OCEANS PROTECTION PLAN

Developing machine-learning emulators for harbour-scale ocean prediction

Michael Dunphy, Maxim Krassovski
Institute of Ocean Sciences, Fisheries and Oceans Canada, Sidney, BC

14 April 2026

Government of Canada Gouvernement du Canada

Canada

Emulators for port models

- We have six NEMO-based port models across Canada that support electronic navigation and drift prediction
  - Tidally dominated with a complex coastline, and most include a major river

- Each port is the 5th or 6th level of dynamic downscaling across ECCC and DFO, eg:

GIOPS →	RIOPS →	CIOPS-W →	Salish Sea→	**South Salish Sea →**	**Vancouver Harbour**
1/4°	1/12°	1/36°	500m	**150m**	**20m**

- Current effort is to explore the art of the possible with machine learning emulators
  - *This is early exploration / work-in-progress*

Emulators for port models

Daily model runs

- 1 x 24h pseudo-analysis: 1 day/day
- 4 x 48h forecasts: 8 days/day
- Vancouver Harbour uses 6x core-hours as South Salish Sea

Exploration 1: emulate VH20

- Retains NEMO-based solution at 150 m
- Could save 6/7ths of the compute

Exploration 2: emulate forecasts

- Retains NEMO-based initial condition for forecasts
- Could save 8/9ths of the compute

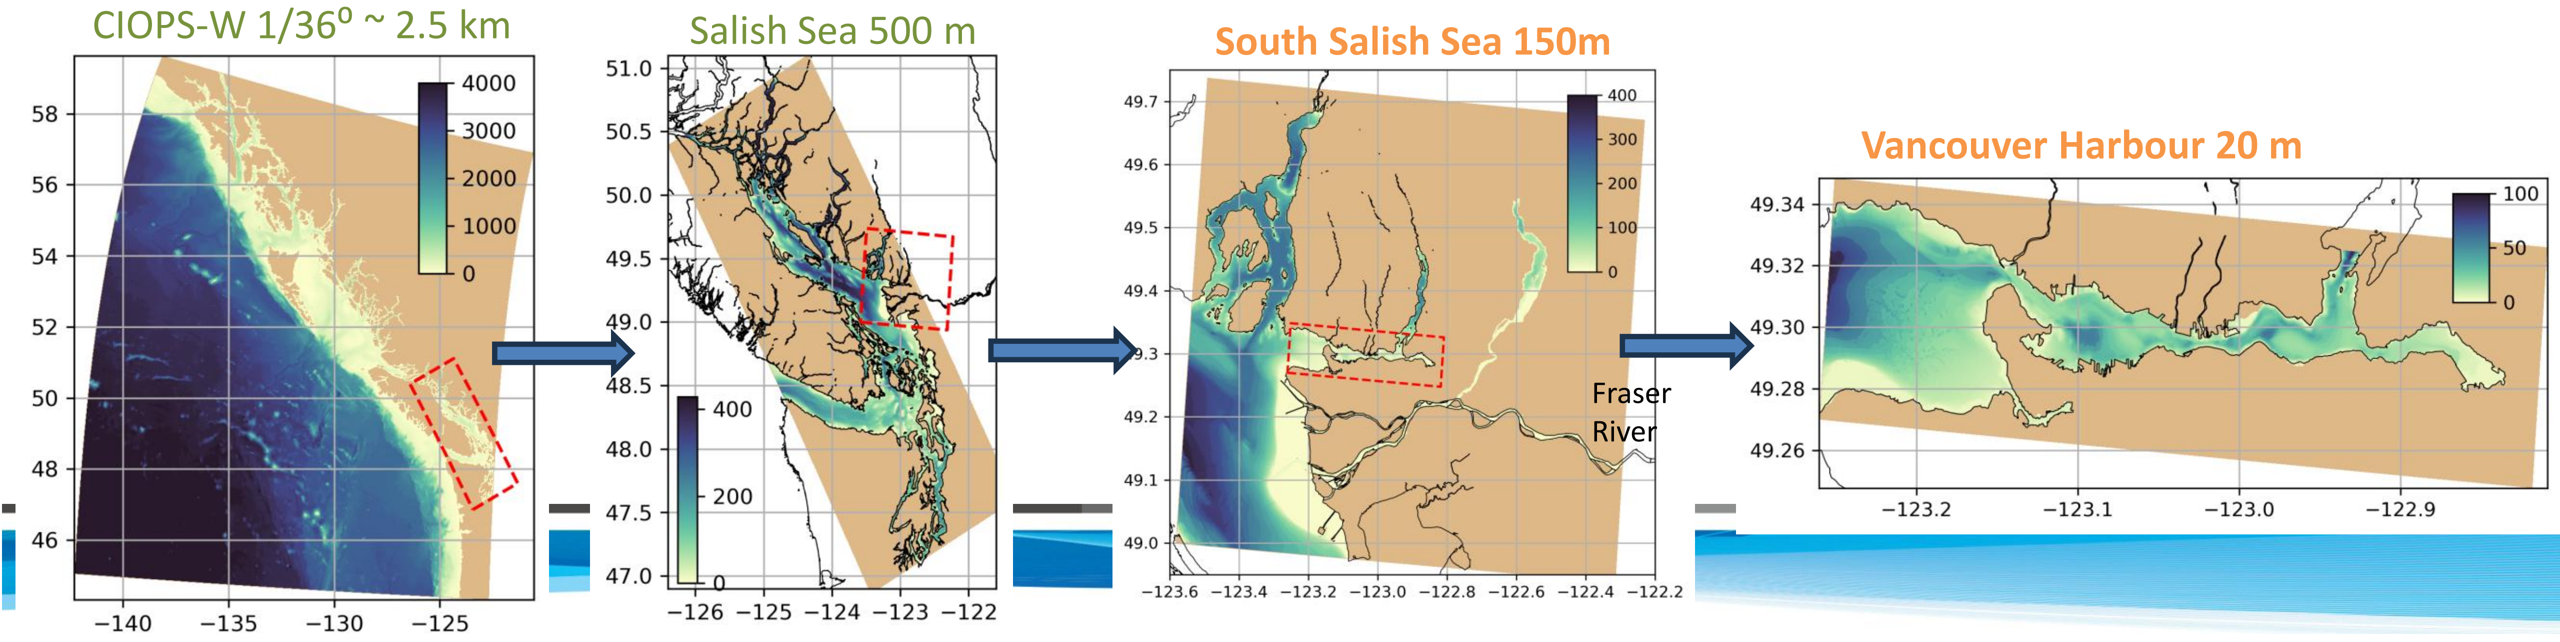

- Exploration 1: VH20 from SSS150 via super-resolution
  - SSH only
- Exploration 2: VH20 from SS500 via neural operator
  - Barotropic mode
- We've coarsened by 5x for faster iteration

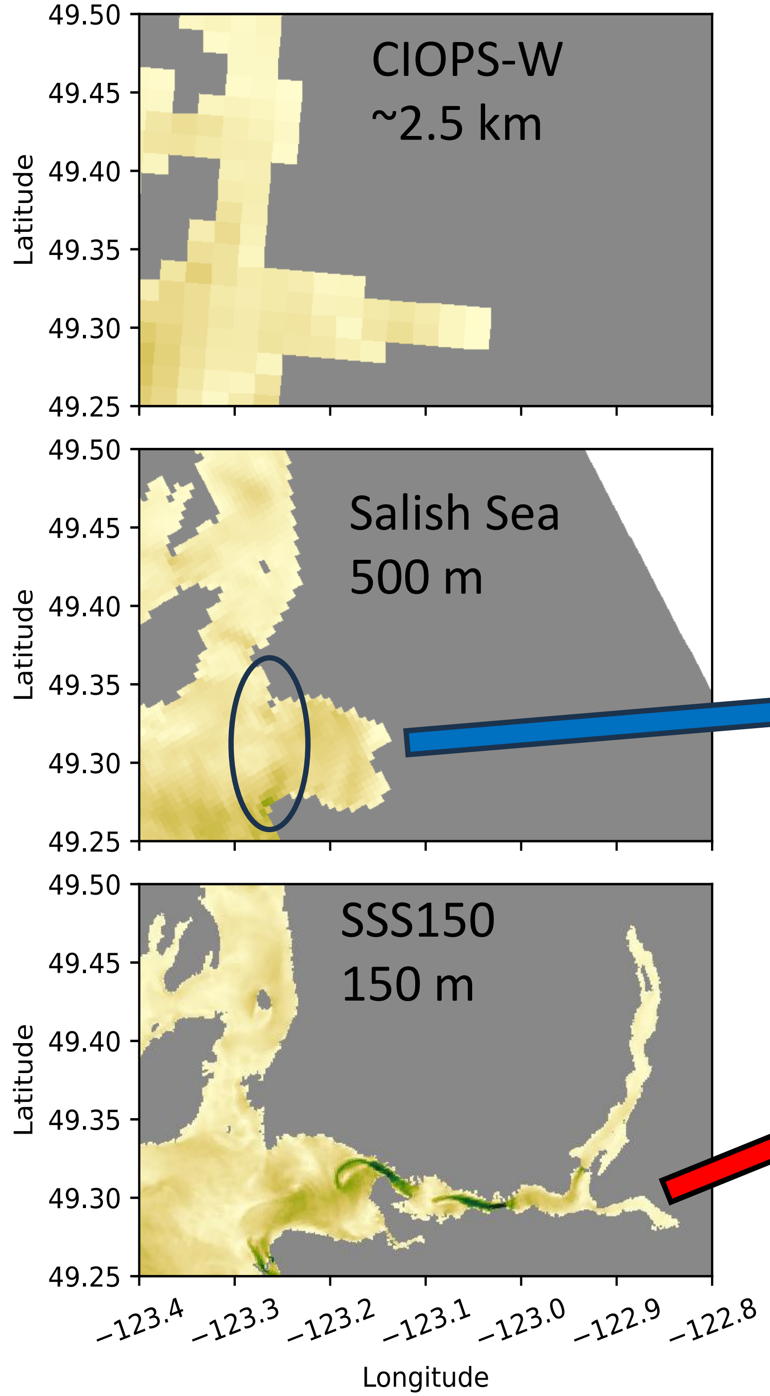

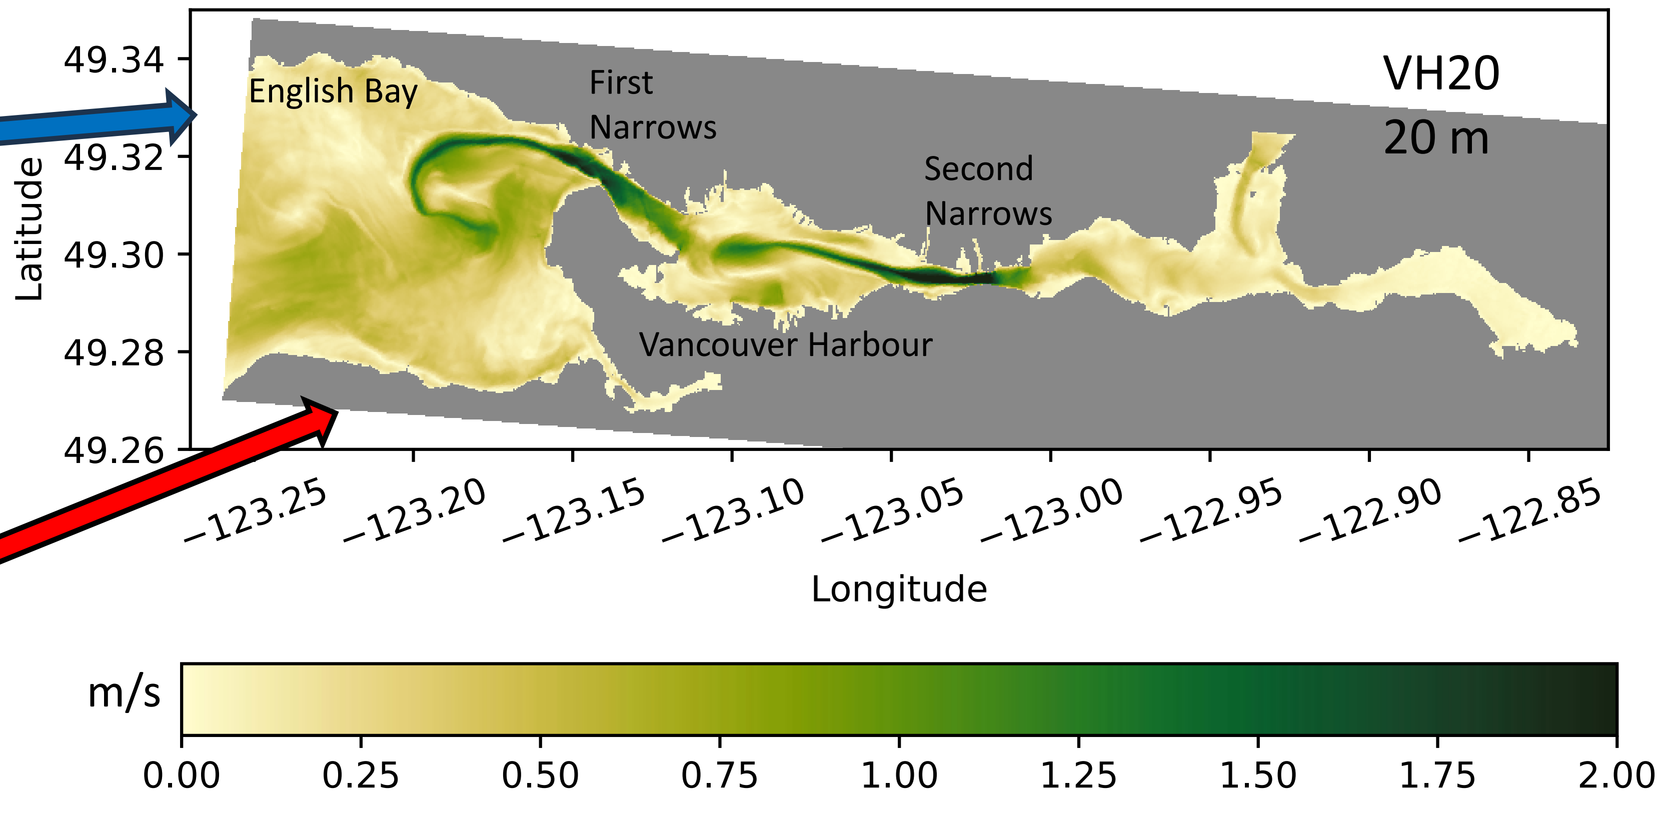

1. Res-U-net

- Emulate VH20 driven by SSS150 per snapshot
- Input channels, 6 or 7:
  - land mask, bathy (20m), upscaled bathy (150), Coriolis
  - Hourly ssh, ssh_ib
  - **Tide**
- Output: ssh hourly @ 20m
- Loss function MSE on ocean points

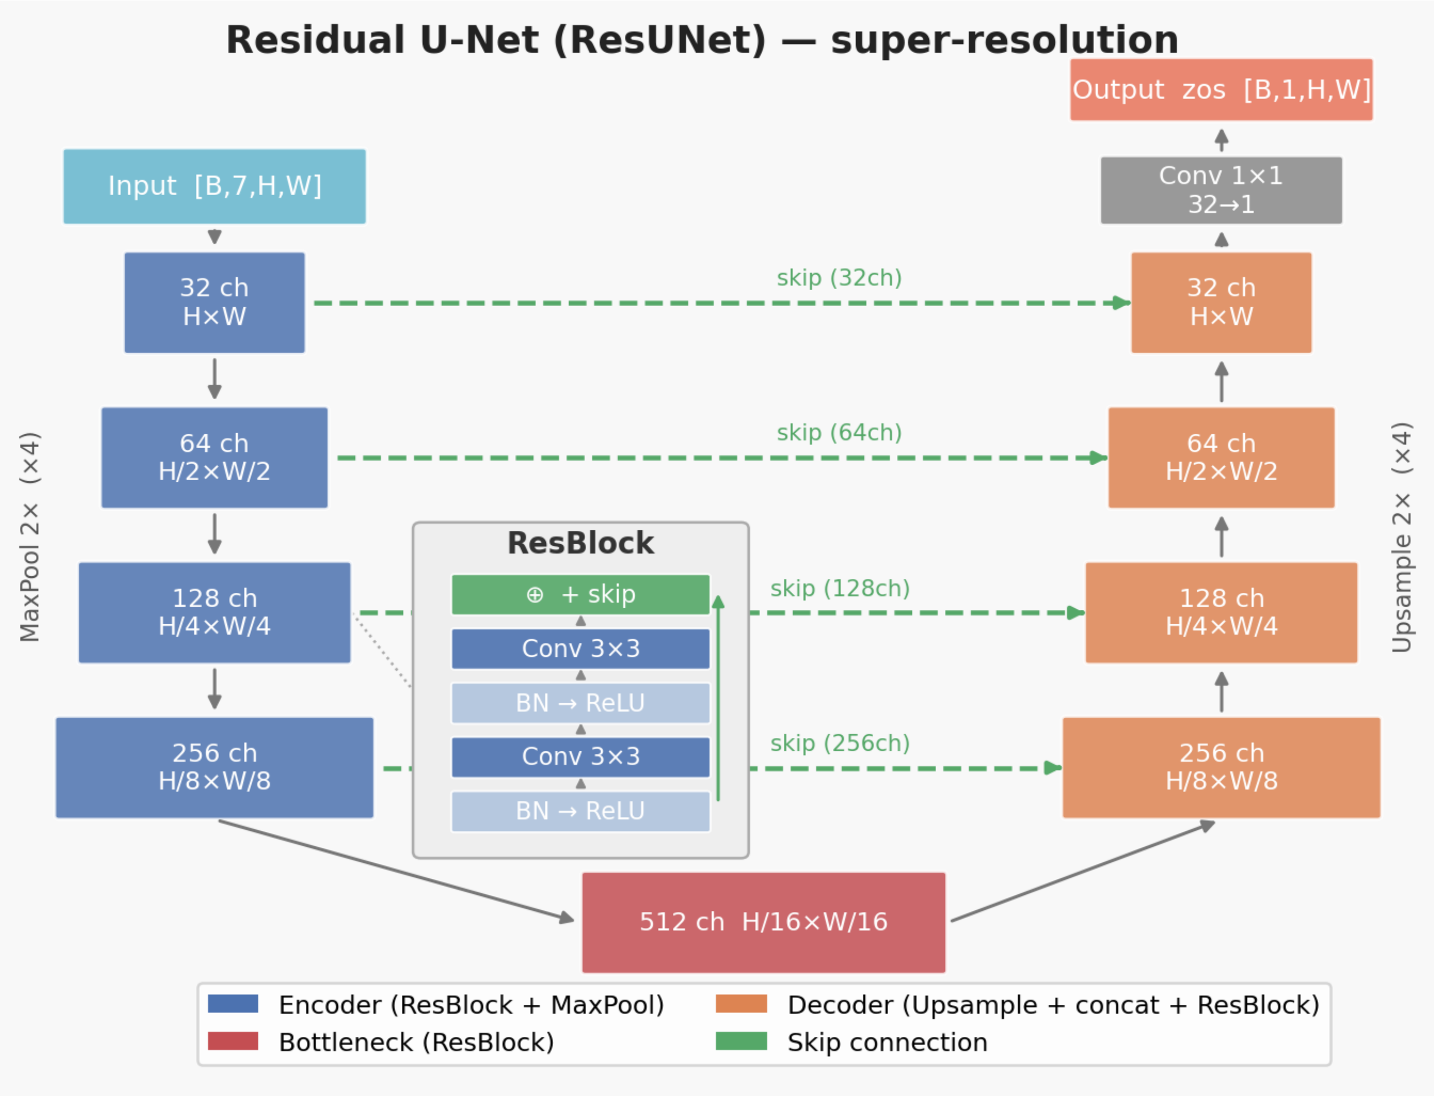

1. Res-U-net - results

Test	X5 rmse	X1 rmse
No ML (interpolation)	19.0 mm	19.5 mm
SSH	4.5 mm	
SSH + 3h lookback	**3.8 mm**	17.6 mm
Detided SSH	13.7 mm	
Tide as 7th channel	19.2 mm	

At x5 coarsened,

- Full SSH with 3h lookback is the best run
- Detiding did not help, nor did providing tide as an extra channel

At native resolution,

- Larger error due to tide phase mismatch between SSS150 and VH20 at east end of model
- ResUNet at x5 works around this with a physically wide receptive field; the RF narrows at native resolution
- ResUNet seems poorly suited to learn how to adjust a broad phase mismatch

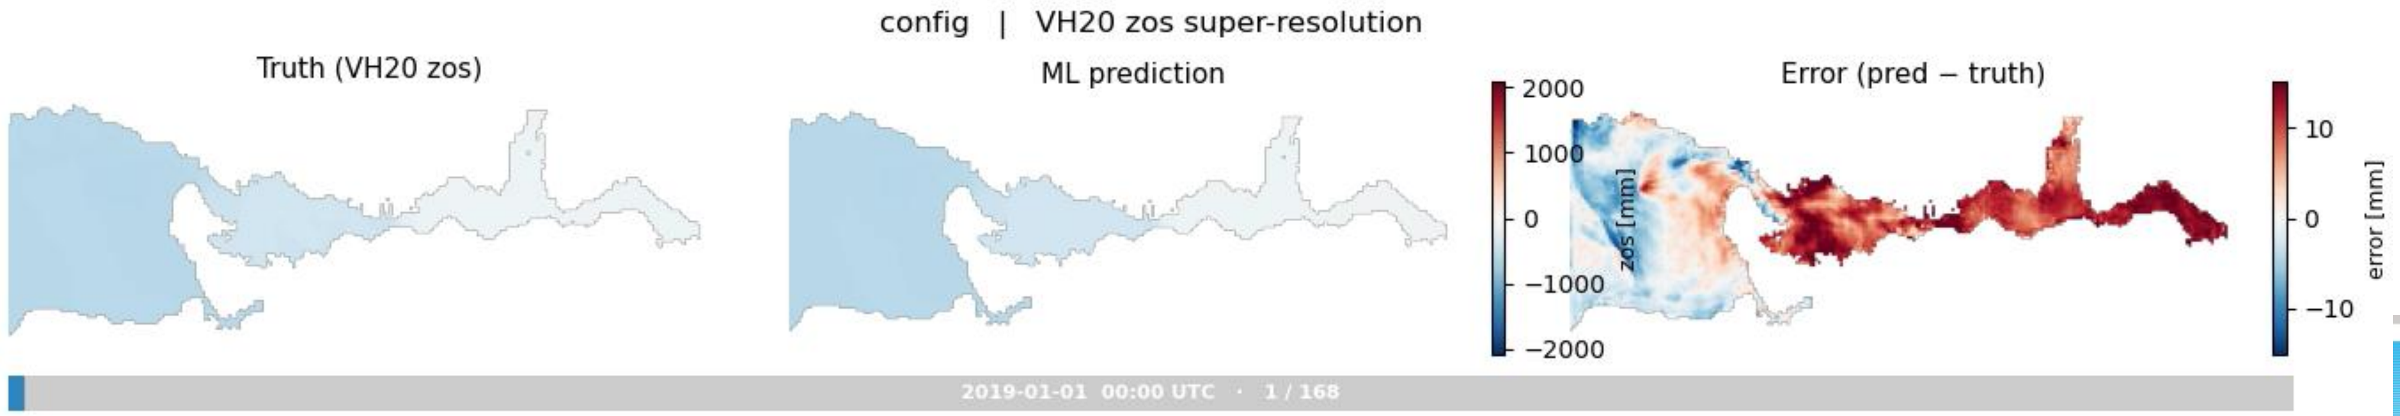

1. Res-U-net - tidal analysis

- Tidal maps look reasonable
- Some differences at east end
- M2 scale 91-98 cm
- K1 scale 86-91 cm

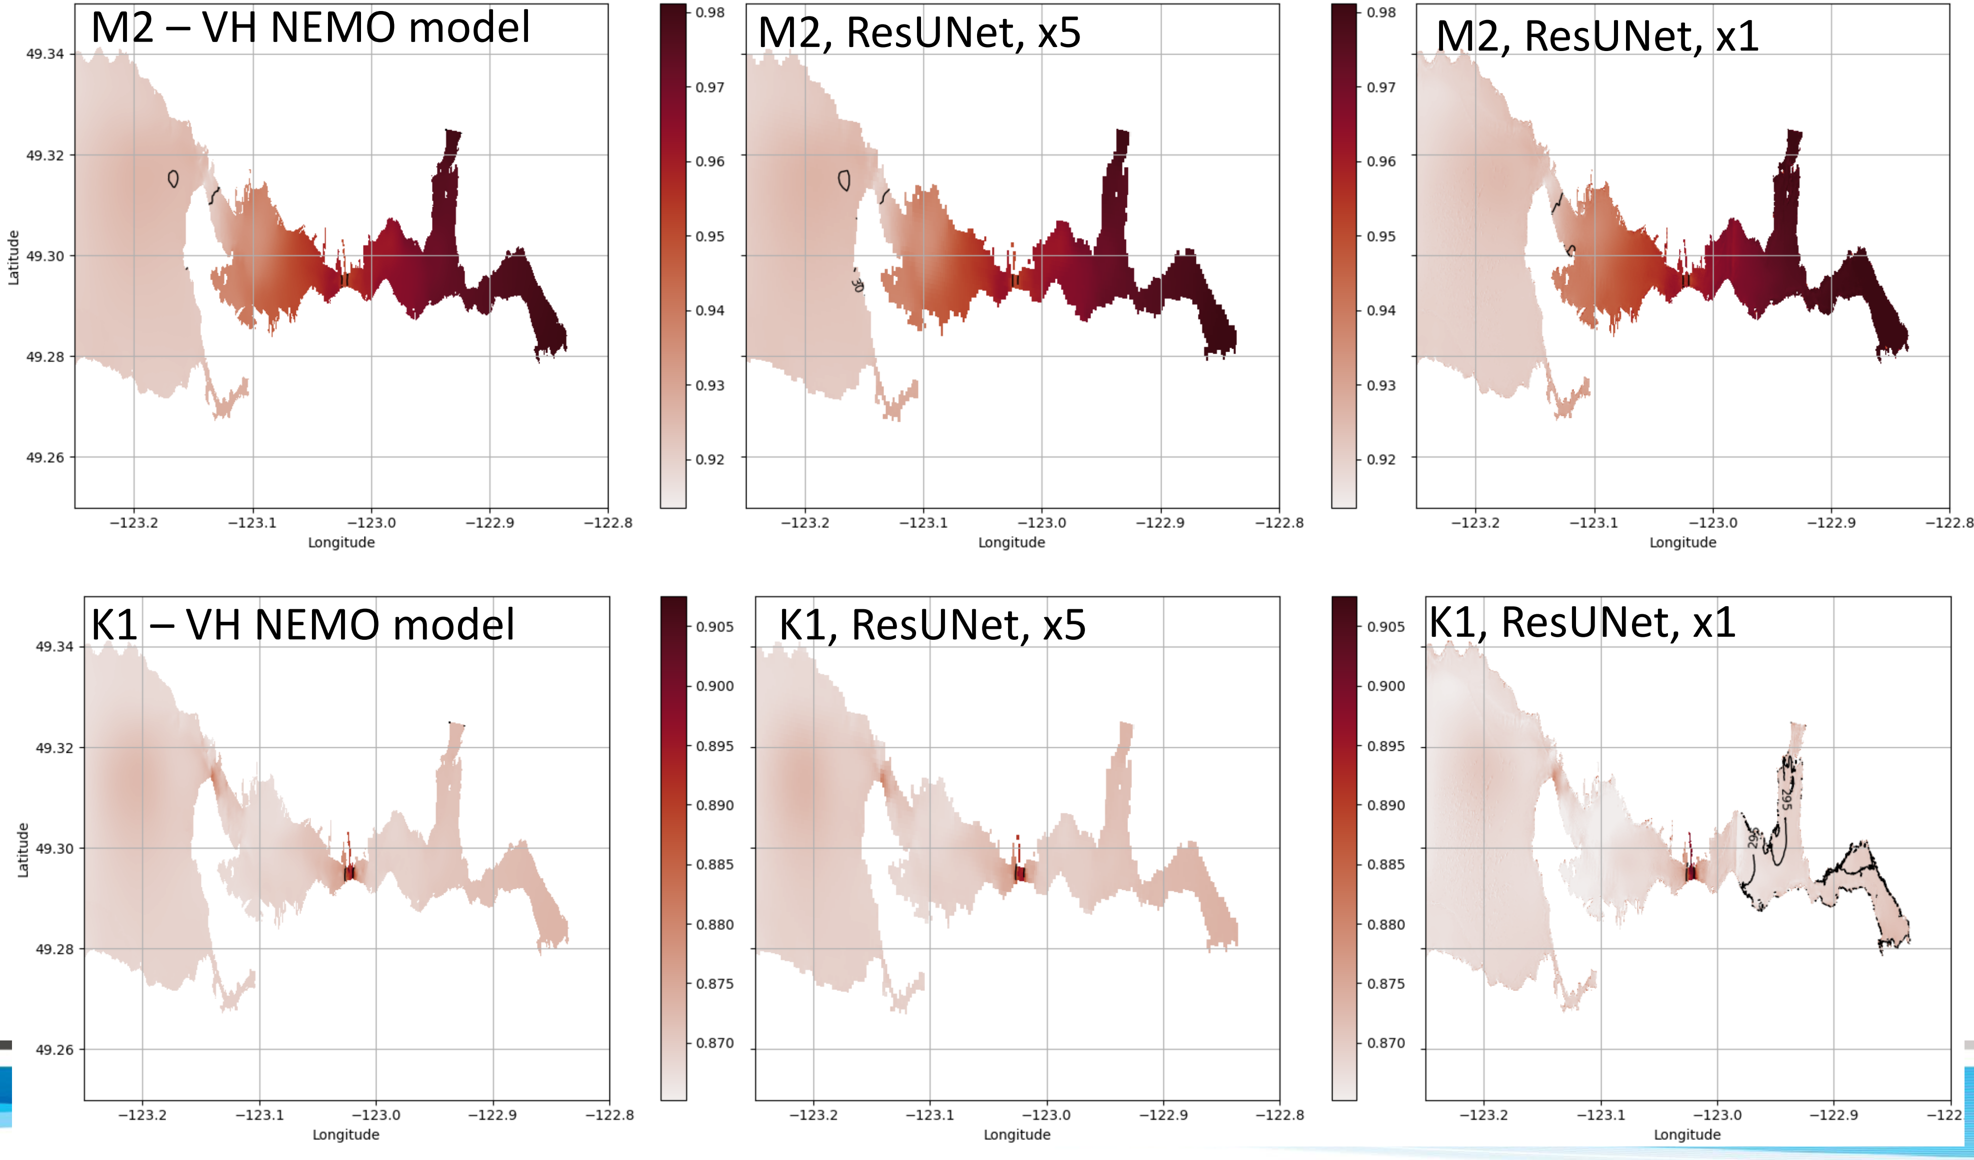

2. Autoregressive {Fourier, Laplace} Neural Operator

- Neural operator learns

  $X^{t+1} = F(x^t, \text{boundary}, \text{surface})$
- Barotropic: ssh, ubar, vbar
- Trained using
  - SS500 patch at West
  - Surface forcing: P, u10, v10
  - VH20 snapshots
- No input from SSS150

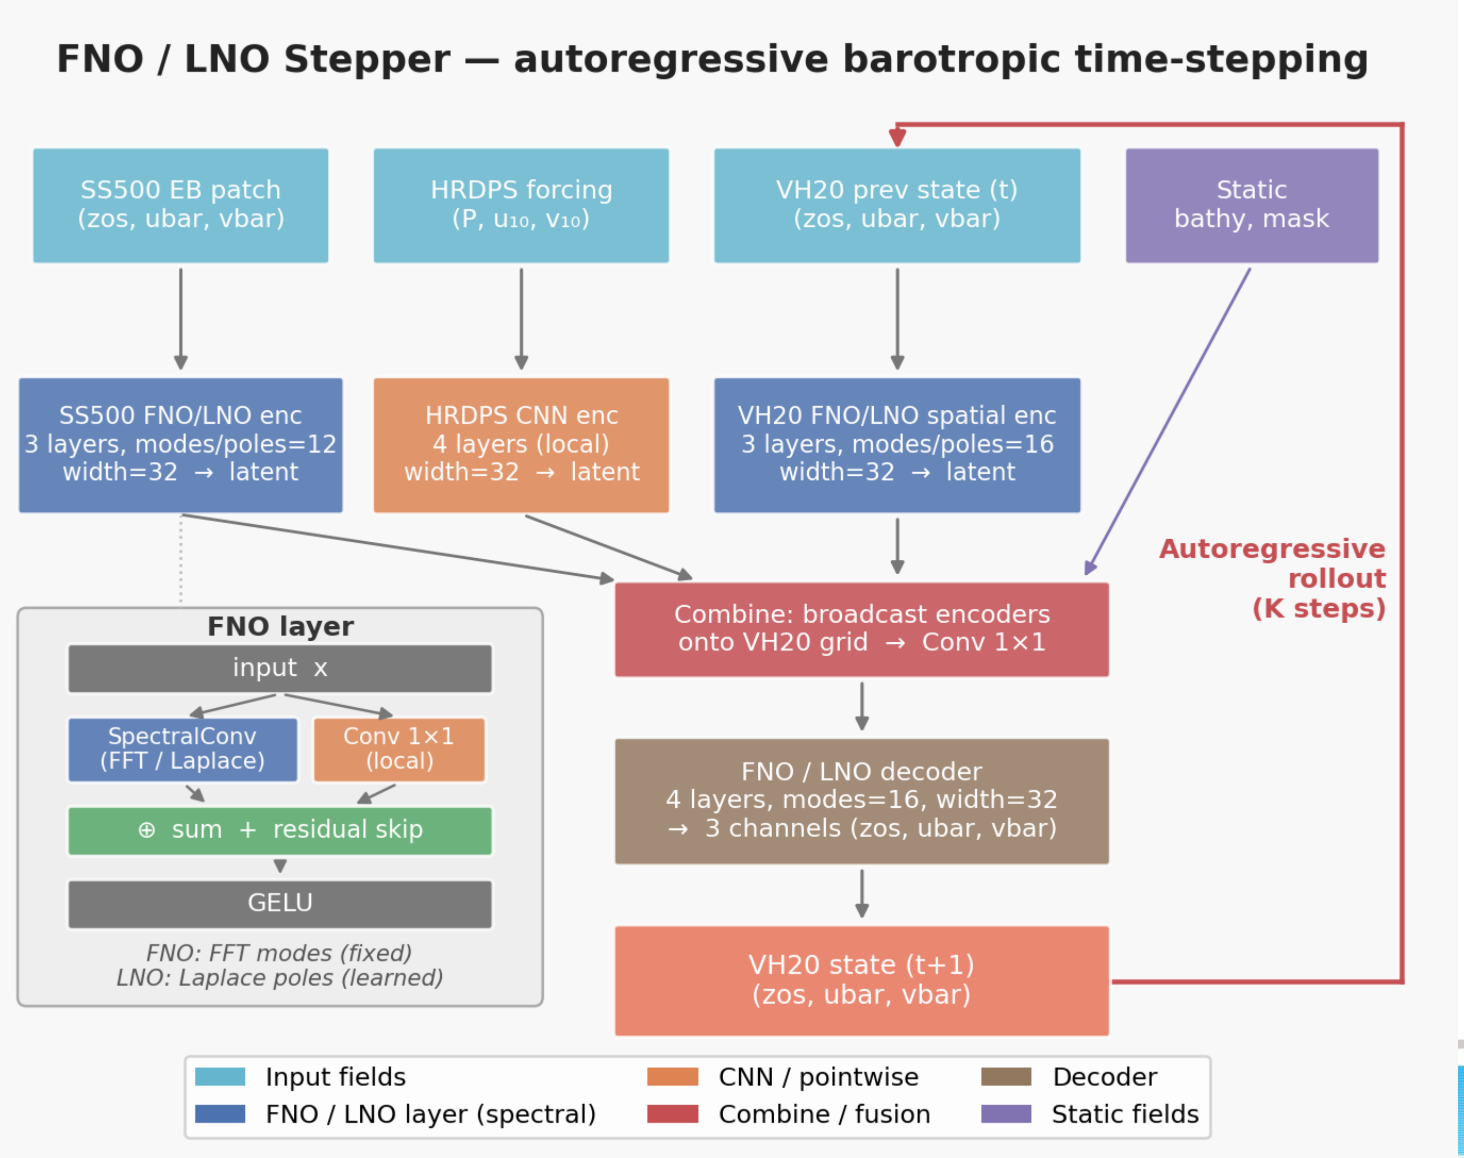

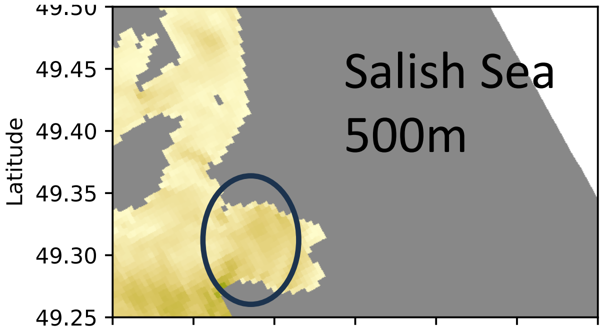

2. FNO vs LNO layer

- Similar design
- LNO has learnable poles instead of fixed Fourier modes
- GELU is Gaussian Error Linear Unit activation function
- Sigma <= 0 enforced

- Performance:
  - LNO forfeits FFT

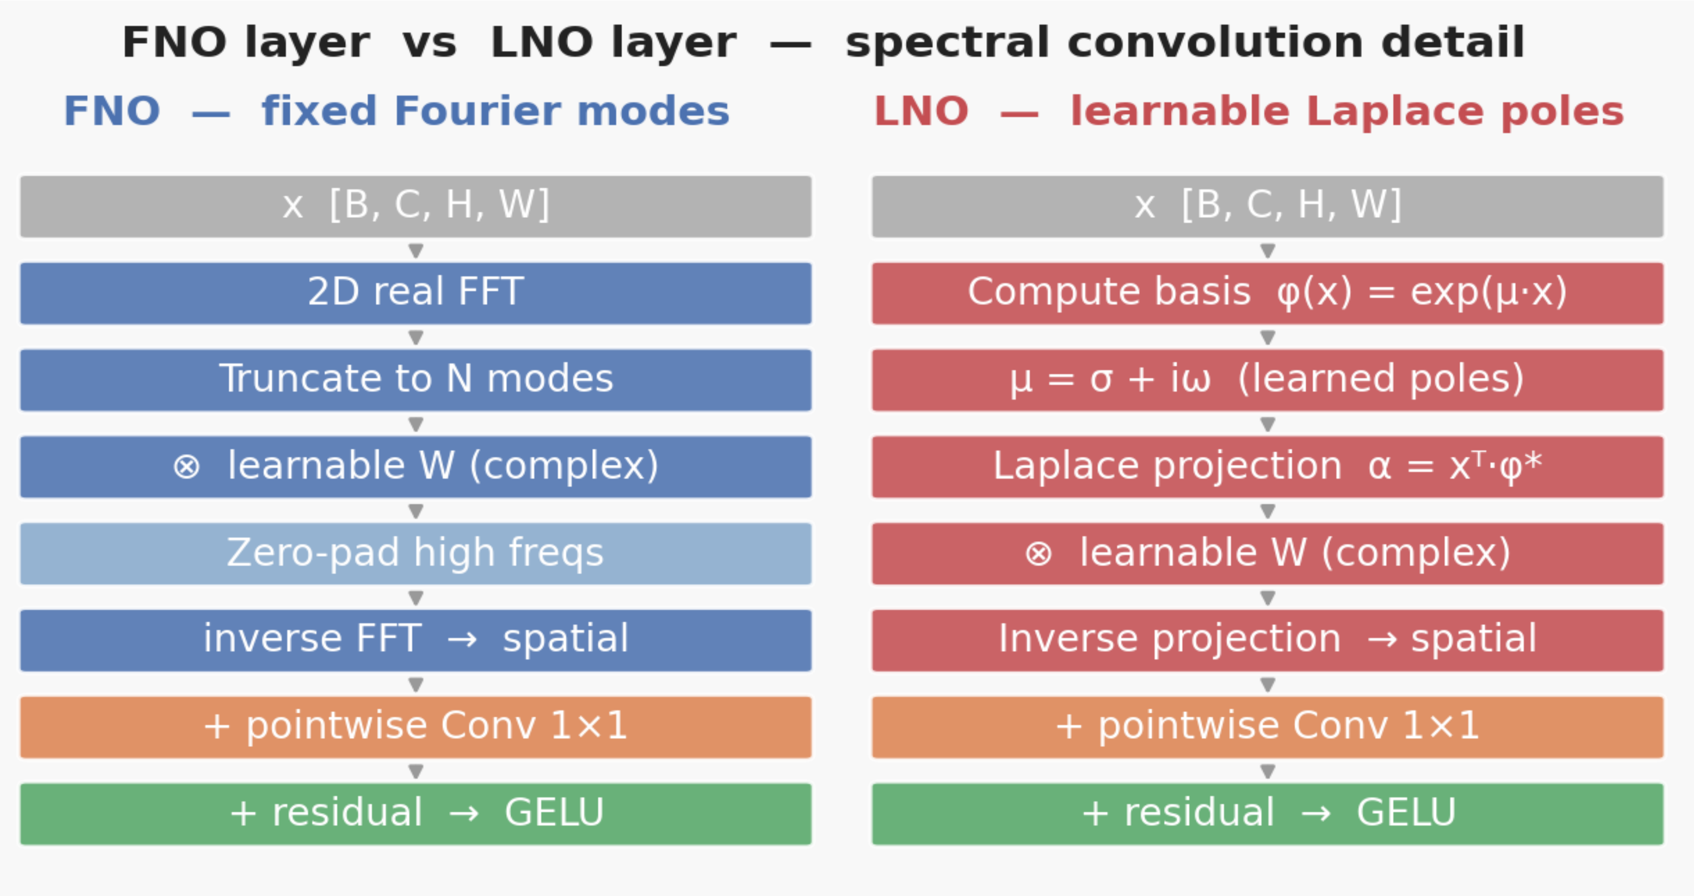

2. FNO Results

Test	X5 rmse	Training
SSH	106 mm	Teacher forcing
SSH	90.6 mm	Rollout training, K=6
SSH – IB	89.2mm	Rollout training, K=6
SSH – IB, 24 modes	94.9 mm	Rollout training, K=6
SSH – IB	83.4 mm	Curriculum rollout K=1..23
SSH – IB, +tide clock	**59.6 mm**	Curriculum rollout K=1..23

- Training mode important; curriculum rollout best
- More Fourier modes (16→24) did not help
- Detiding (not shown) did not help
- Tide clock helped significantly

- Prediction initializes at each 00Z

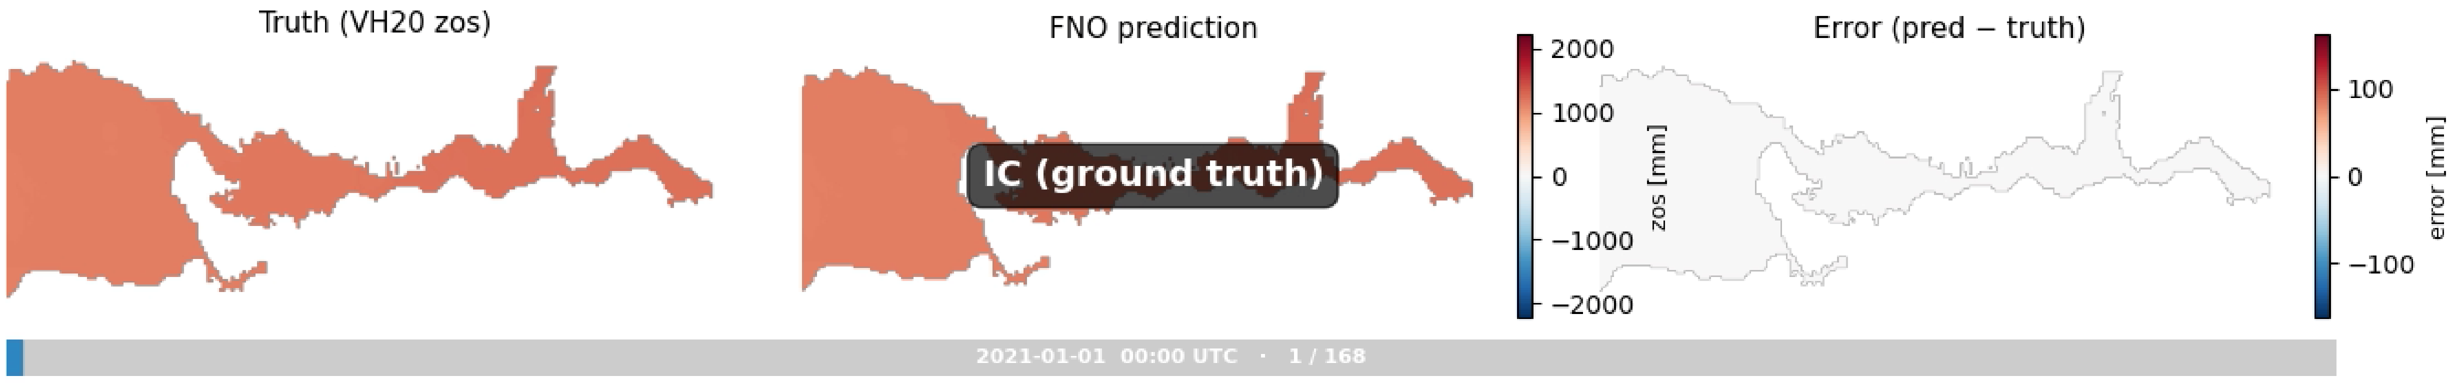

2. LNO results

Test	X5 rmse	Notes
(5b) SSH without SSH_IB	90.5 mm	Rollout training, K=6
(5d) Add tide clock	**61.7 mm**	Rollout training, K=6

- Performance comparable to FNO
- Tide clock helped significantly

- Prediction initializes at each 00Z

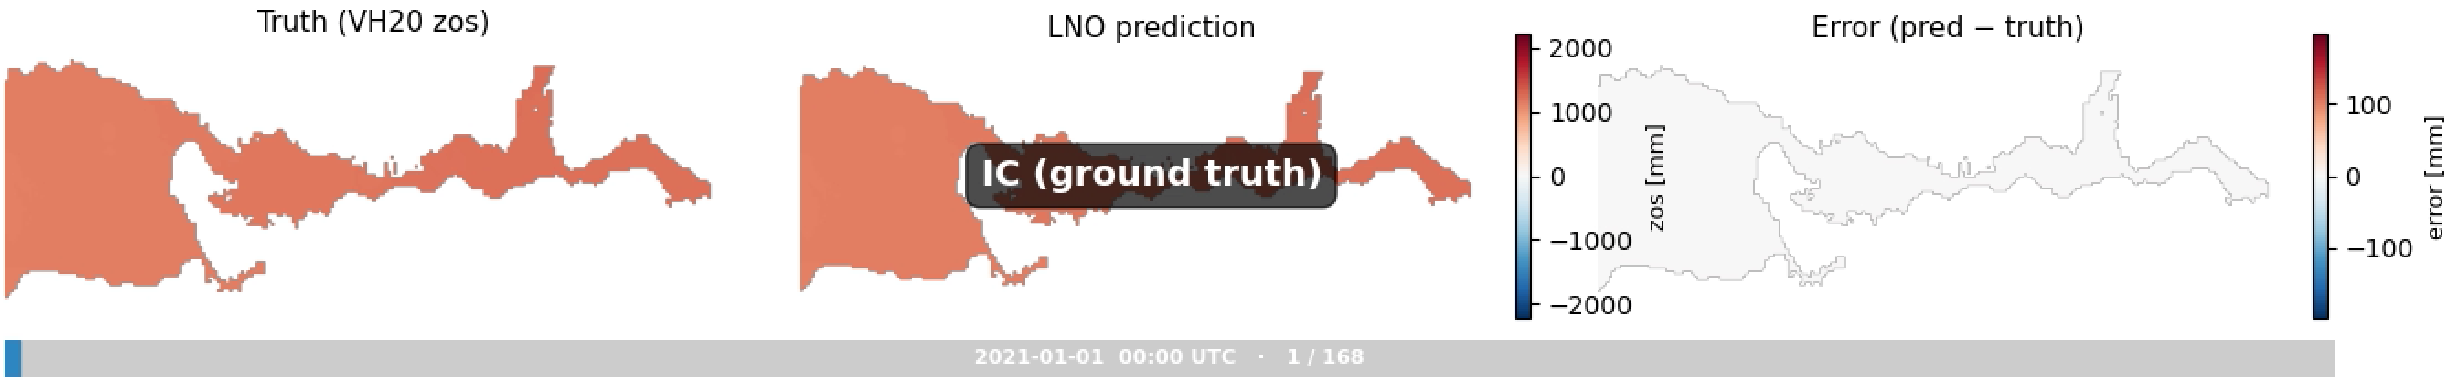

2. FNO & LNO – water level analysis

Characterizing the errors

- Synthesize 1yr of ssh
- Compare to 5 tidal stations

Station	NEMO	FNOx5	LNOx5	NEMO	FNOx5	LNOx5	NEMO	FNOx5	LNOx5
	Non-tidal crmse (mm)			Tidal crmse (mm)			Total WL crmse (mm)		
Sandy Cove	41	58	65	52	61	54	64	83	83
Kitsilano	44	60	67	54	62	60	68	85	88
Ambleside	43	59	65	56	62	60	69	85	87
Vancouver	42	58	66	52	61	54	63	83	84
Port Moody	45	62	67	66	72	69	78	95	95
Average	43	59	66	56	64	59	68	86	87

Non-tidal:
FNO better?

Tidal:
LNO better?

Total: similar

2. FNO and LNO – tidal analysis

- Tidal maps not smooth
- LNO smoother than FNO
- Both noisier than NEMO
- The variation is small, order 1 cm
  - M2 scale 91-98 cm
  - K1 scale 86-91 cm
- Other constituents are similar

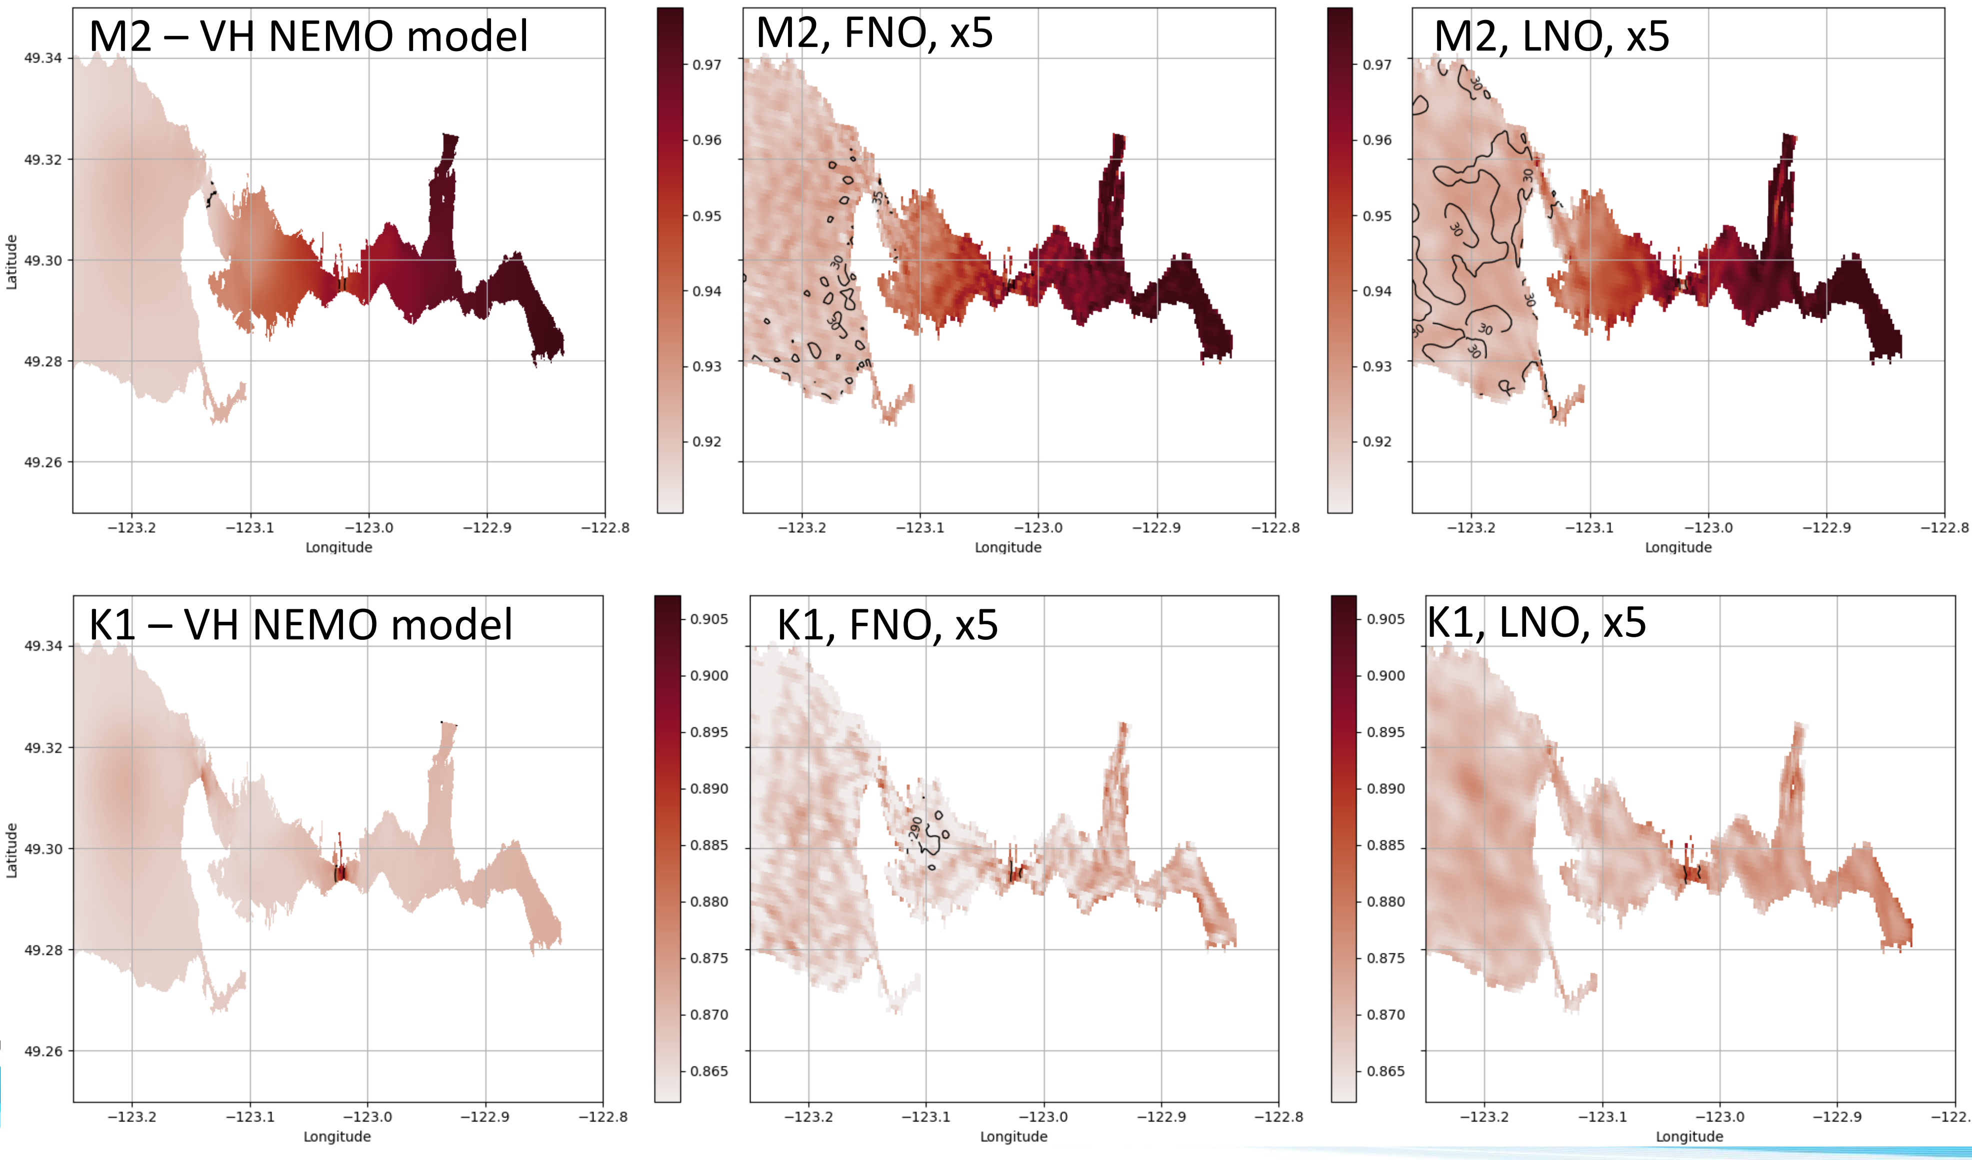

2. FNO and LNO – water level spectrum

- Black is Obs, Green is NEMO, Red is FNO, Yellow is LNO

- Concatenation of 24h rollouts
  - Expect noise at 1 cpd, not seeing it

- Error spectrum for FNO & LNO above that of NEMO for all frequencies

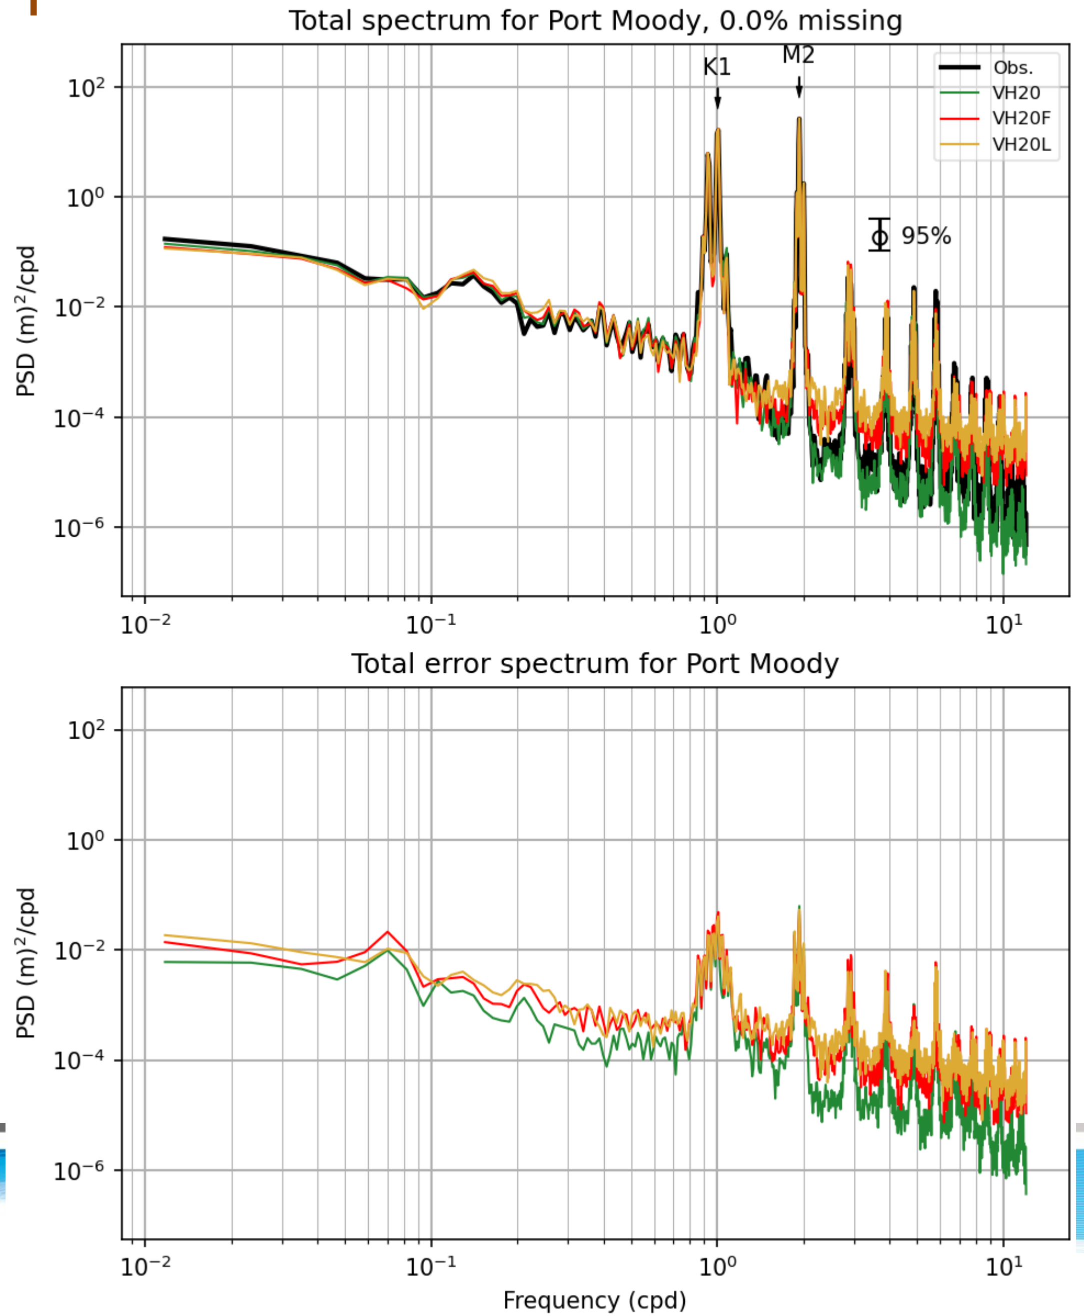

Summary

- Res-U-Net:
  - At x5, reduces approx. 19.5mm to 3.8mm -- bulk of fine signal recovered
  - At x1, rmse is 17.2 mm (partly due tide-phase & receptive field issue)

- FNO & LNO at x5
  - Comparable performance, 59.6 vs 61.7 mm rmse
  - Including tidal clock valuable
  - Tidal constituent maps have “noise”; smoother for LNO than FNO
  - Tentative: FNO better at non-tidal WL, while LNO better at tidal WL

Next steps: refinements

- Res-U-Net:
  - Expand receptive field for x1
- FNO & LNO:
  - Test at full resolution
  - Curriculum training for LNO; more poles
  - Test increment form: $X^{t+1} = x^t + F(x^t, \text{forcing})$ vs. $X^{t+1} = F(x^t, \text{forcing})$
- Loss function:
  - Penalize gradient errors
  - Explore PINN constraints

Next steps: exploration

- Other methods; leading options may be:
  - Graph Neural Operator (via ANEMOI)
  - Shifted Window (SWIN) Transformer
  - Koopman Neural Operator
  - LSTM with EOFs
  - Conv LSTM

- Expand to 3D
  - Surface current important for our application

END