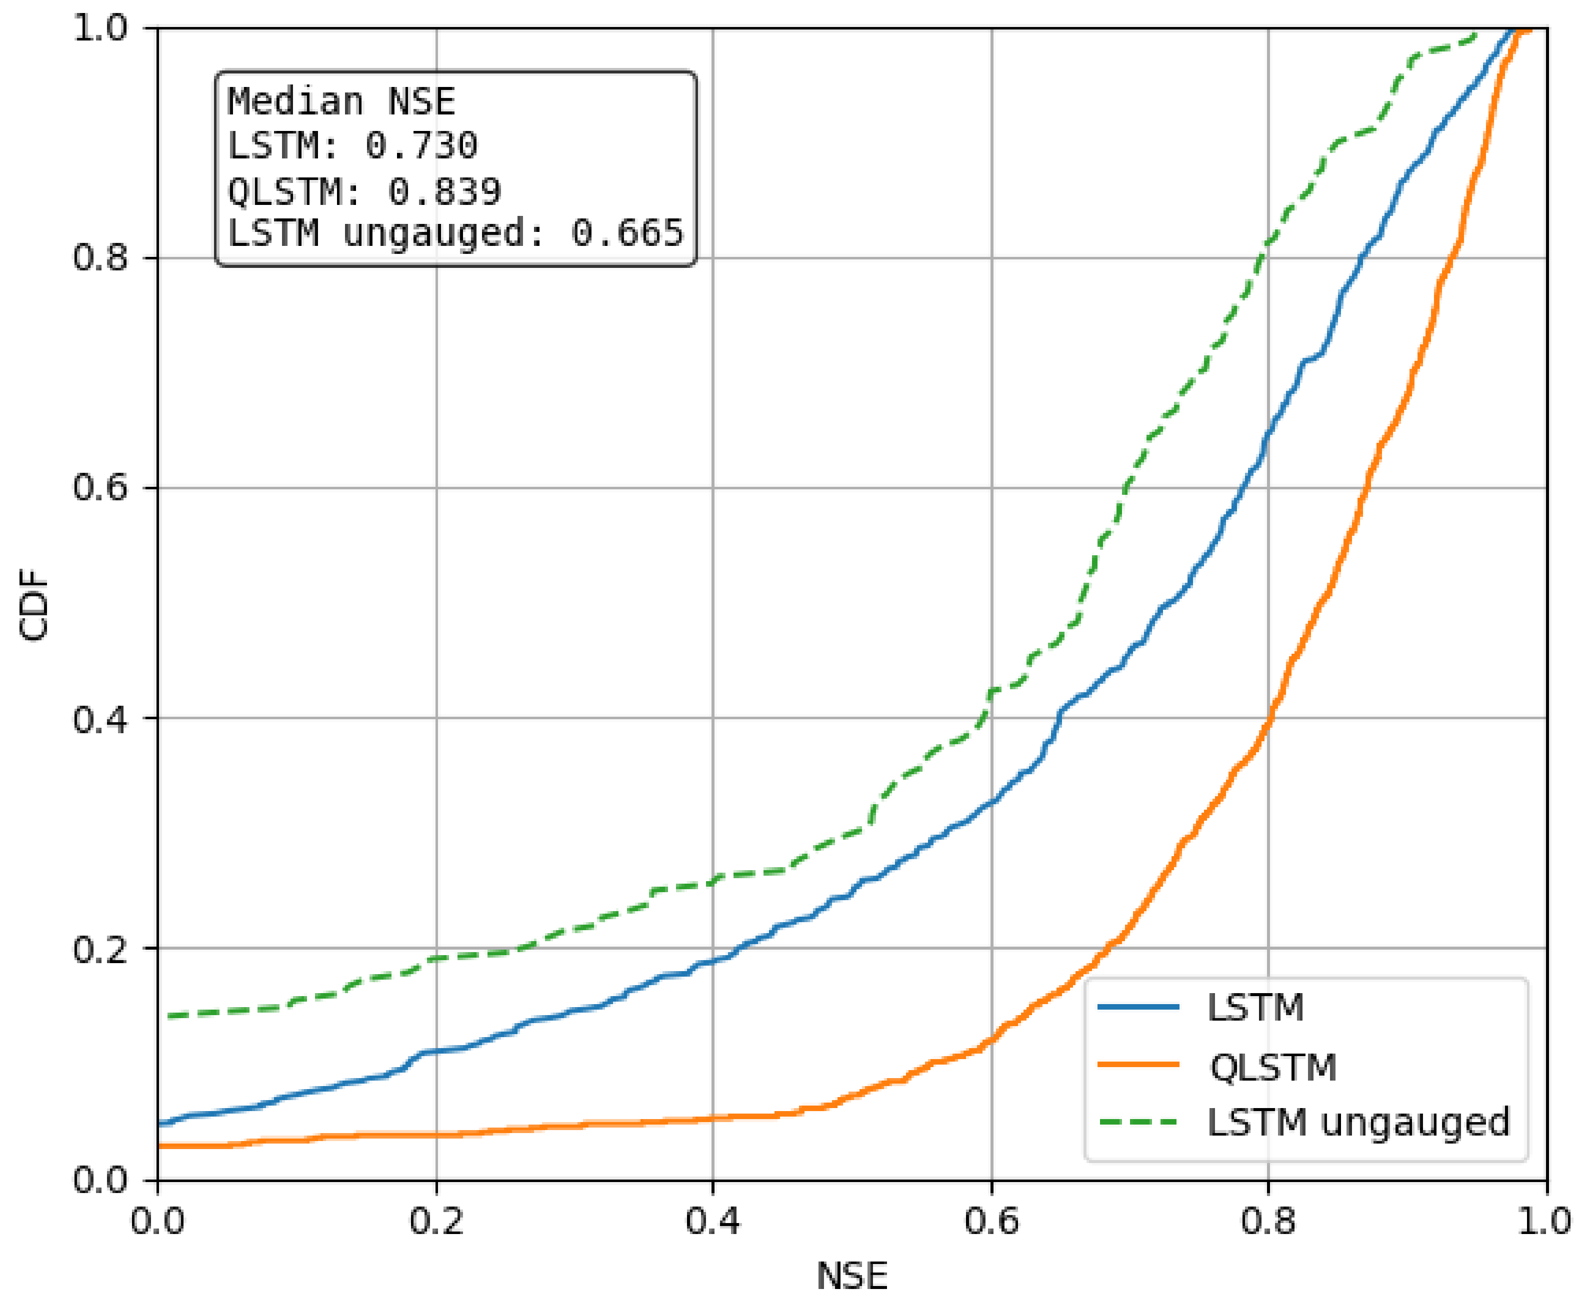

As before, plus skill of a QLSTM with streamflow data integration 50% of the time (50% hold-out), evaluated with train/test split in time only.

Skill Distribution across Canada (QLSTM)

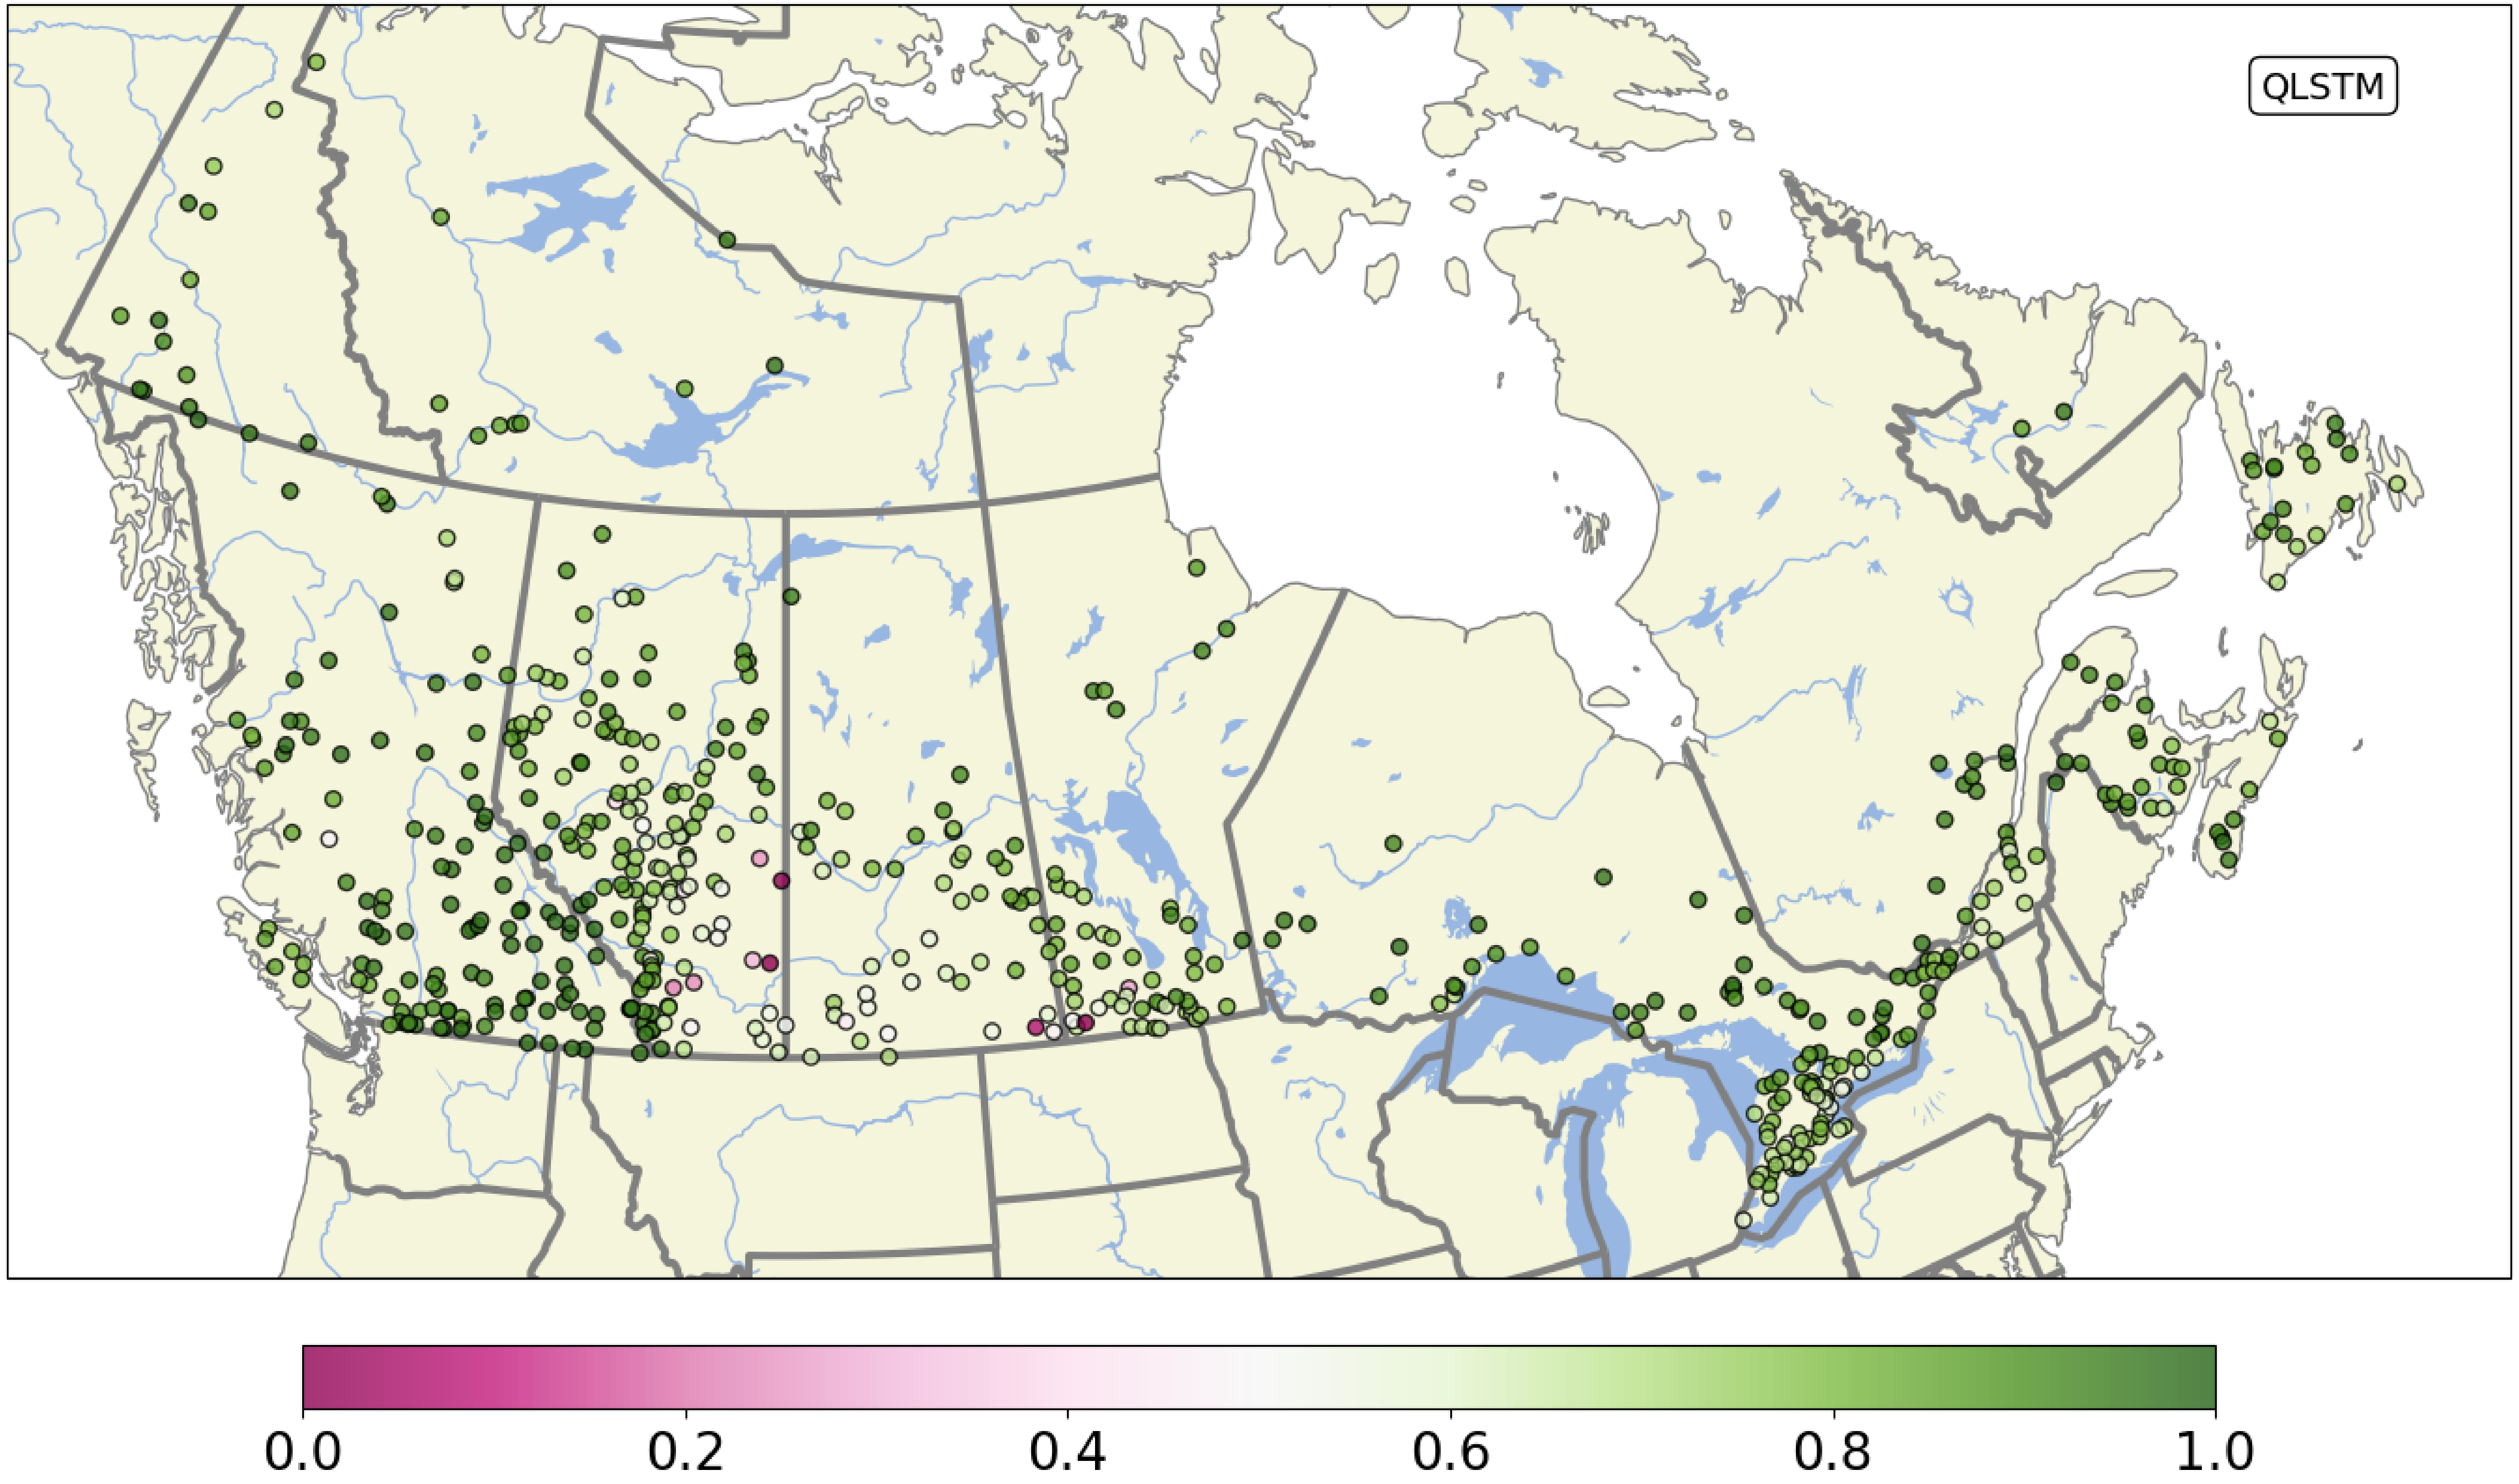

QLSTM Skill Distribution

- Very strong skill in snow-dominated catchments
- Prairies have weakest skill (like most models)
- Southern Ontario may have too much human interference

Gauge Data Integration (Assimilation)

ML models are trained using our <u>custom dataset</u>:

- Weather data mainly based on AgERA5
- **Fine-tuning with CaPA significantly improves skill**
- Static geographic attributes are derived from Canada1Water
- HYDAT gauge data from WSC

QLSTM: Data Integration

- QLSTMs use past streamflow observations as additional predictor
- Significantly increases skill, if streamflow observations are available

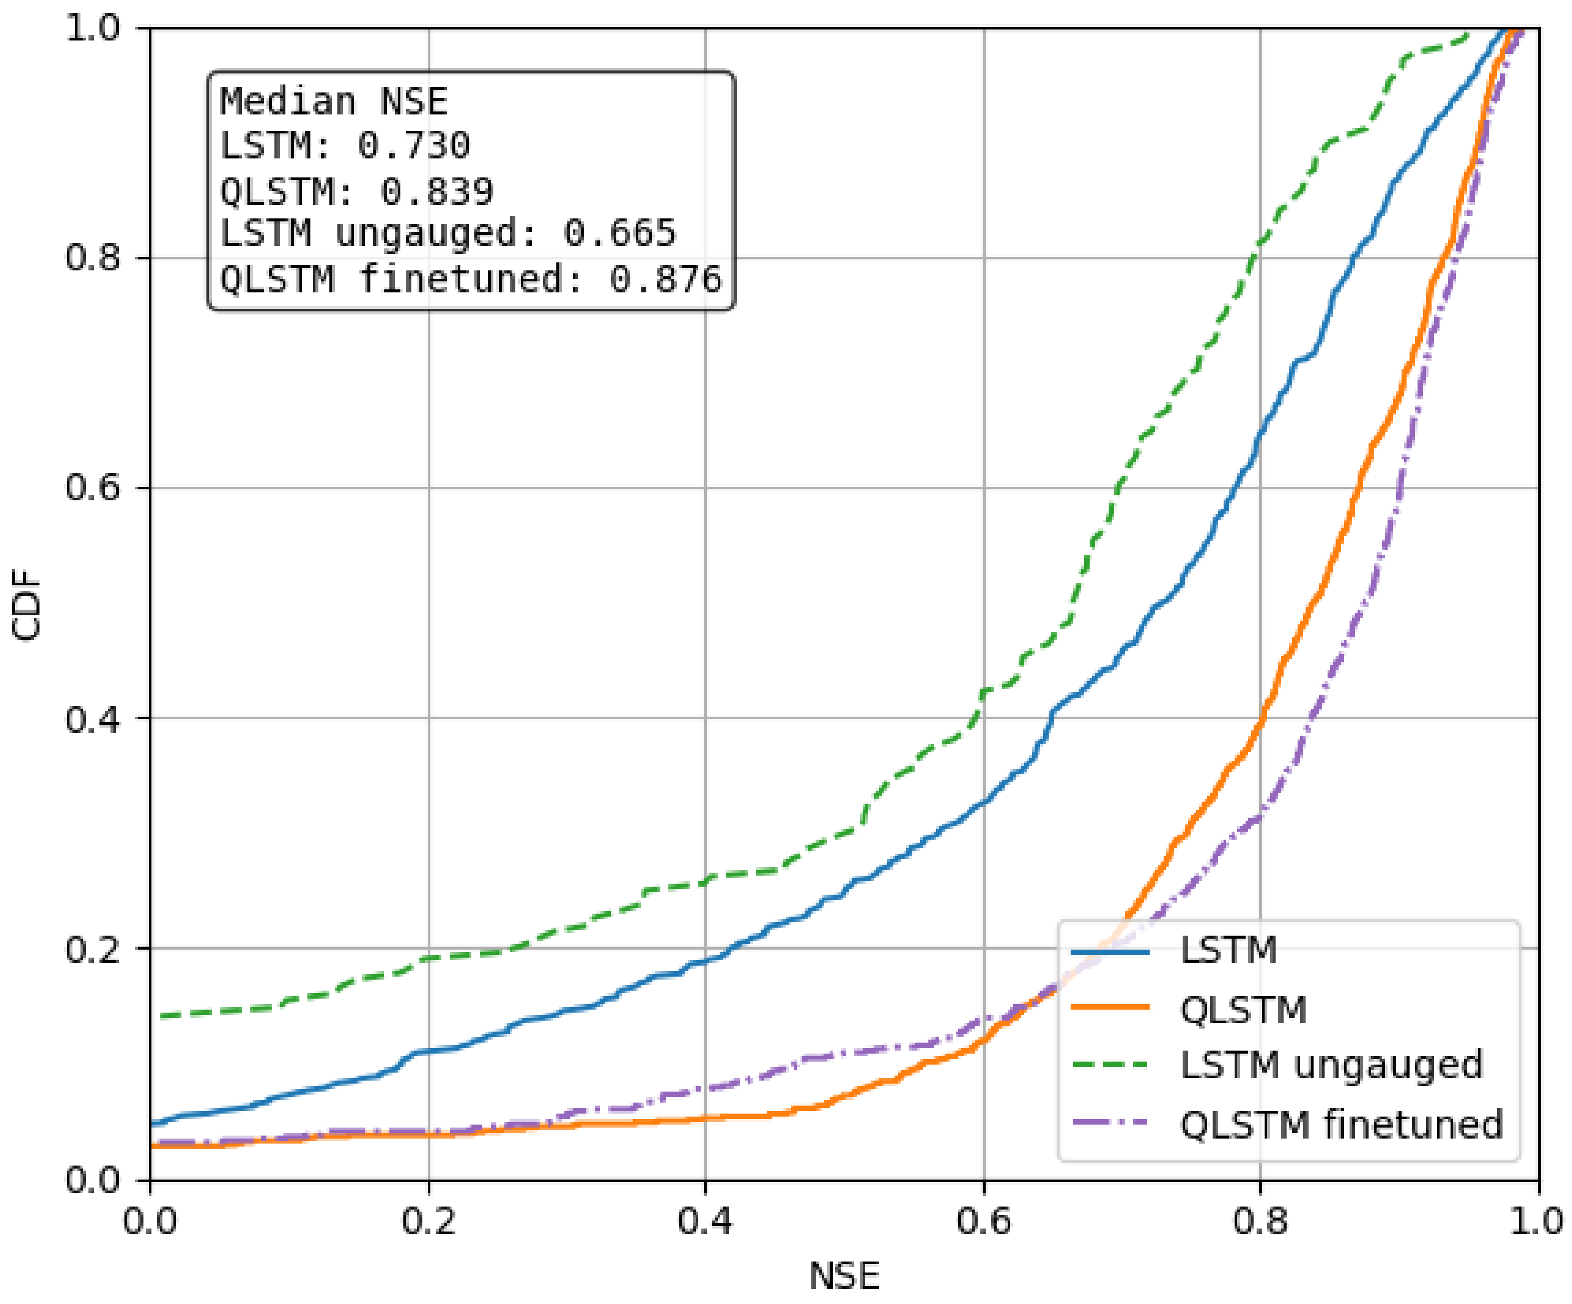

As before, with skill of a standard QLSTM with streamflow data integration (50% hold-out), and a QLSTM fine-tuned using CaPA.

HydroSphereAI Features

- Fine-tuning with augmented, custom datasets
- Custom forecast end-points and synthetic gauges

- Integration of real-time observed streamflow data integration (data assimilation)
- Real-time rain gauges

- Flood forecasting with hourly time stepping
- High/low flow alerts
- (Sub-)seasonal drought forecasting

- Snowpack forecasts
- Soil moisture forecasts
- (and streamflow)

Example: a Seasonal Simulation with Snow

1-year historical simulation, with observations and historical normals:

- **Alpine headwaters between Banff & Jasper**
- Snow accumulation over winter, followed by runoff peak in spring (freshet)
- ML models capture seasonal behavior well
- Forced by historical climate (AgERA5)

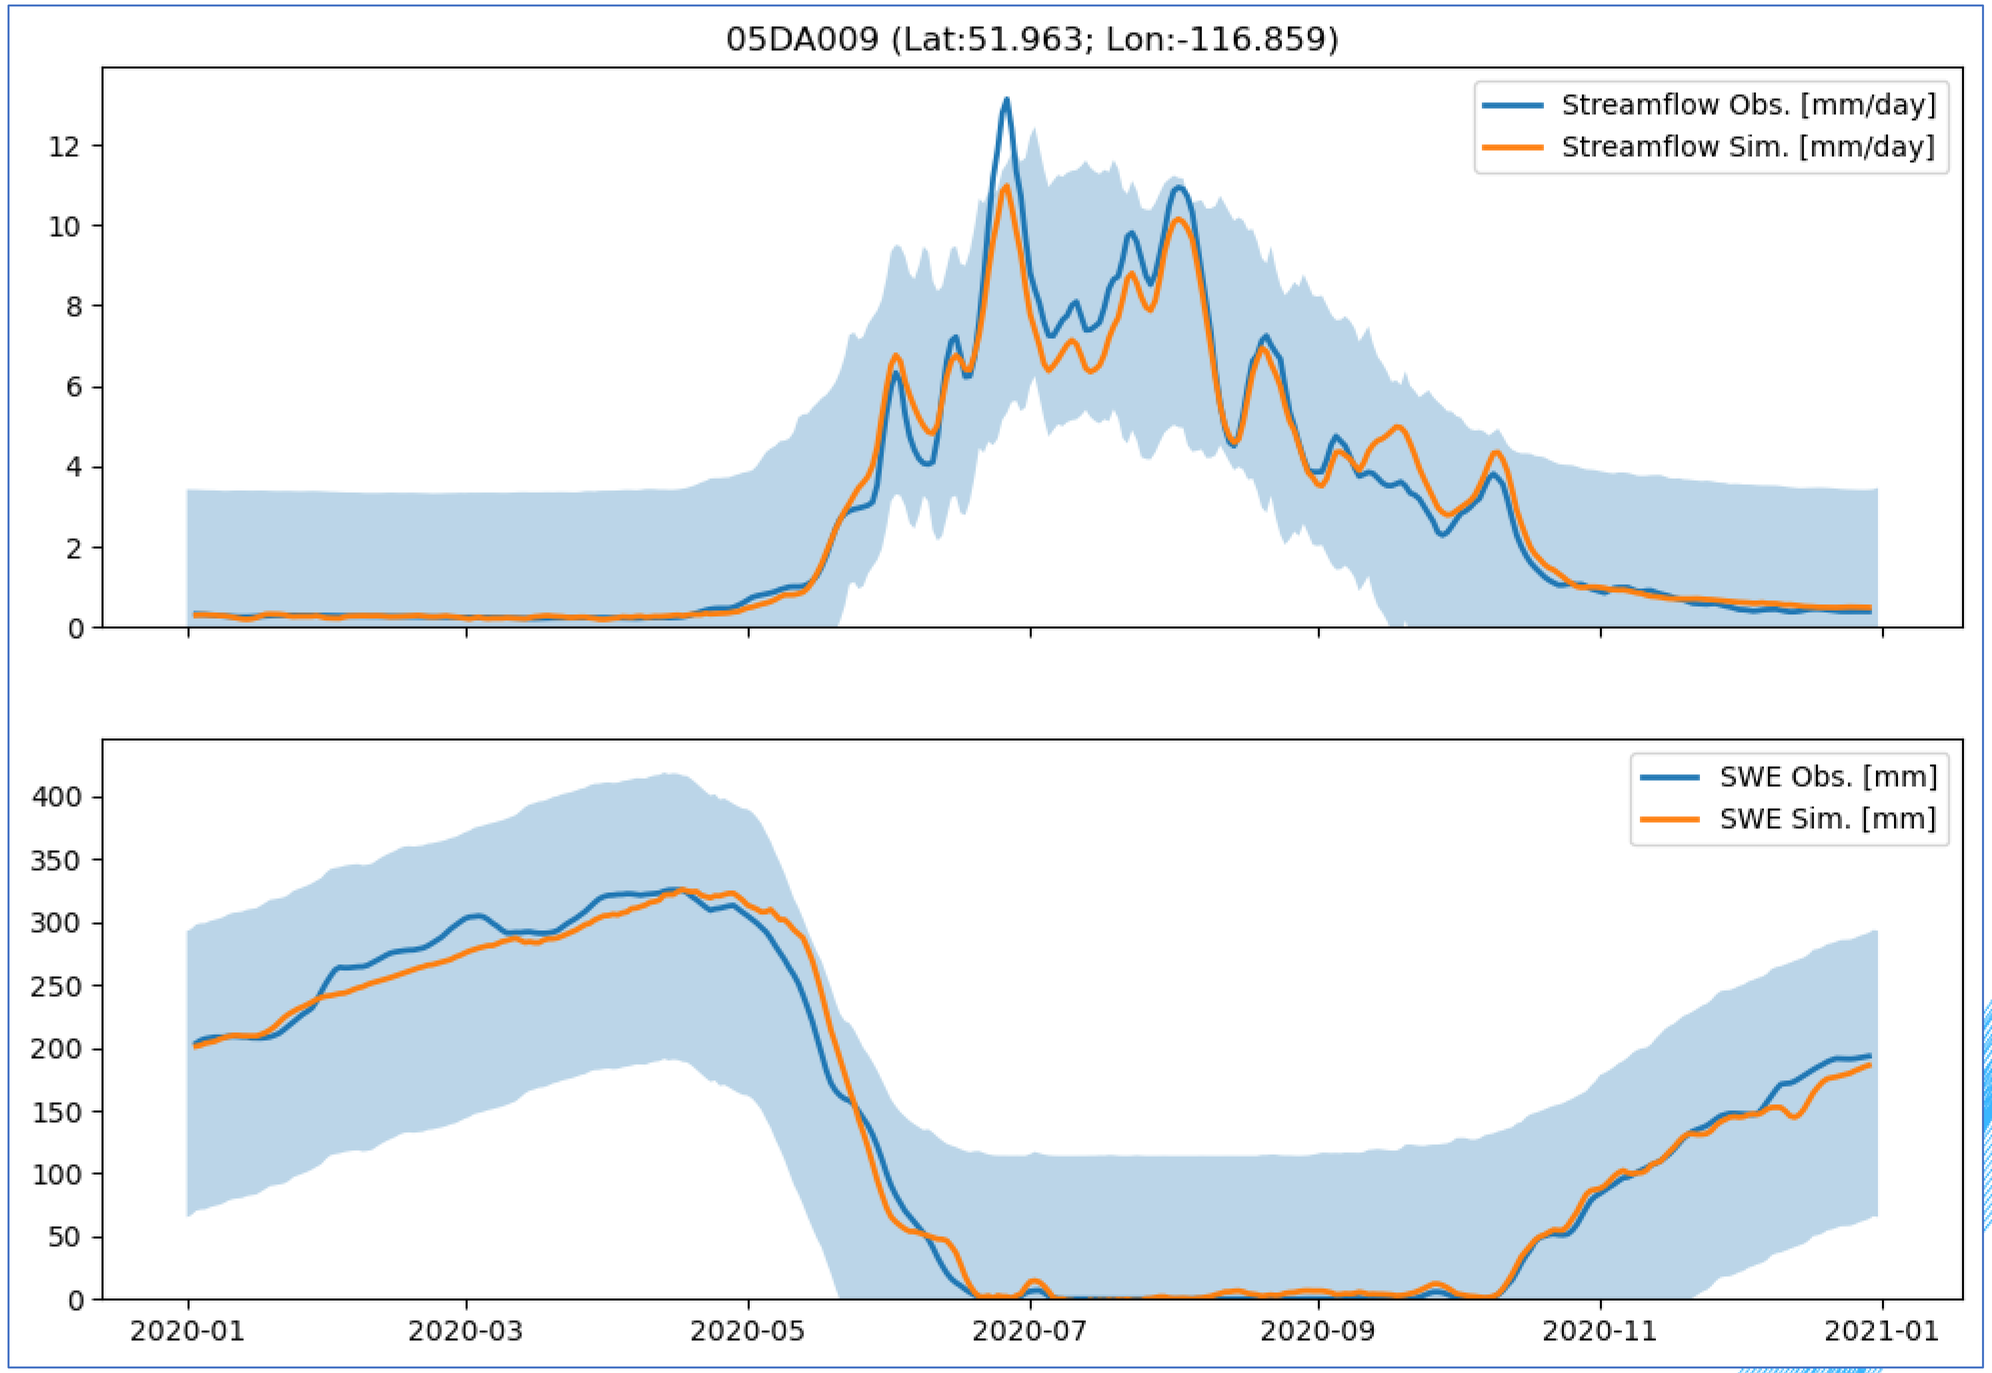

One year of historical simulation (hindcast) with observed streamflow and snow from ERA5-Land; observed historical range in light blue (±σ).

Example: a Seasonal Simulation with Snow

1-year historical simulation, with observations and historical normals:

- **A watershed in eastern Quebec**
- Snow accumulation over winter, followed by runoff peak in spring (freshet)
- ML models capture seasonal behavior well
- Forced by historical climate (AgERA5)

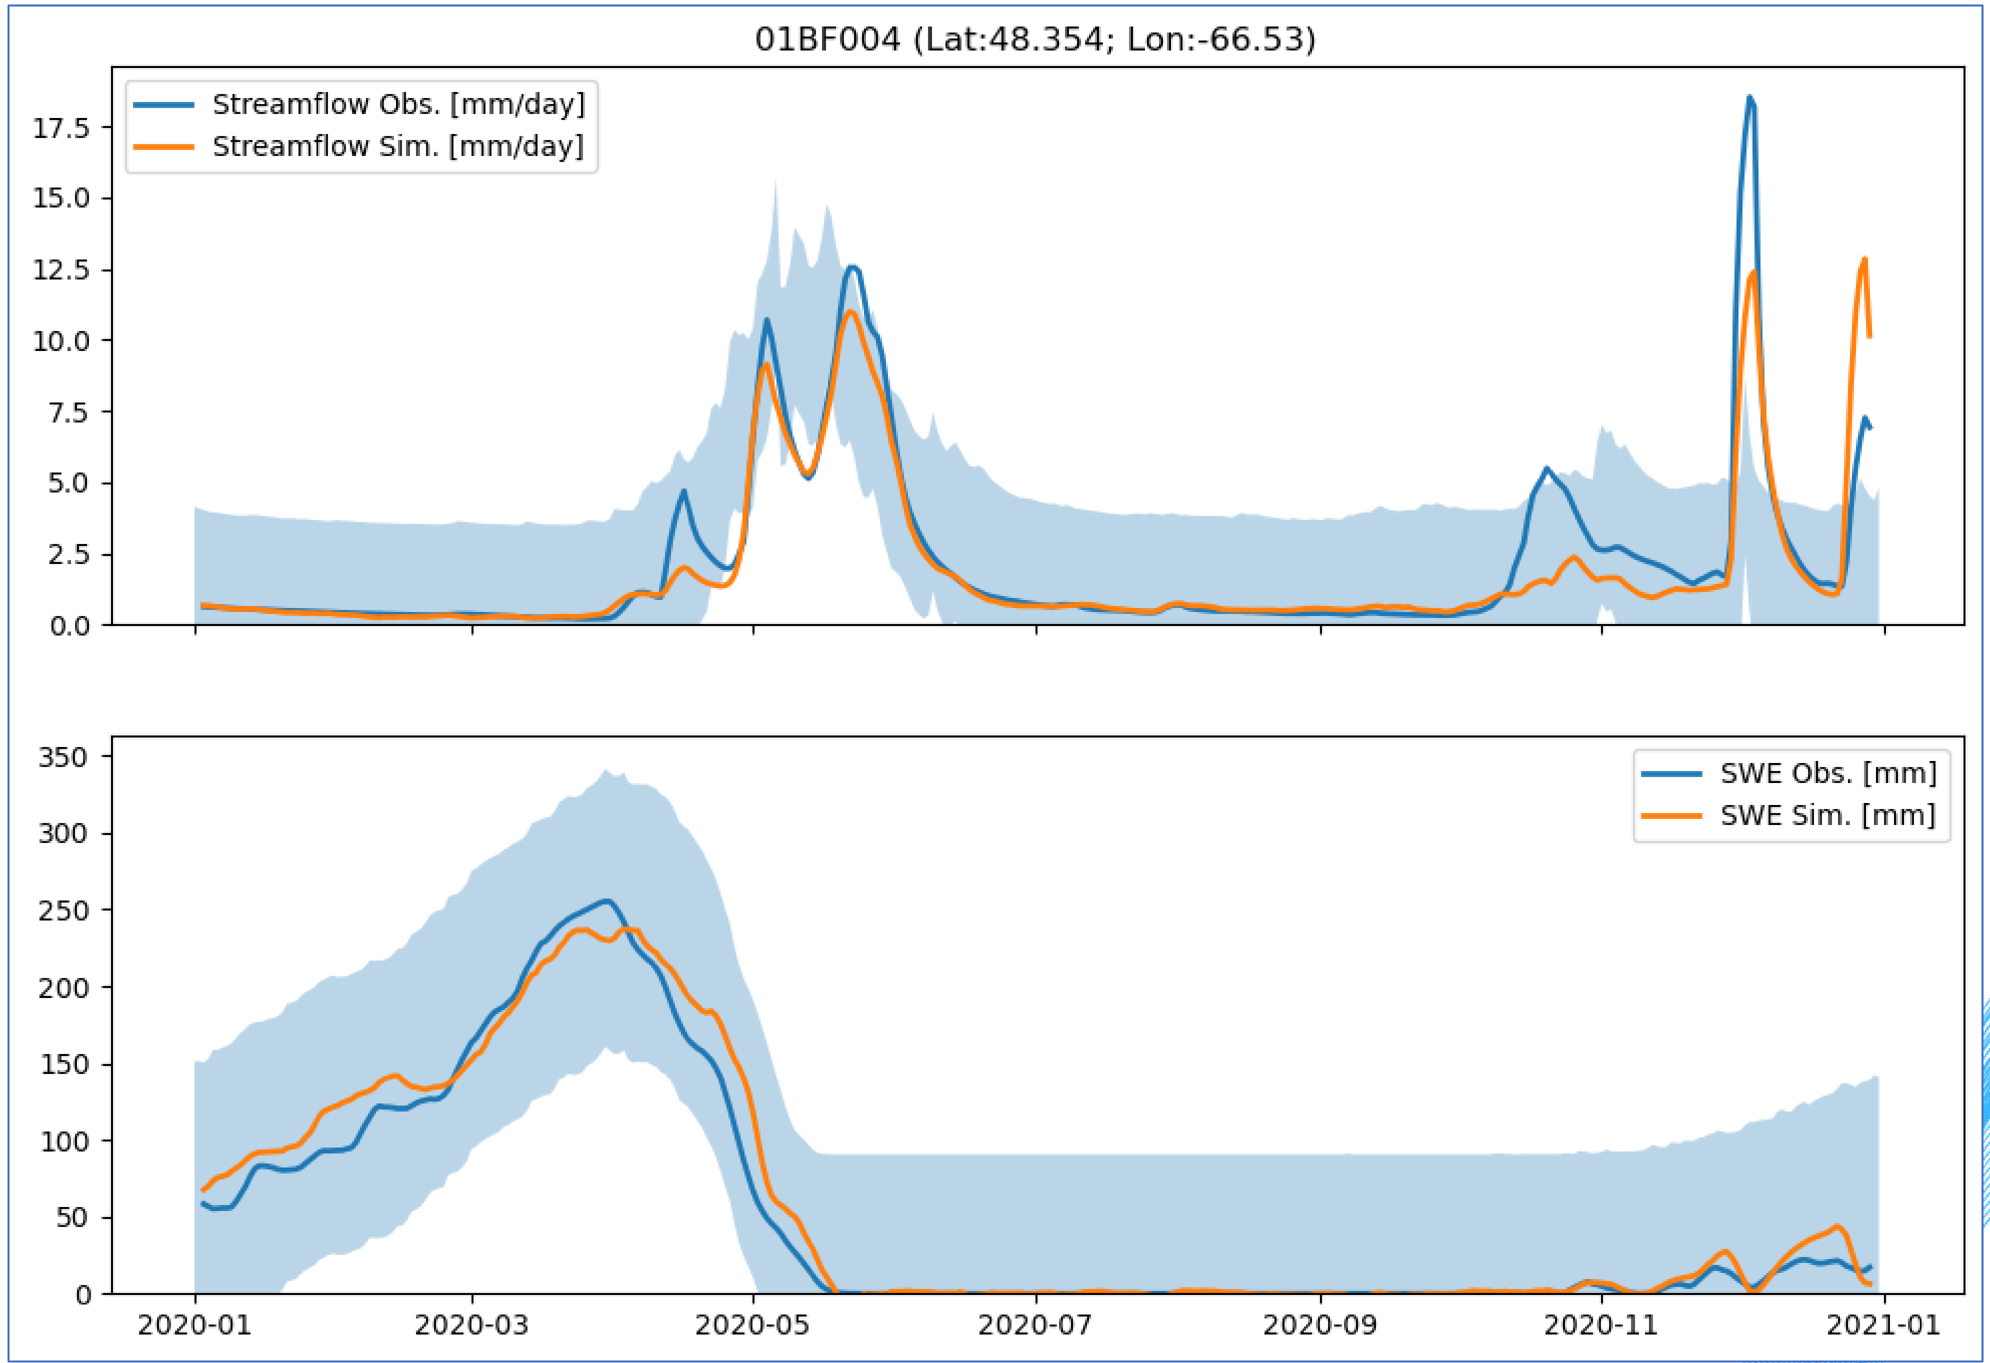

One year of historical simulation (hindcast) with observed streamflow and snow from ERA5-Land; observed historical range in light blue (±σ).

Short-range Hourly Forecasts

- Captures short-term variability by simulating multiple time-scales (MTS models)
- Hourly models achieve performance comparable to daily (median NSE: 0.71).
- Useful for extreme events and **flood forecasting**
- Can be **integrated with real-time rain gauges** and streamflow data

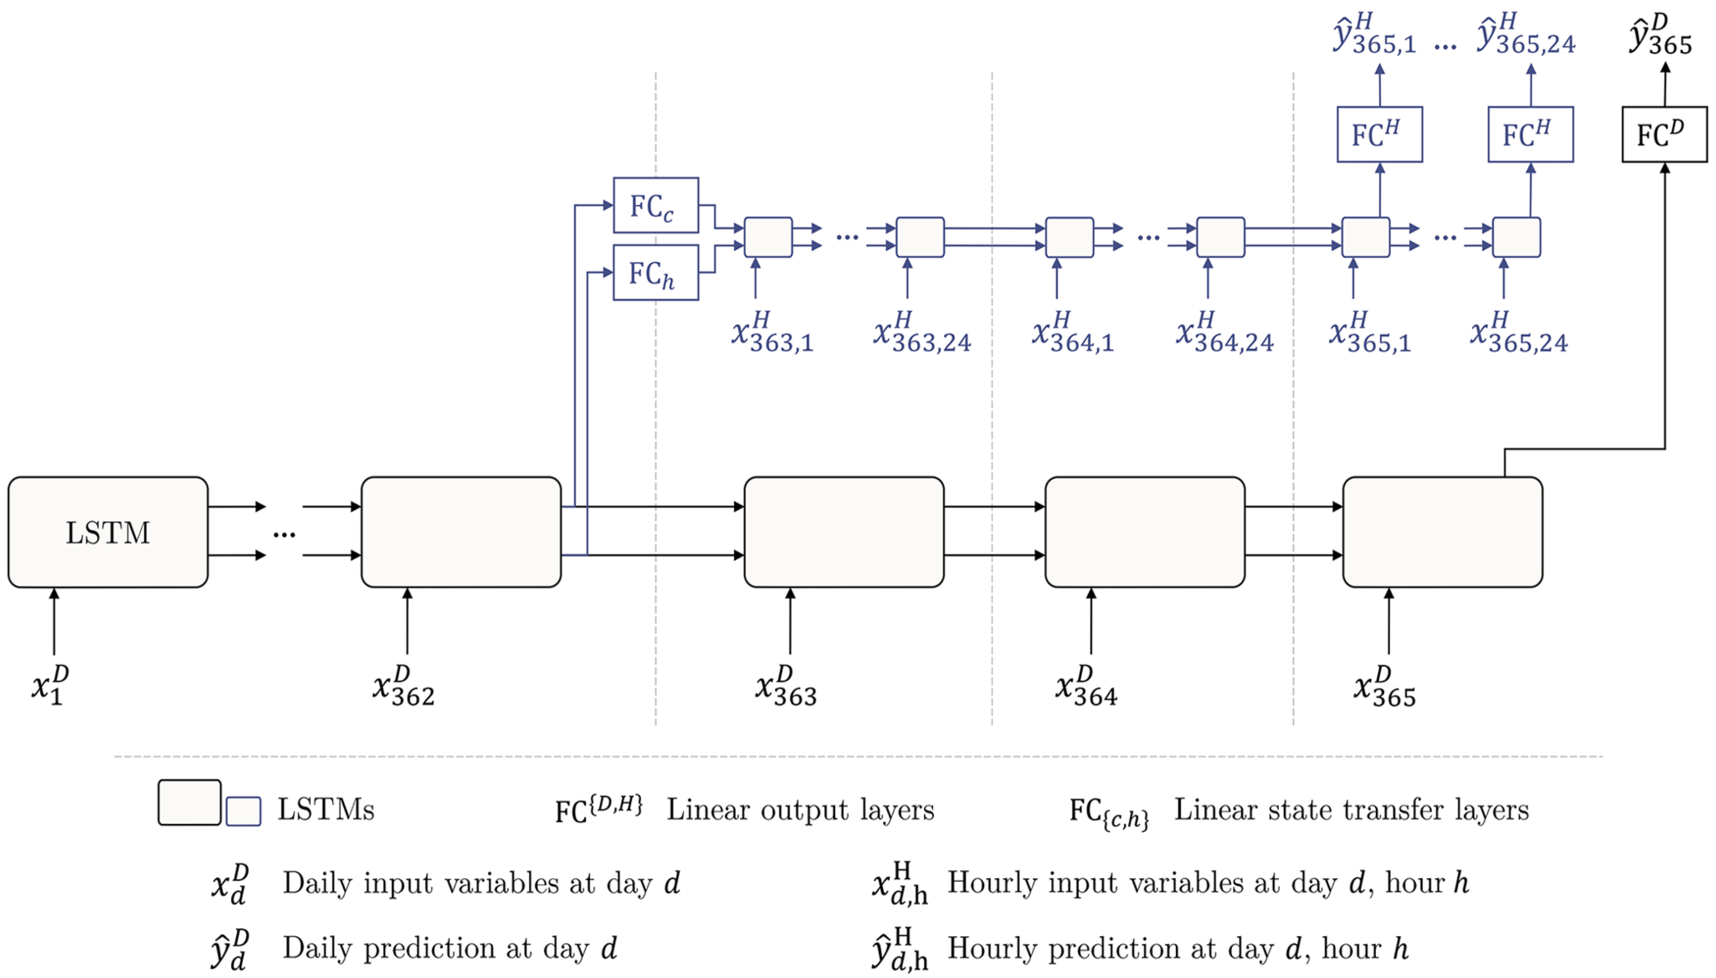

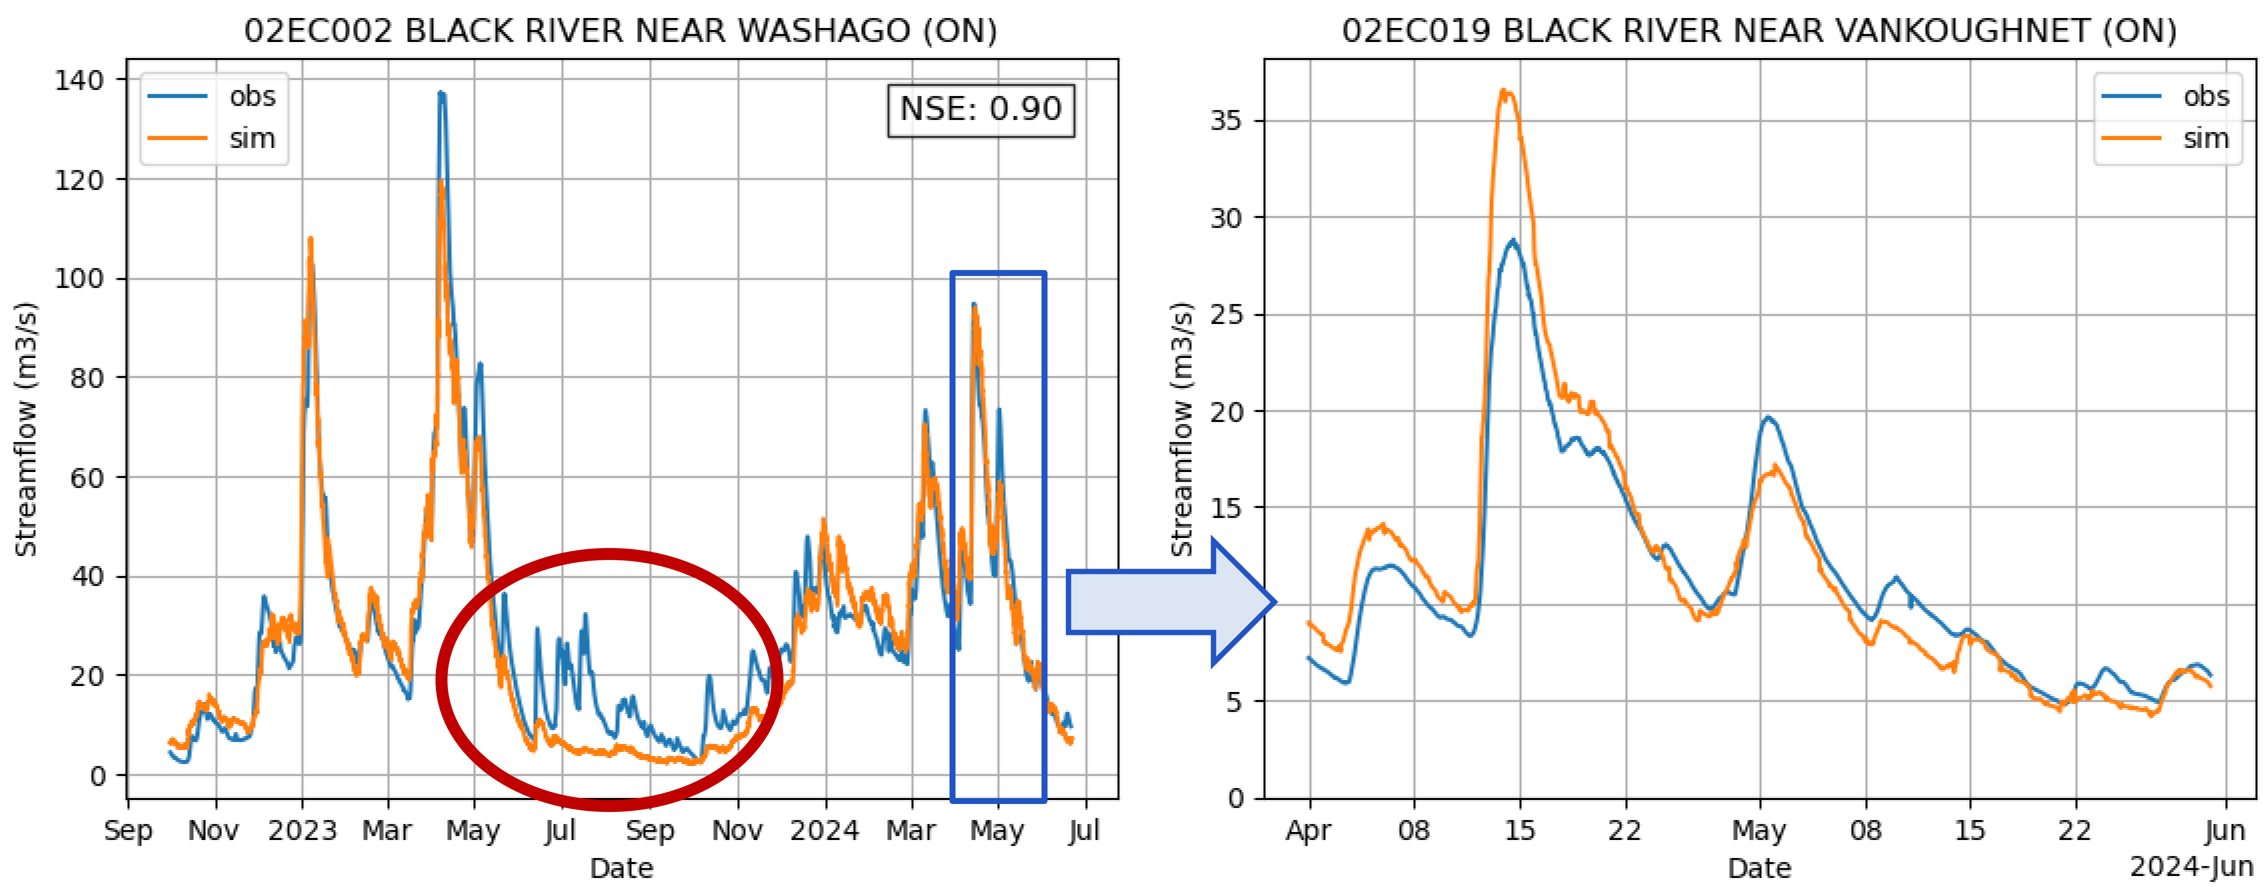

Top: Structure of the Multi-Time-Scale models: an hourly model is initialized from and constrained to a daily parent model (Gauch et al. 2021).

Left: Hourly (MTS) models for a catchment in Northern Ontario for WY2023/24 (left) and Apr & May 2024 (right). Note different gauges!

Canada-wide Deployment

Aquanty is already running forecasts for all HYSETS watersheds and has static attributes and forcing for all HyDat catchments.

Deployment across Canada will be very straightforward:

- ML Models scale very well – computational resources are already available
- Static/geographic attributes available from Canada1Water
- Weather data from ECCC/MSC: CaPA, HRDPS, GEPS
- Larger rivers will need routing

Routing for Large Basins

Routing for a large headwater subbasin of the Peace River.

Smoky River at Watino (07GJ001), ~50,000 km^2

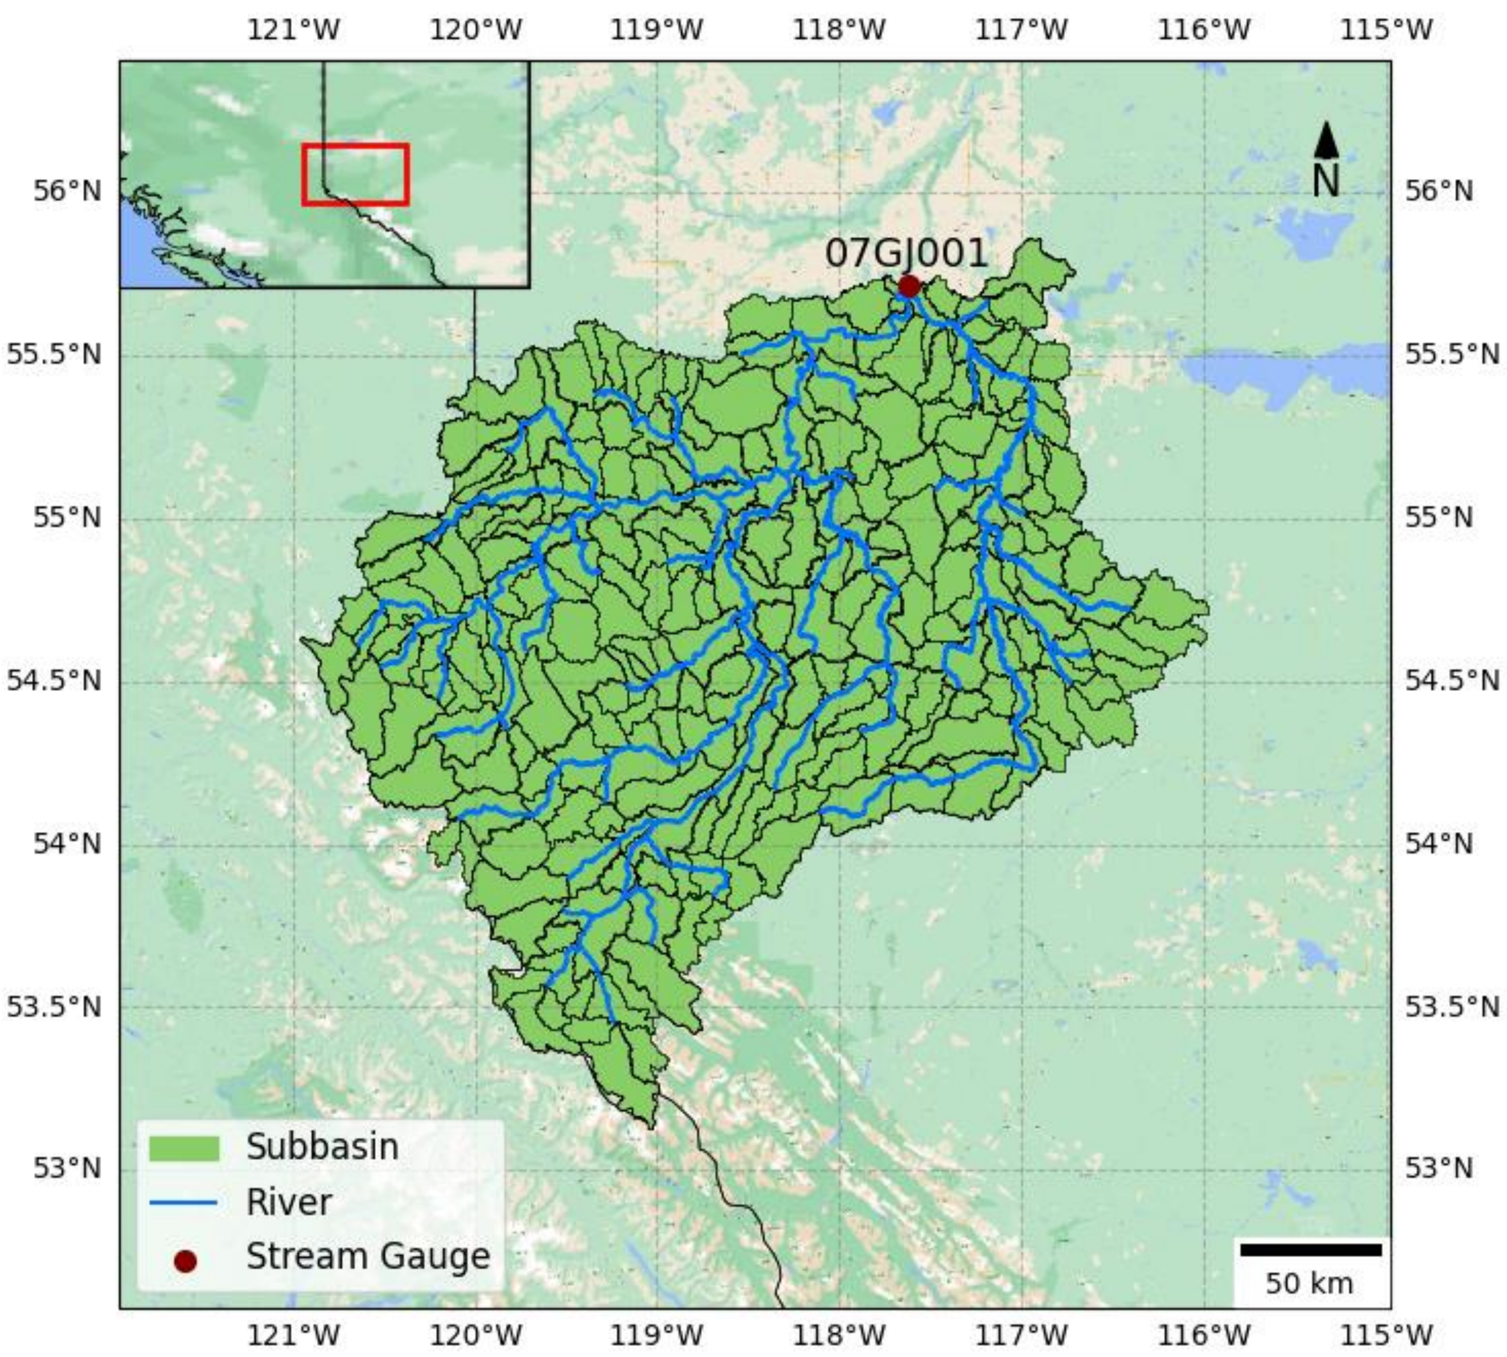

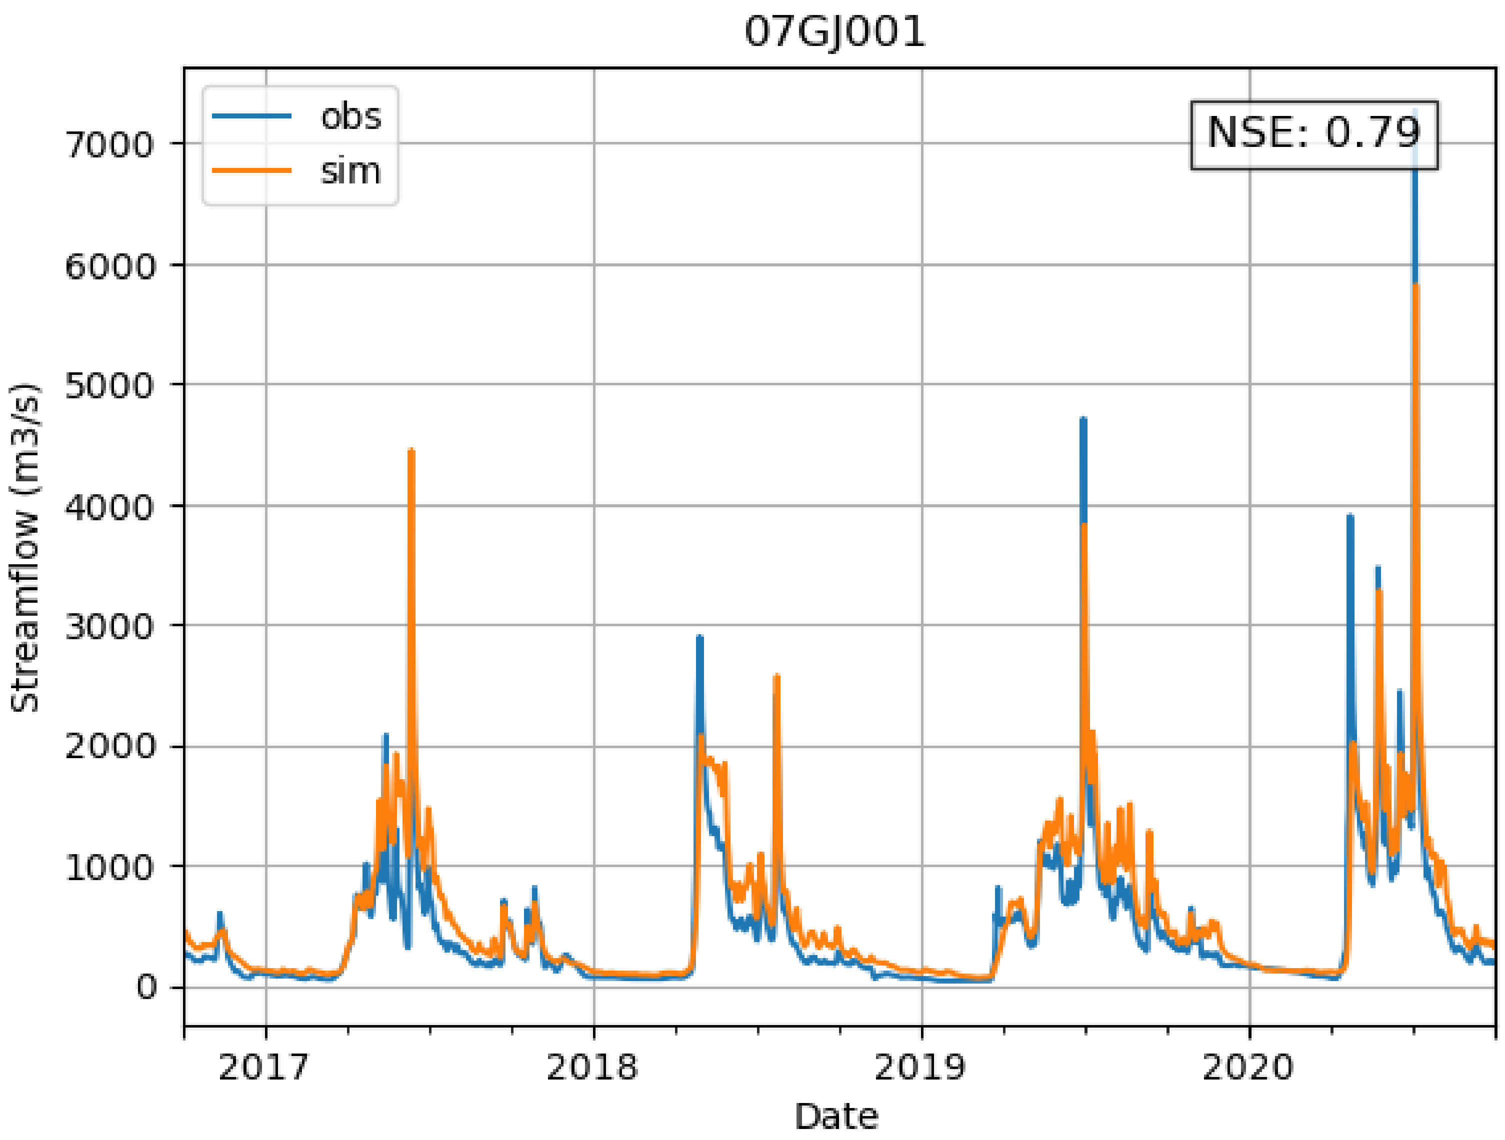

Simulation skill of the routing model; only showing test period (excluding calibration period)

- The basin is discretized into a complete and non-overlapping set of subbasins
- Each subbasin (< 2000 km^2) is simulated using an LSTM
- LSTM forecasts are connected using a routing model

From Model to Forecast

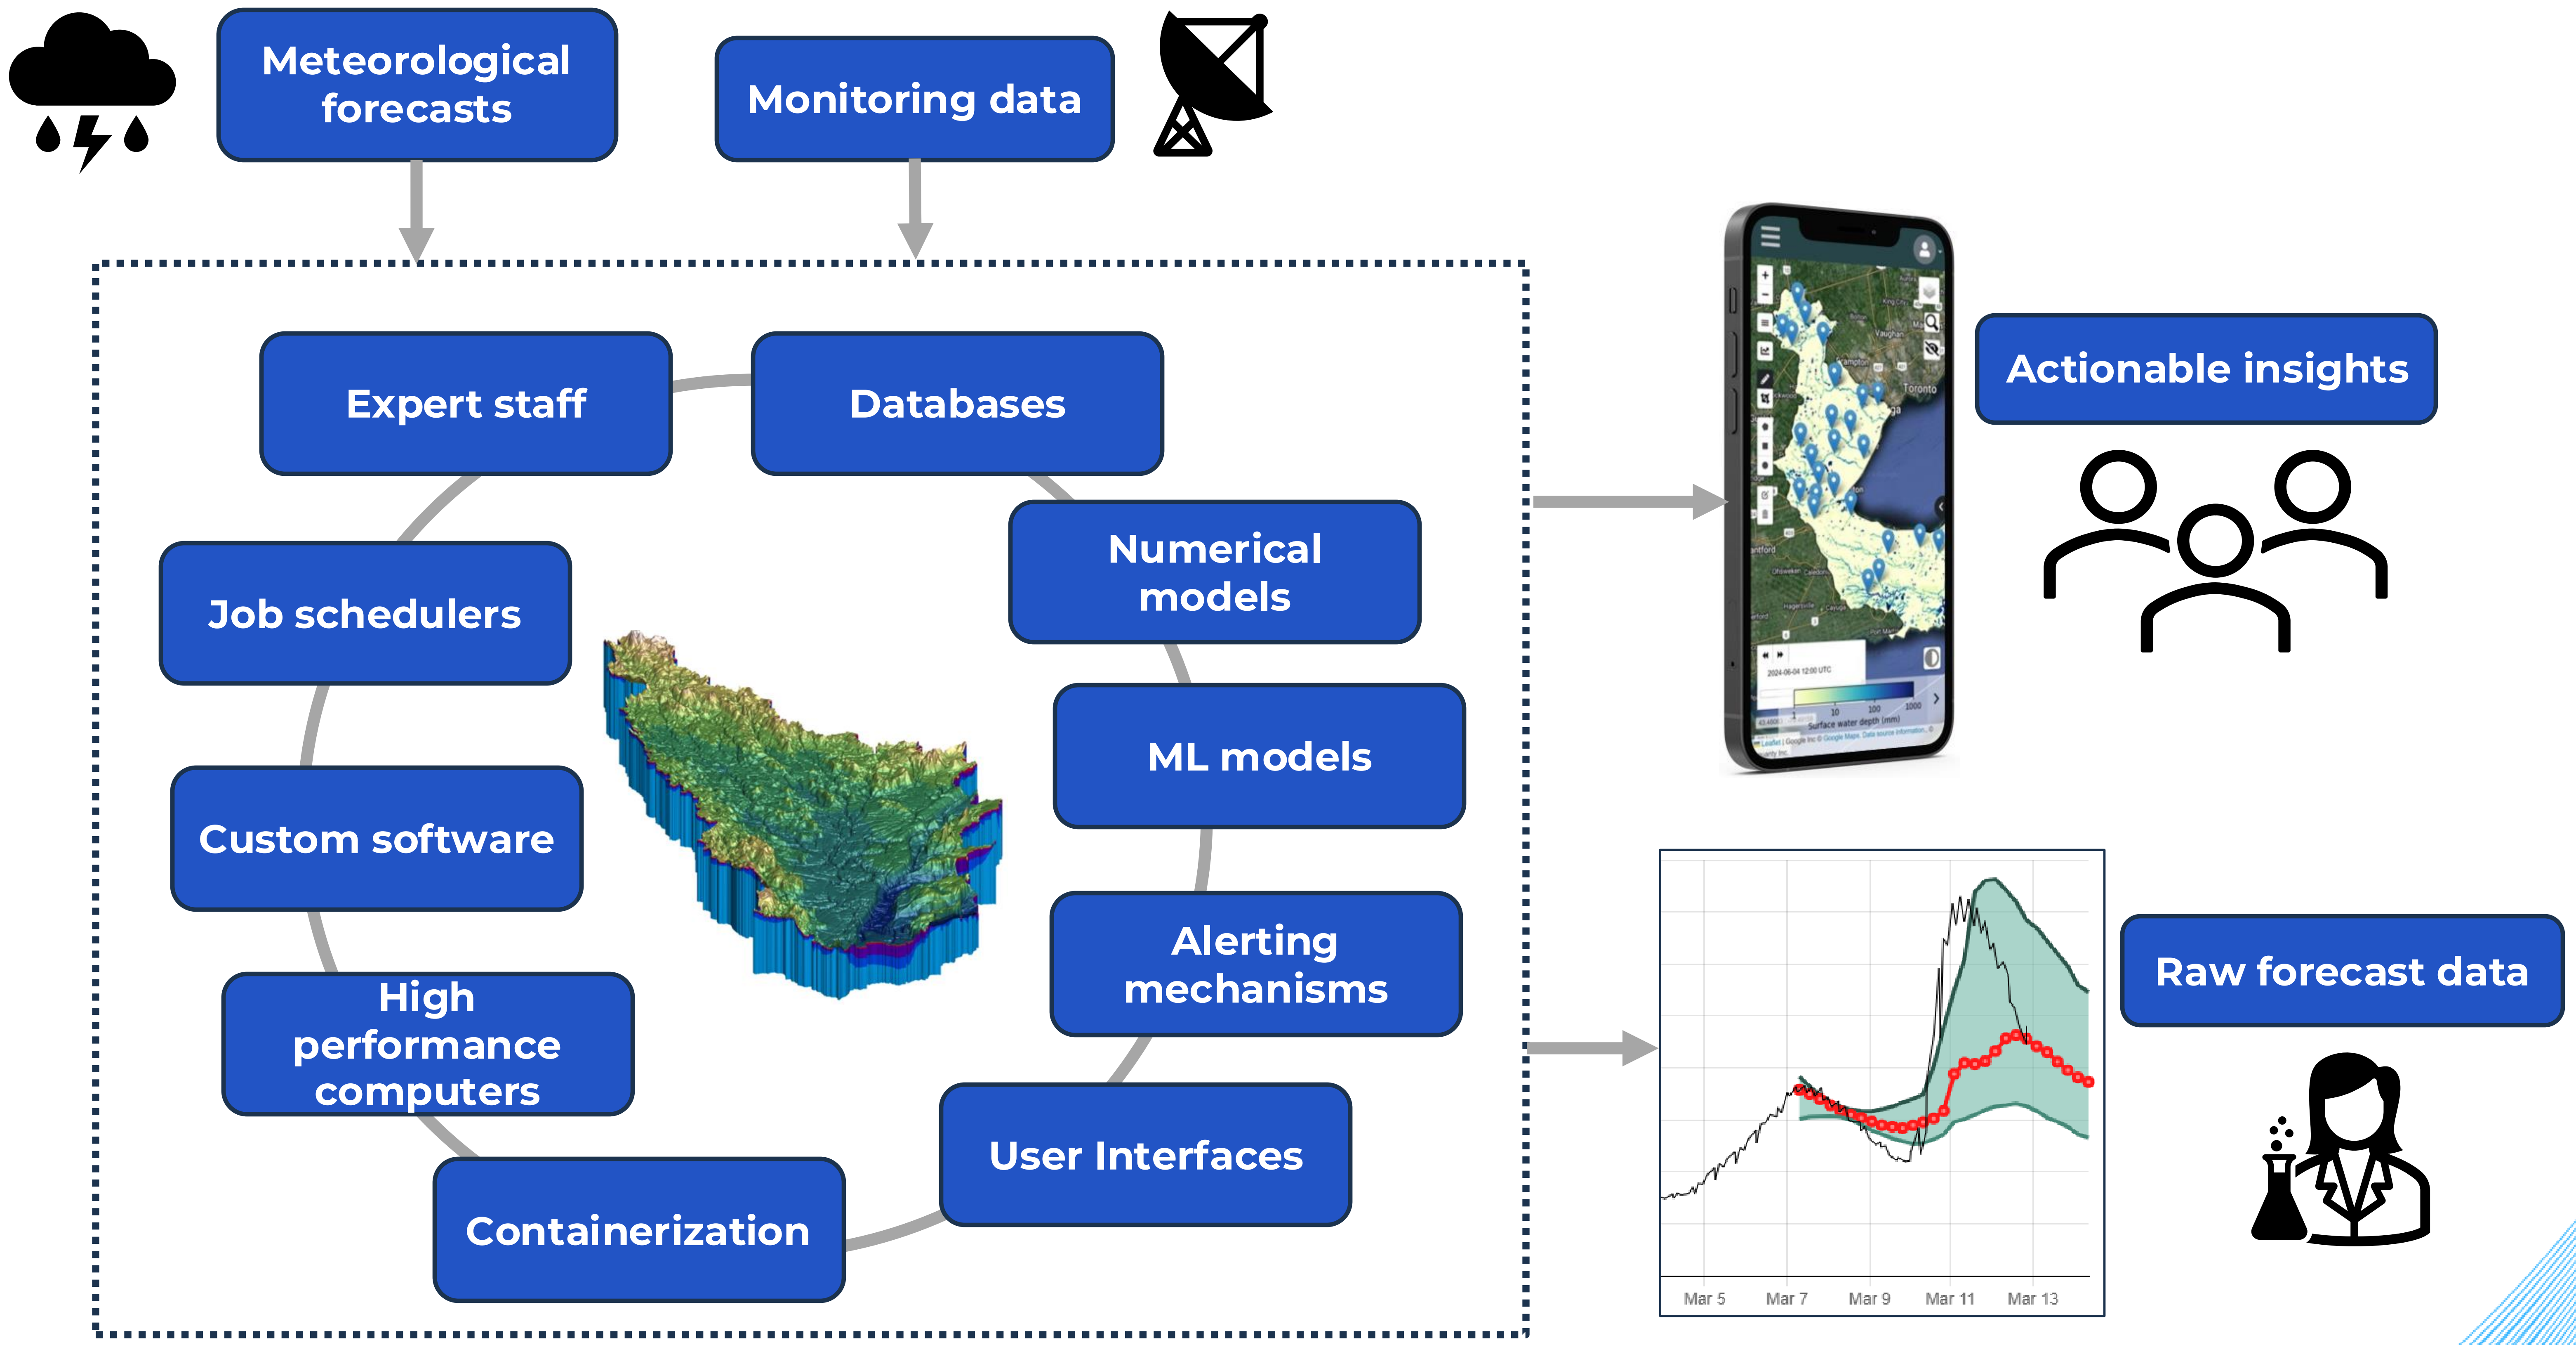

Beta-Test with Conservation Ontario

aquanty

Operational deployment with all Conservation Authorities in ON

- At all WSC stations < 2000 km^2 in 36 Ontario watersheds

Conservation ONTARIO
Natural Champions

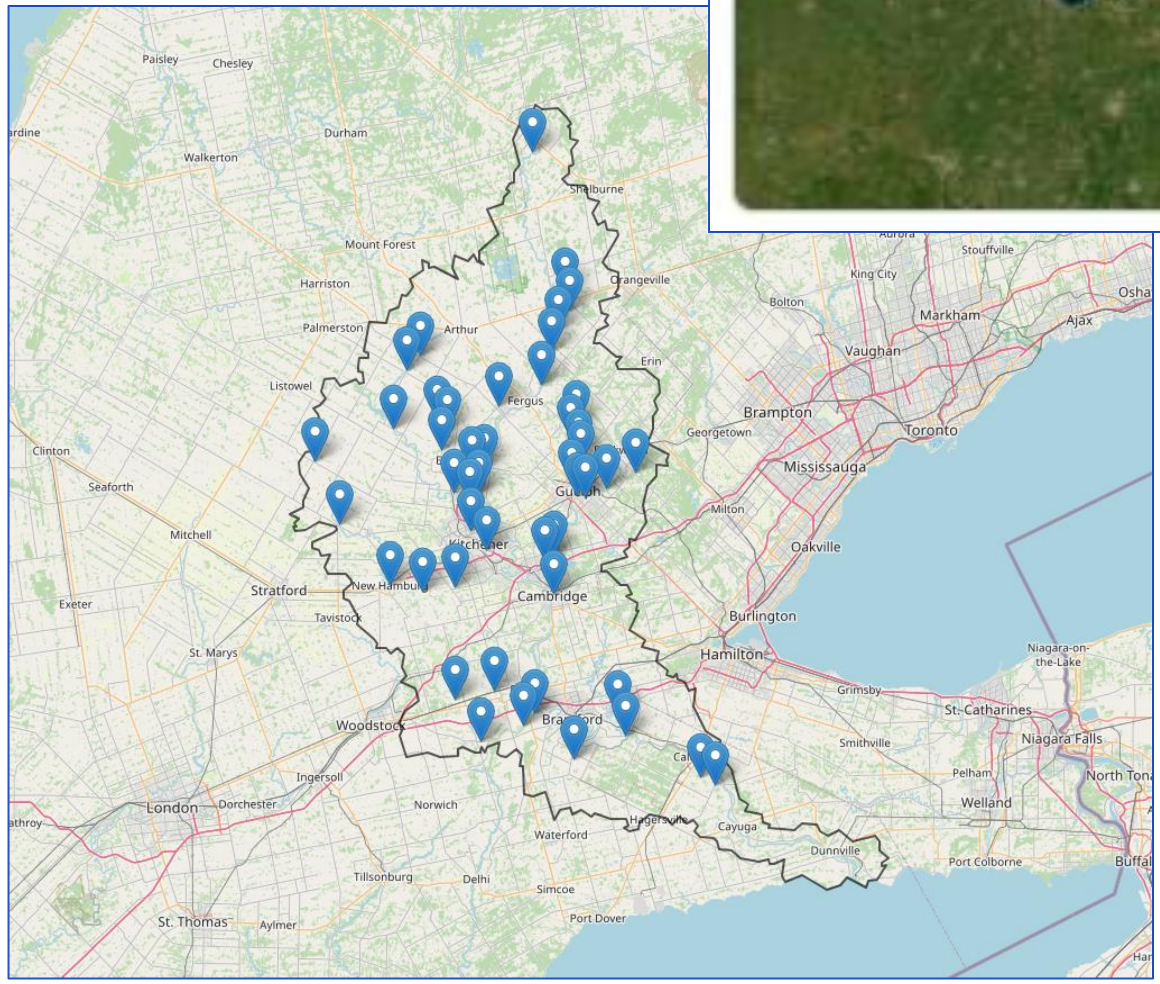

Active Participation by CAs

- 19 CAs are participating and actively using the forecasts

Daily ML-based Forecasts

- Fully automated processing
- 16 day forecasts based on GEPS

HydroClimateSight

aquanty

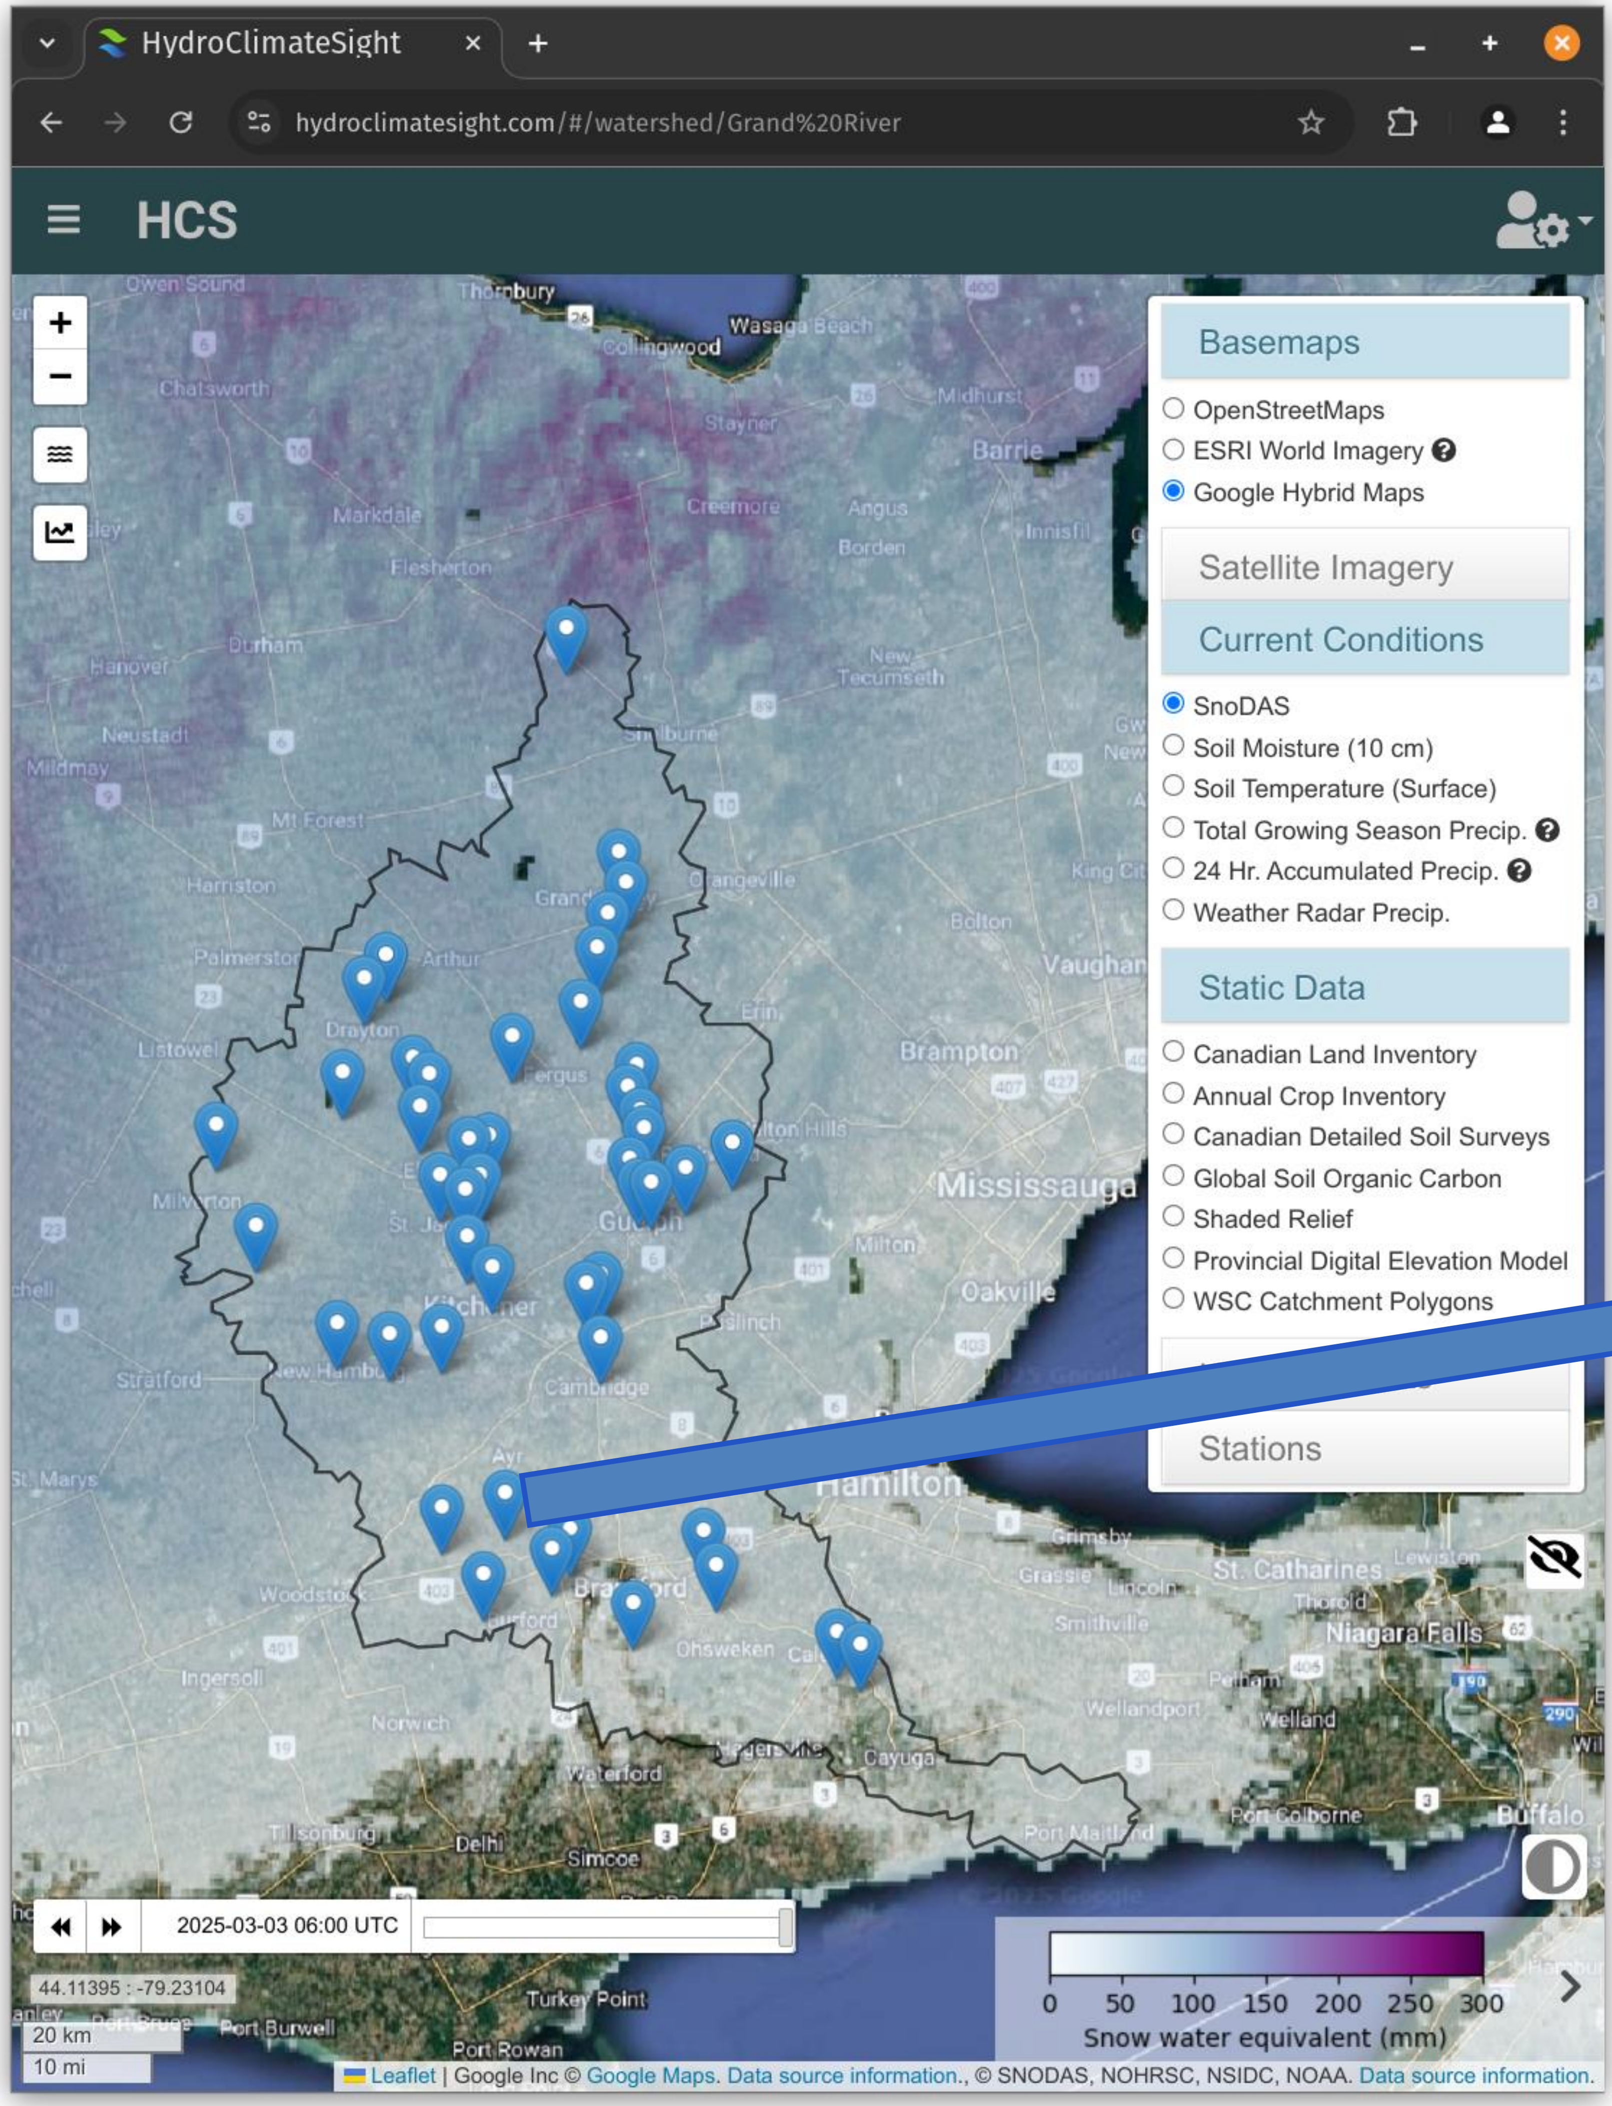

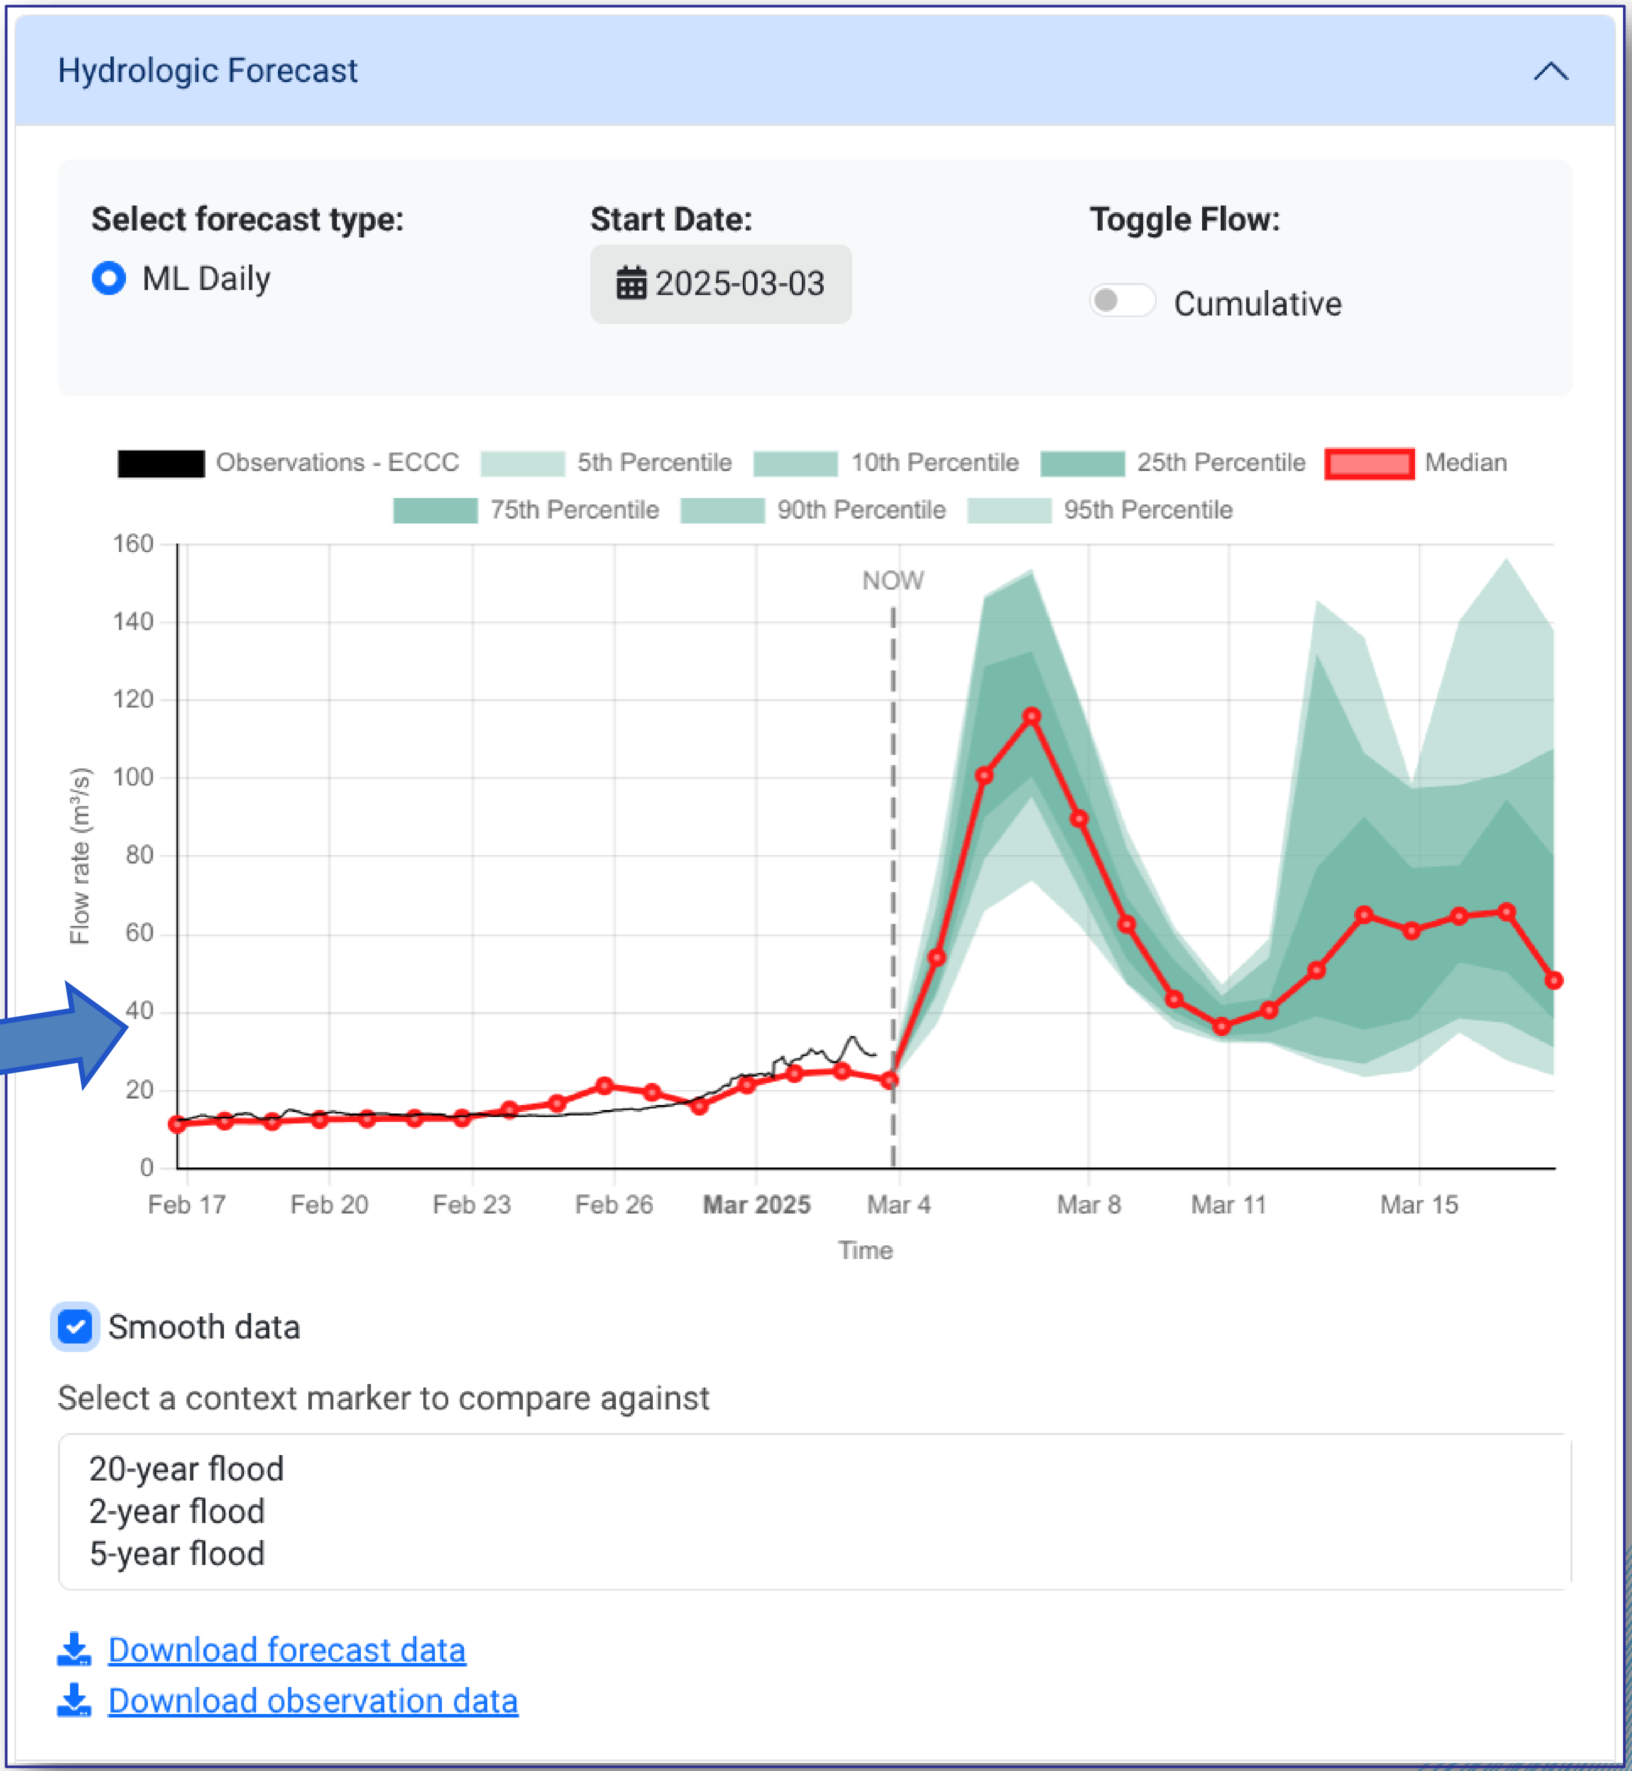

Forecast Initialization

HRDPA+HRDPS

GEPS

Model → Hindcast → Today → Model → Forecast →

Data Integration (assimilation) can improve spin-up significantly!

- 1-year spin-up forced by CaPA and HRDPS analysis from ECCC
- Forecasts from GEPS (ensemble), HRDPS, HRRR, NAM and CFSv2 (ensemble) supported

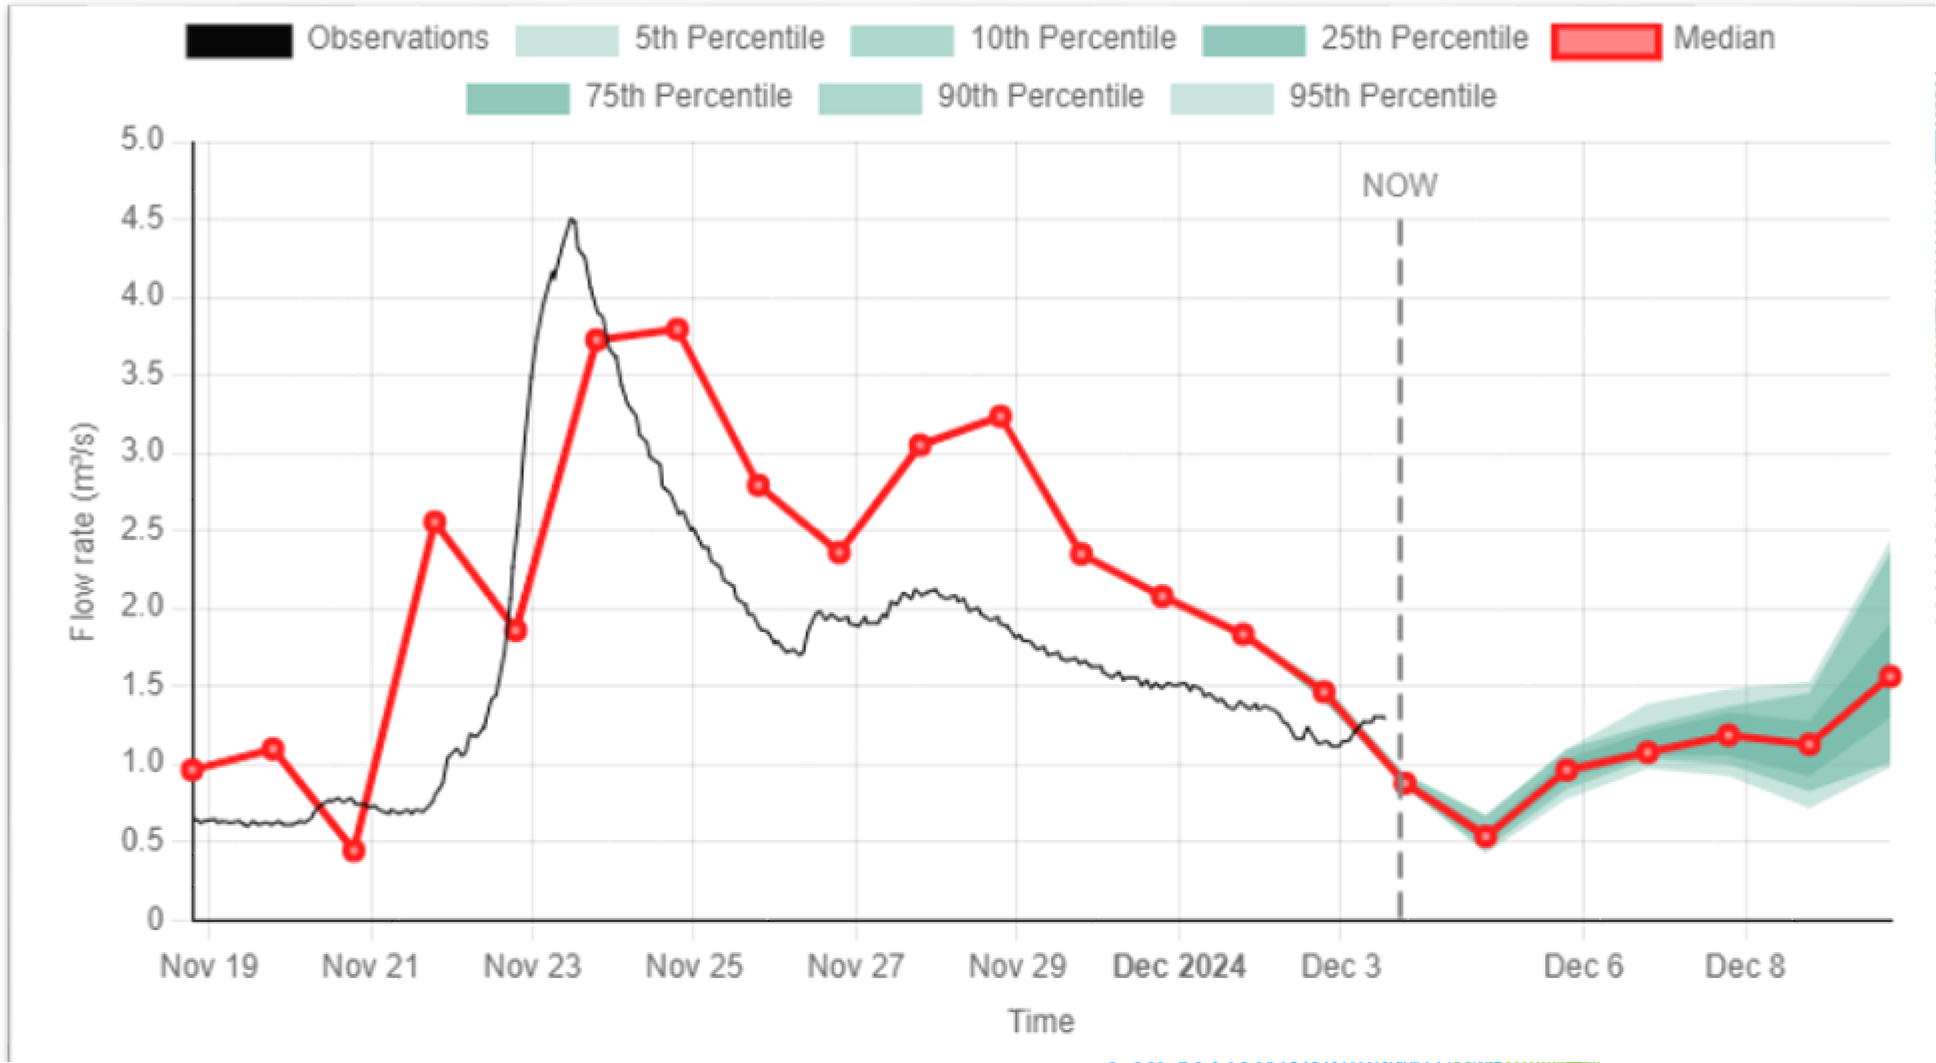

From Training to Forecast

Seamless weather and climate data integration from calibration, over near-real-time spin-up to actual (ensemble) forecast

Historical Weather & Climate Data

- Model training and evaluation
- Historical reanalysis
- e.g. CanGrid, ERA5(L)

Near-real-time Weather Data

- Model spin-up
- Data Assimilation
- Digital Twin
- e.g. CaPA, HRDPS Analysis

Numerical Weather Forecast

- Driving hydrological forecasts
- … including ensembles
- e.g. HRDPS, GEPS, CFS

Data Ingestion and Processing

A generic scheduler automatically launches workflows

Robust data downloads

- CaPA
- HRDPS
- CalDAS
- SnoDAS
- ERA5
- WSC
- Kisters
- Custom data sources

Data transformation

- QA/QC
- Extract variables
- Infilling
- Reprojection
- NetCDF/Zarr conversion
- Cleanup

Storage

- Local SSD for immediate usage
- Local NAS (HDD) for longer term storage + backups
- Time series database
- Cloud blob storage

Weather Data Processing

aquanty

Weather forecast data processing is highly modular, supporting multiple processing steps (right):

- Long-range forecasts are statistically downscaled and bias-corrected
- Forcing for physical models is lapse-rate adjusted using high-res DEM (critical for snow model)
- ML Model forcing is aggregated to HRU and daily (except hourly models)
- Near-real-time spin-up forcing data follows the same pipeline

Download
- Download grib files from ECCC
- Convert/compute variables

Reprojection
- Clip, reproject & regrid to multiple areas of interest

Bias-correction
- Simple statistical bias-correction
- Lapse-rate adjustment using DEM

Export
- Generate, e.g. geotiff or zarr
- Aggregate over polygons
- Resample in time (e.g. daily)

Interoperability and Scalability

- Globally scalable using either cloud, or cloud + on premise resources
- Focus on integrating with the tools that are already being used
- **Web API access**

Direct Wiski imports (ZRXP)

GIS imports

Direct data transfer

HEC-RAS format (DSS)

Delft FEWS imports

Responsive web design

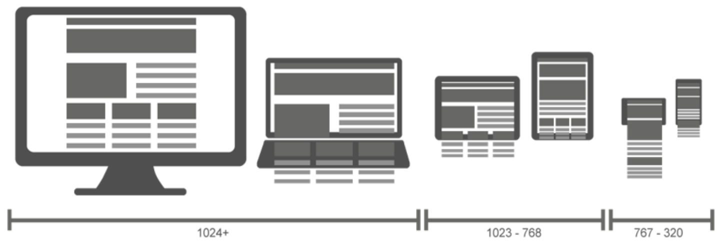

Questions & Roundtable Discussion