Climate Adaptation-Aware Flood Prediction for Coastal Cities Using Deep Learning

Bilal Hassan, Areg Karapetyan, Aaron Chung Hin Chow, Samer Madanat
New York University Abu Dhabi

NeurIPS 2025 Workshop on Tackling Climate Change with Machine Learning

MOTIVATION

Coastal Vulnerability

- 70% of megacities are coastal
- 90% are vulnerable to flooding
- Sea-level rise could increase flood risk **nine-fold** by 2050

MOTIVATION

Adaptation Measures

- Construction of engineered fortifications (e.g., seawalls)
- These engineered structures change shoreline geometry and hydrodynamics
- Requires careful modeling and simulation

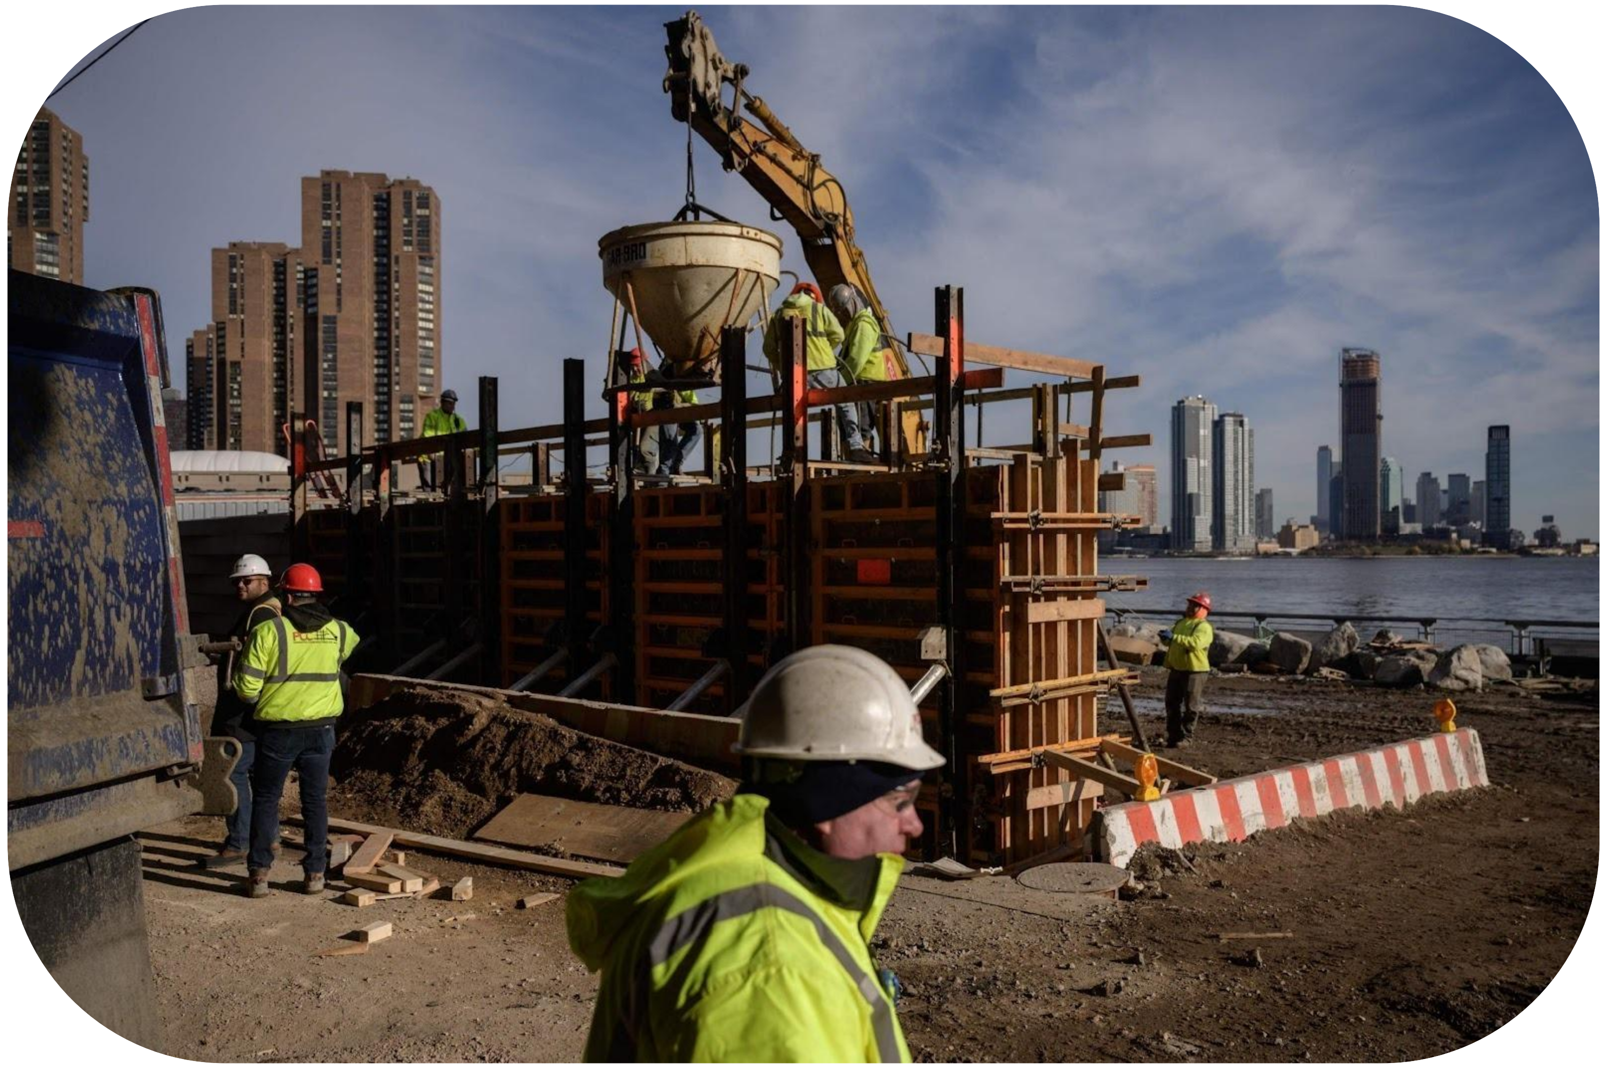

MOTIVATION

Available Tools

- High-resolution modeling is essential to capture urban details
- Physics-based high-fidelity simulators (e.g., Delft3D)
    - Accurate but **extremely slow** and **resource-intensive**

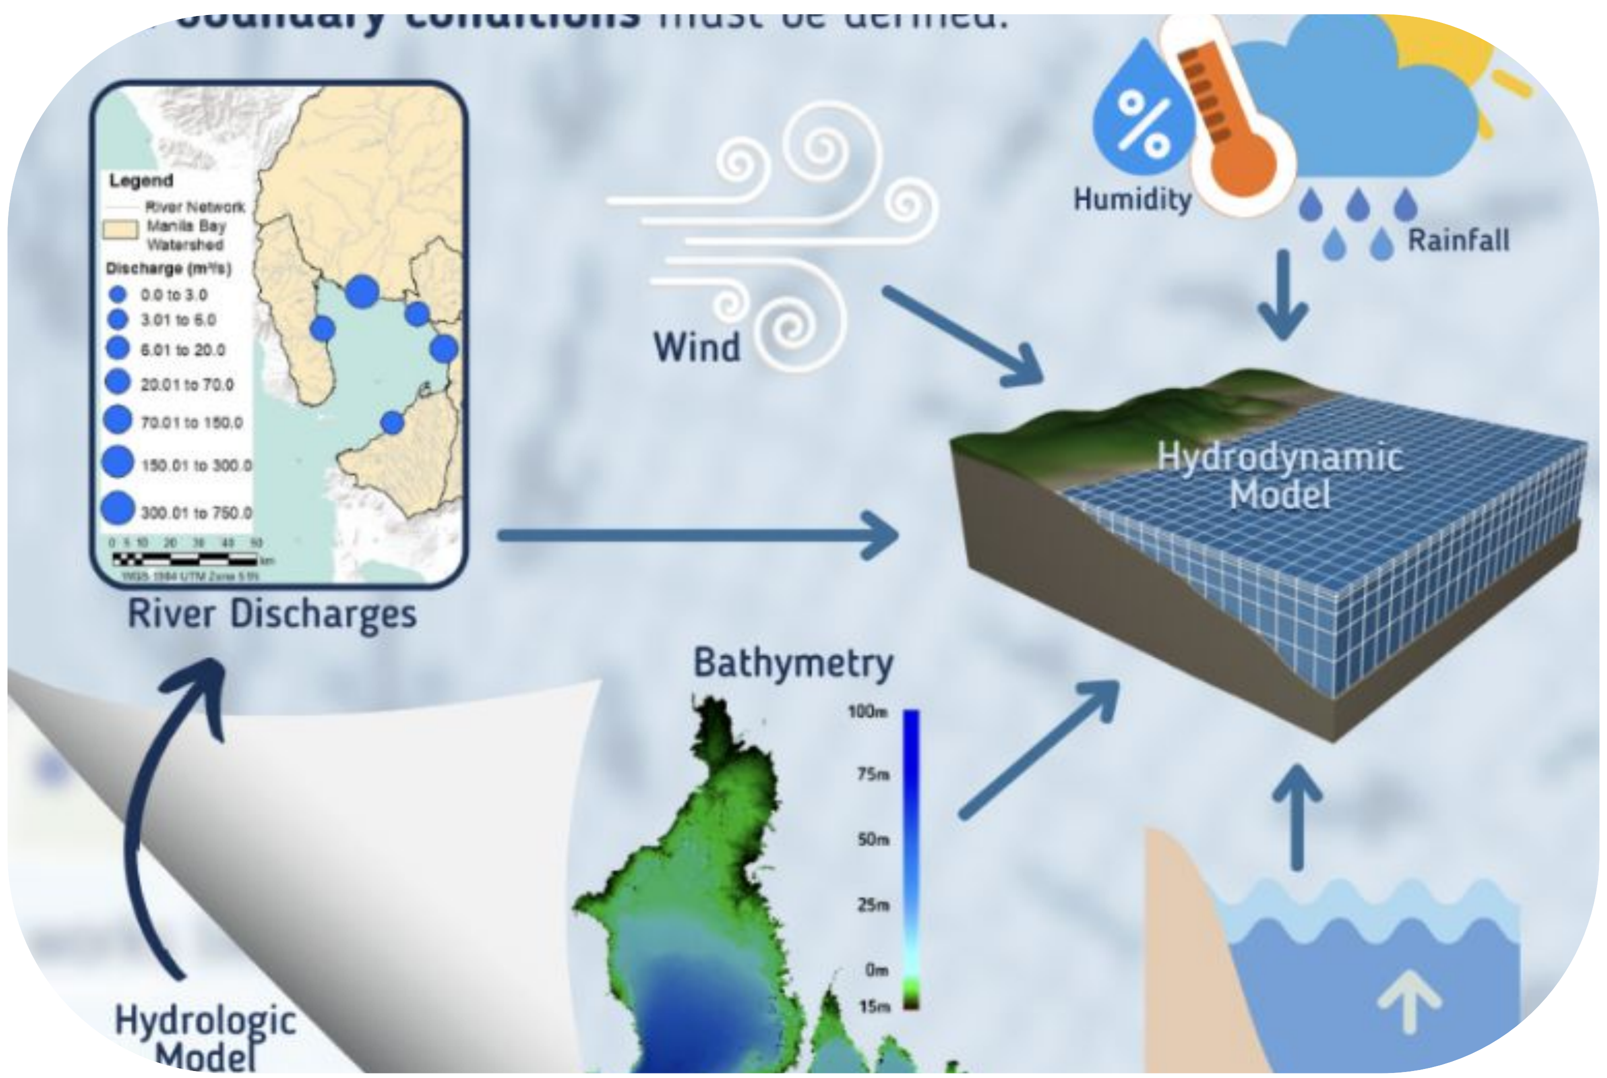

PROBLEM STATEMENT

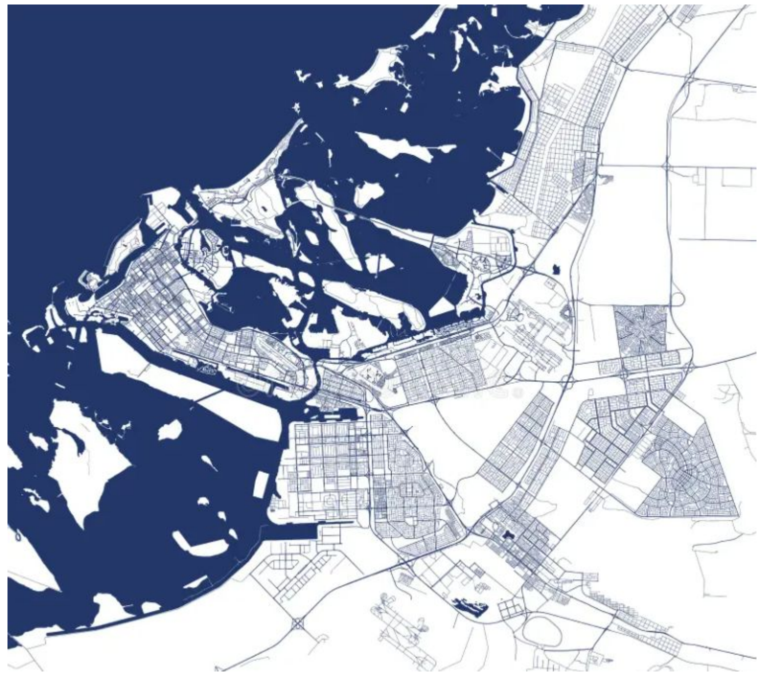

&

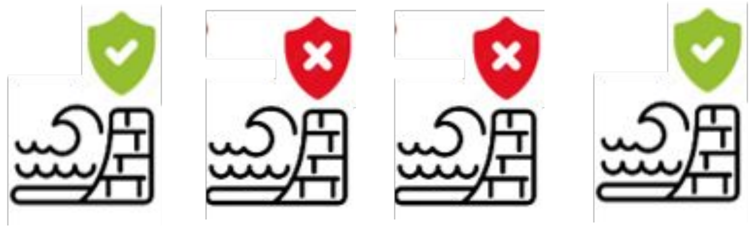

&

SLR Projection

?

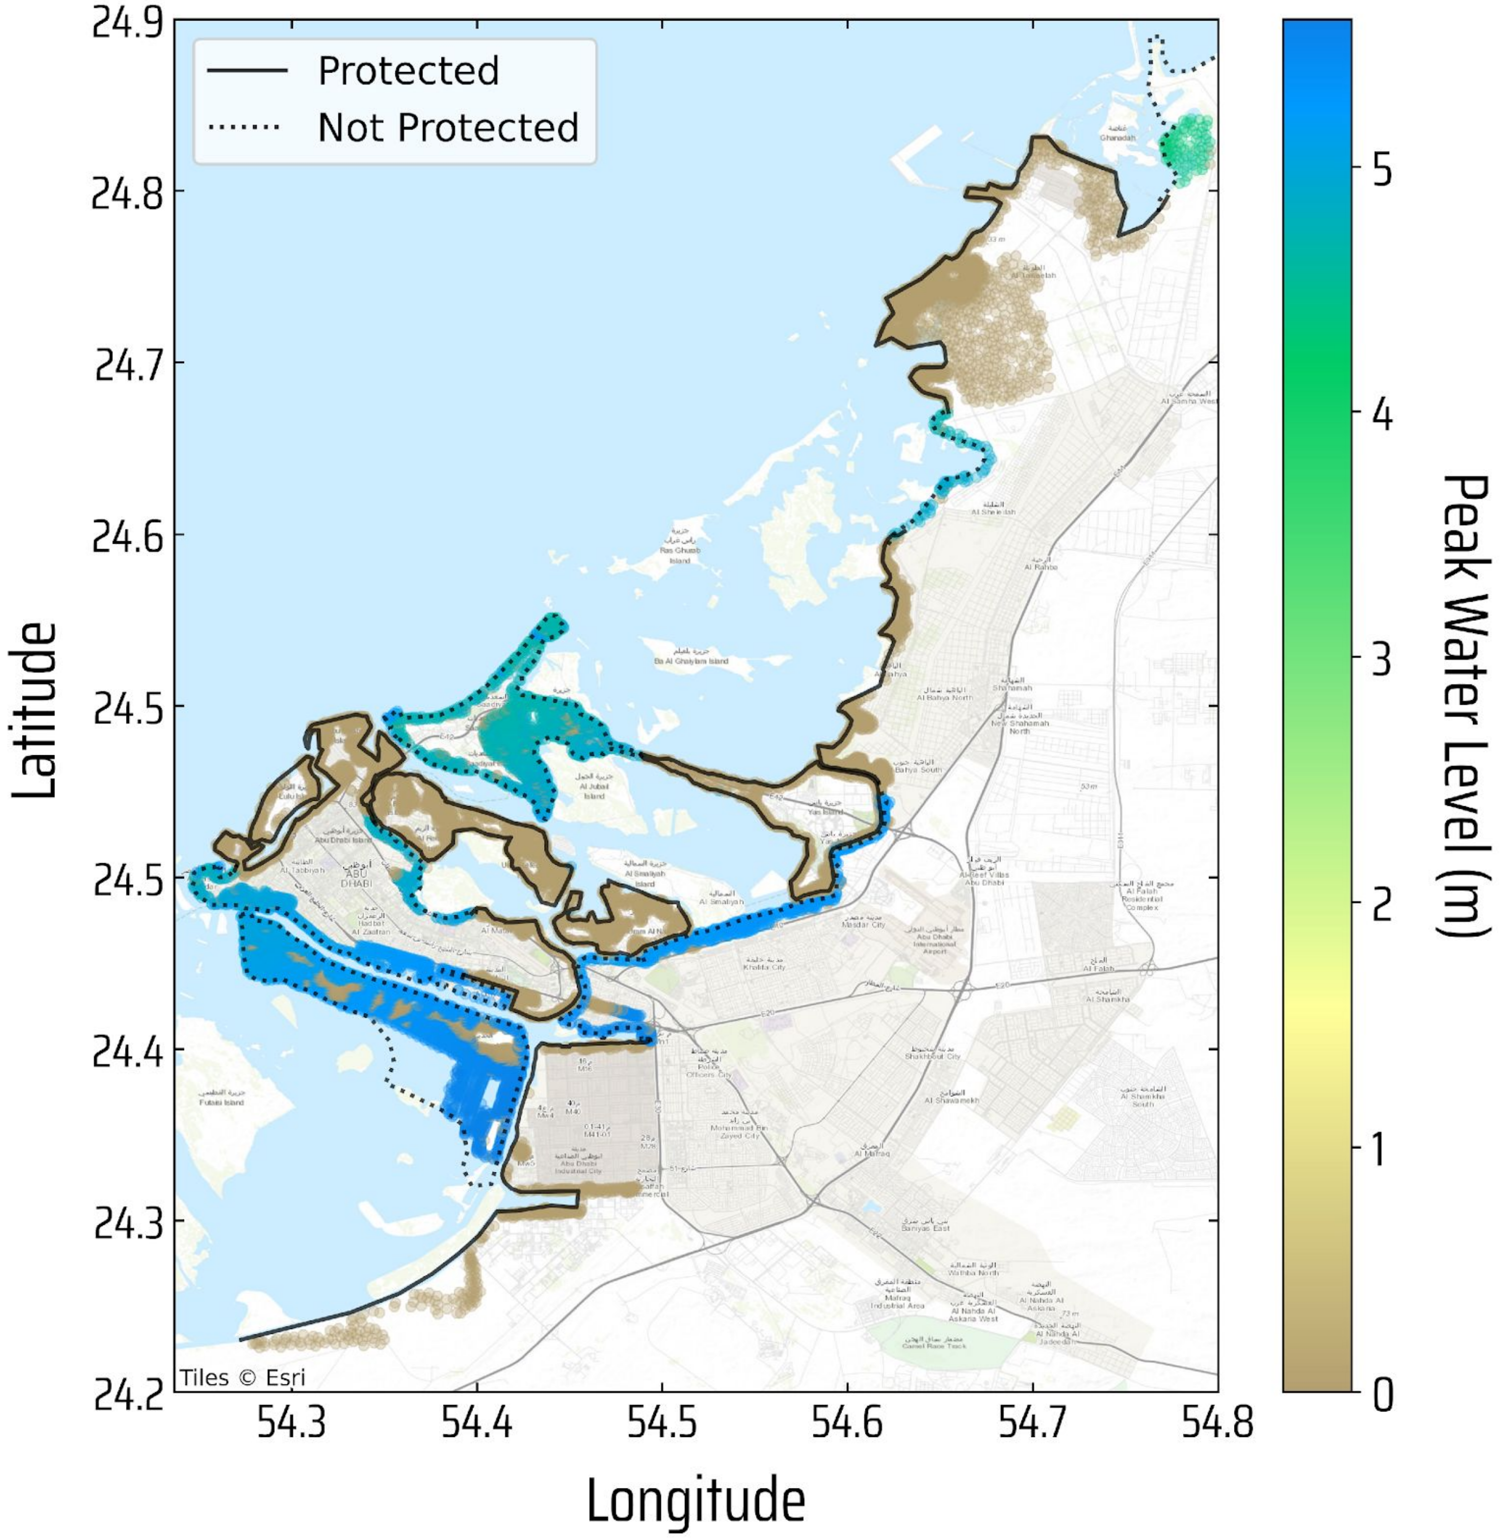

Challenges

Data-scarce Regime

- Few training samples (e.g., 100)
- Historical flood data is not applicable

Dense Prediction

- High-resolution output for detailed inundation mapping
- Tens of thousands output variables

CONTRIBUTIONS

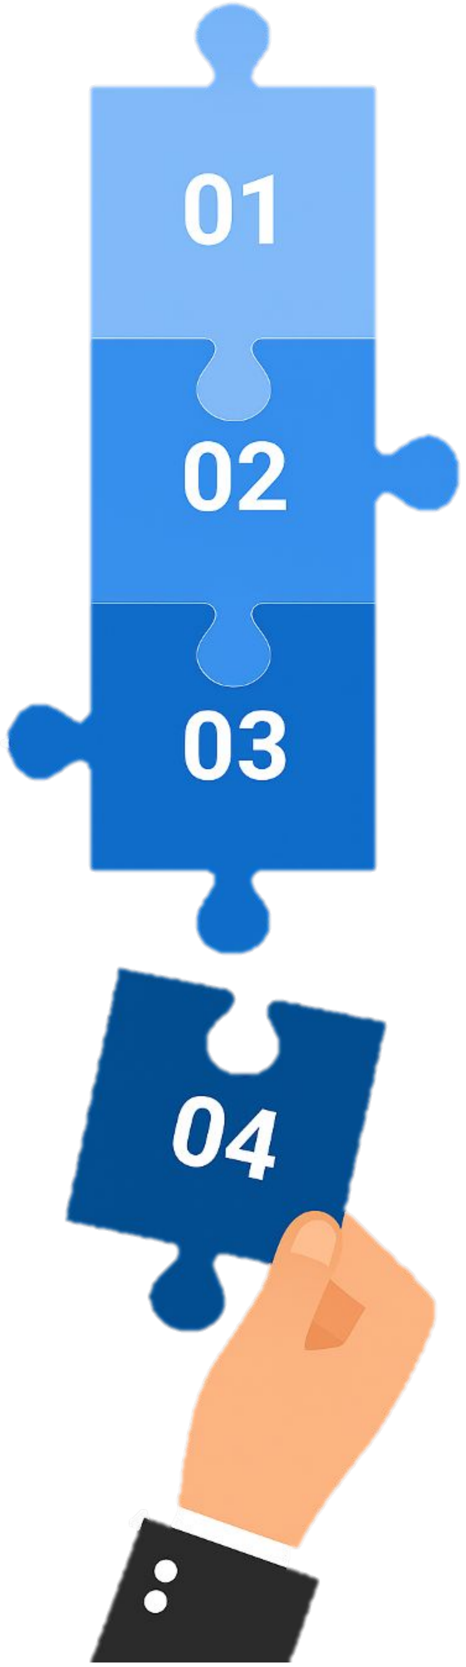

- ❖ Propose CASPIAN-v2, a lightweight CNN for accurate flood prediction under diverse SLR & shoreline protection scenarios.
- ❖ Provide comprehensive datasets for Abu Dhabi and San Francisco covering multiple adaptation strategies.
- ❖ Provide interpretability through explainable AI techniques to support decision-making.

STUDY AREAS

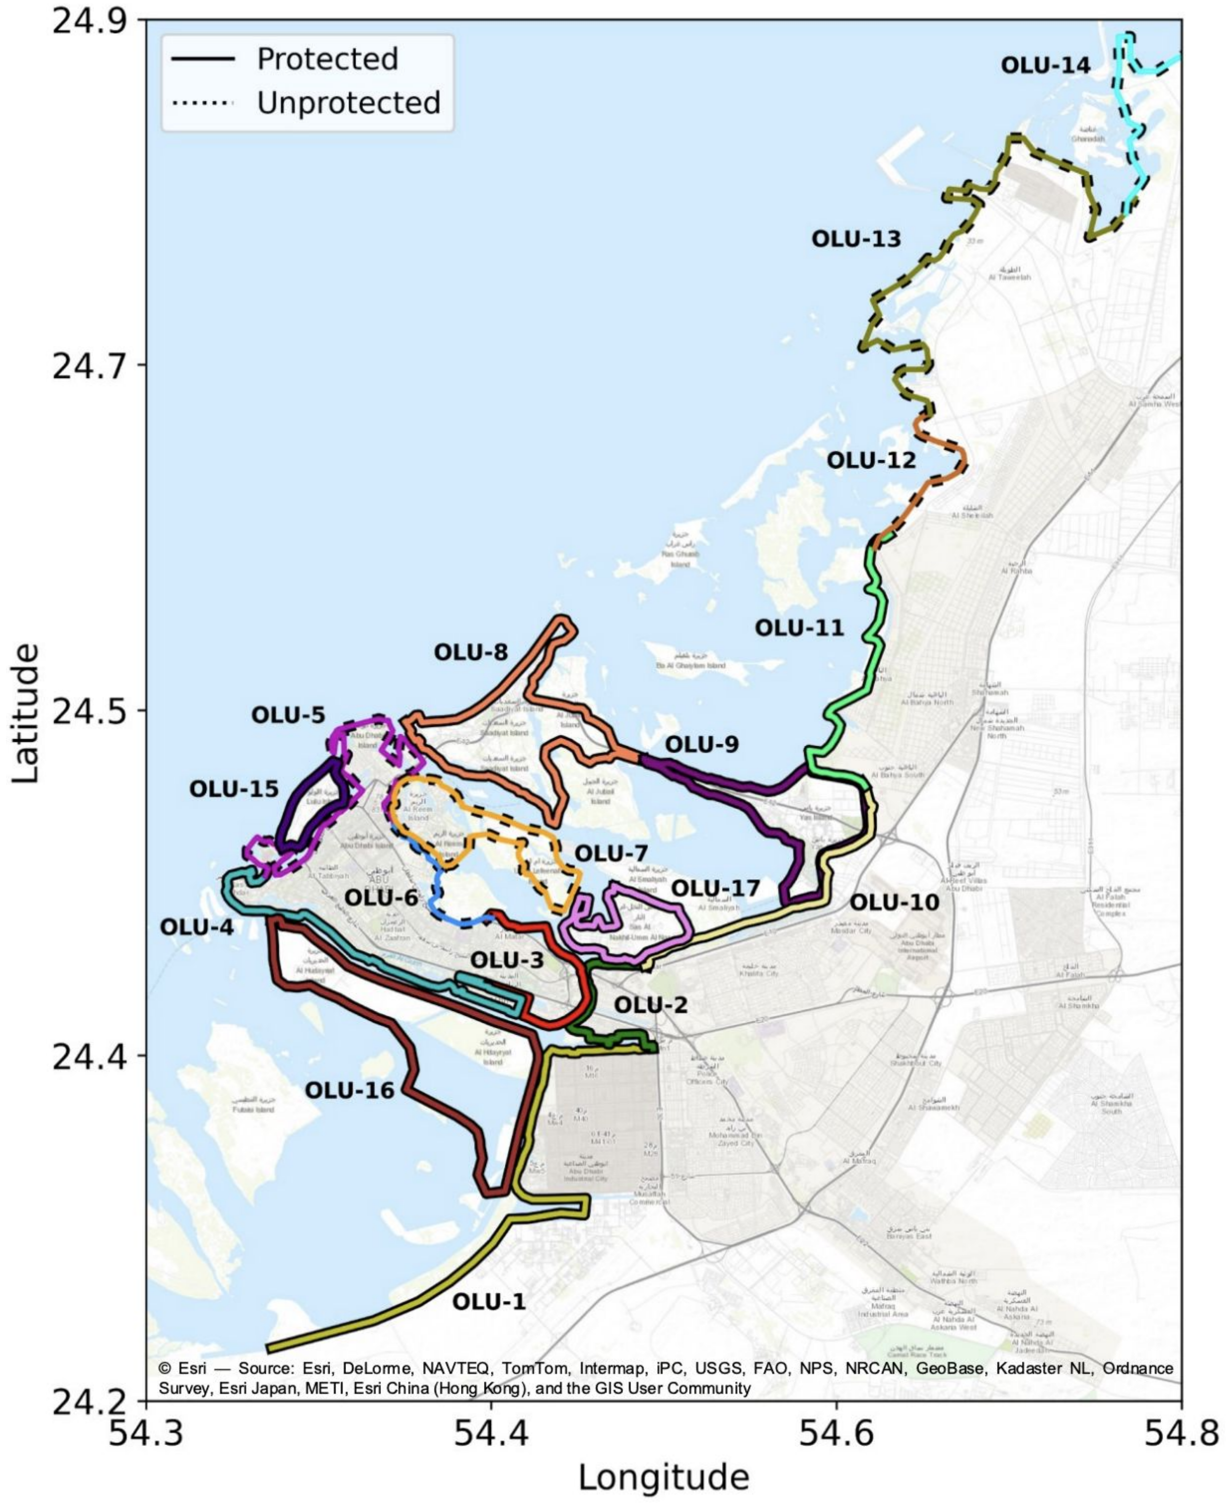

Abu Dhabi, UAE

- Shallow bathymetry amplifies risk
- 0.5 m SLR by mid-century could double flood zones
- 17 OLUs as defined by Abu Dhabi Urban Planning Council

San Francisco Bay, US

- Sheltered inland bay yet subject to tidal flooding
- Sea walls may raise bay water levels by 1 m
- 30 OLUs discretized by morphology & infrastructure

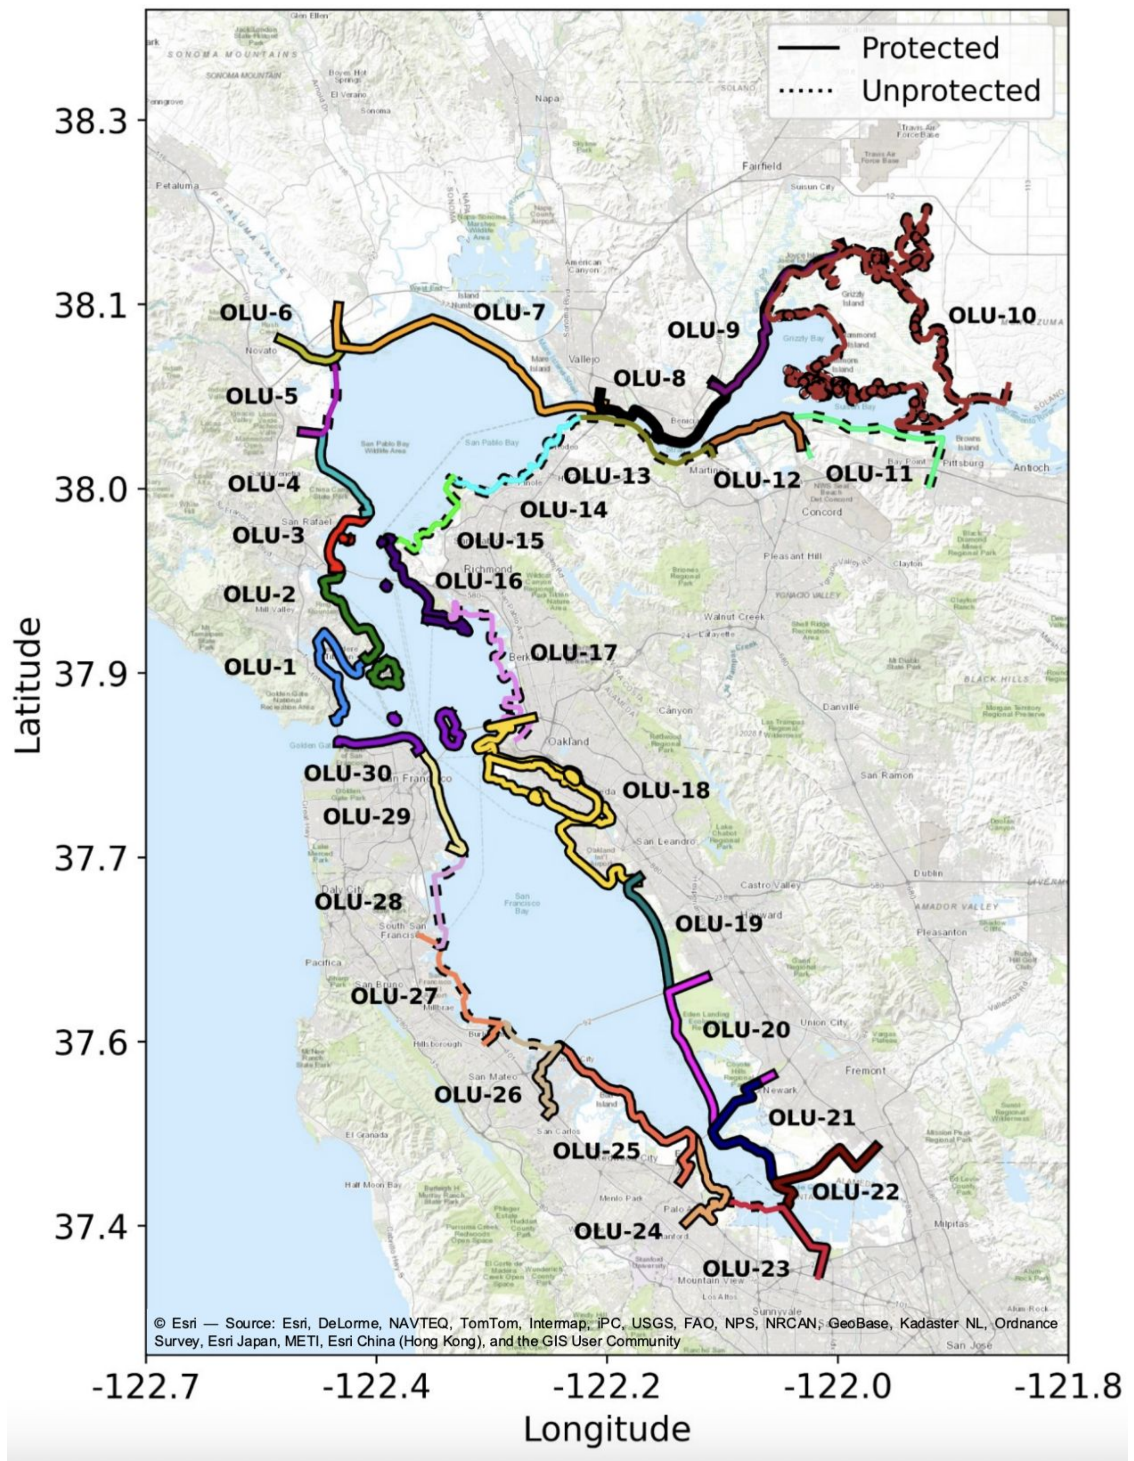

DATASET DETAILS

Region	Set	SLR	Total Scenarios	Training	Validation	Test
AD	Main	0.5m	142	96	10	36
	Holdout	0.5m	32	-	-	32
SF	Main	1.0m	285	225	24	36
	Holdout	1.0m	46	-	-	46
	Generalizability	0.5m	30	20	4	6
		1.5m	30	20	4	6

PROPOSED FRAMEWORK

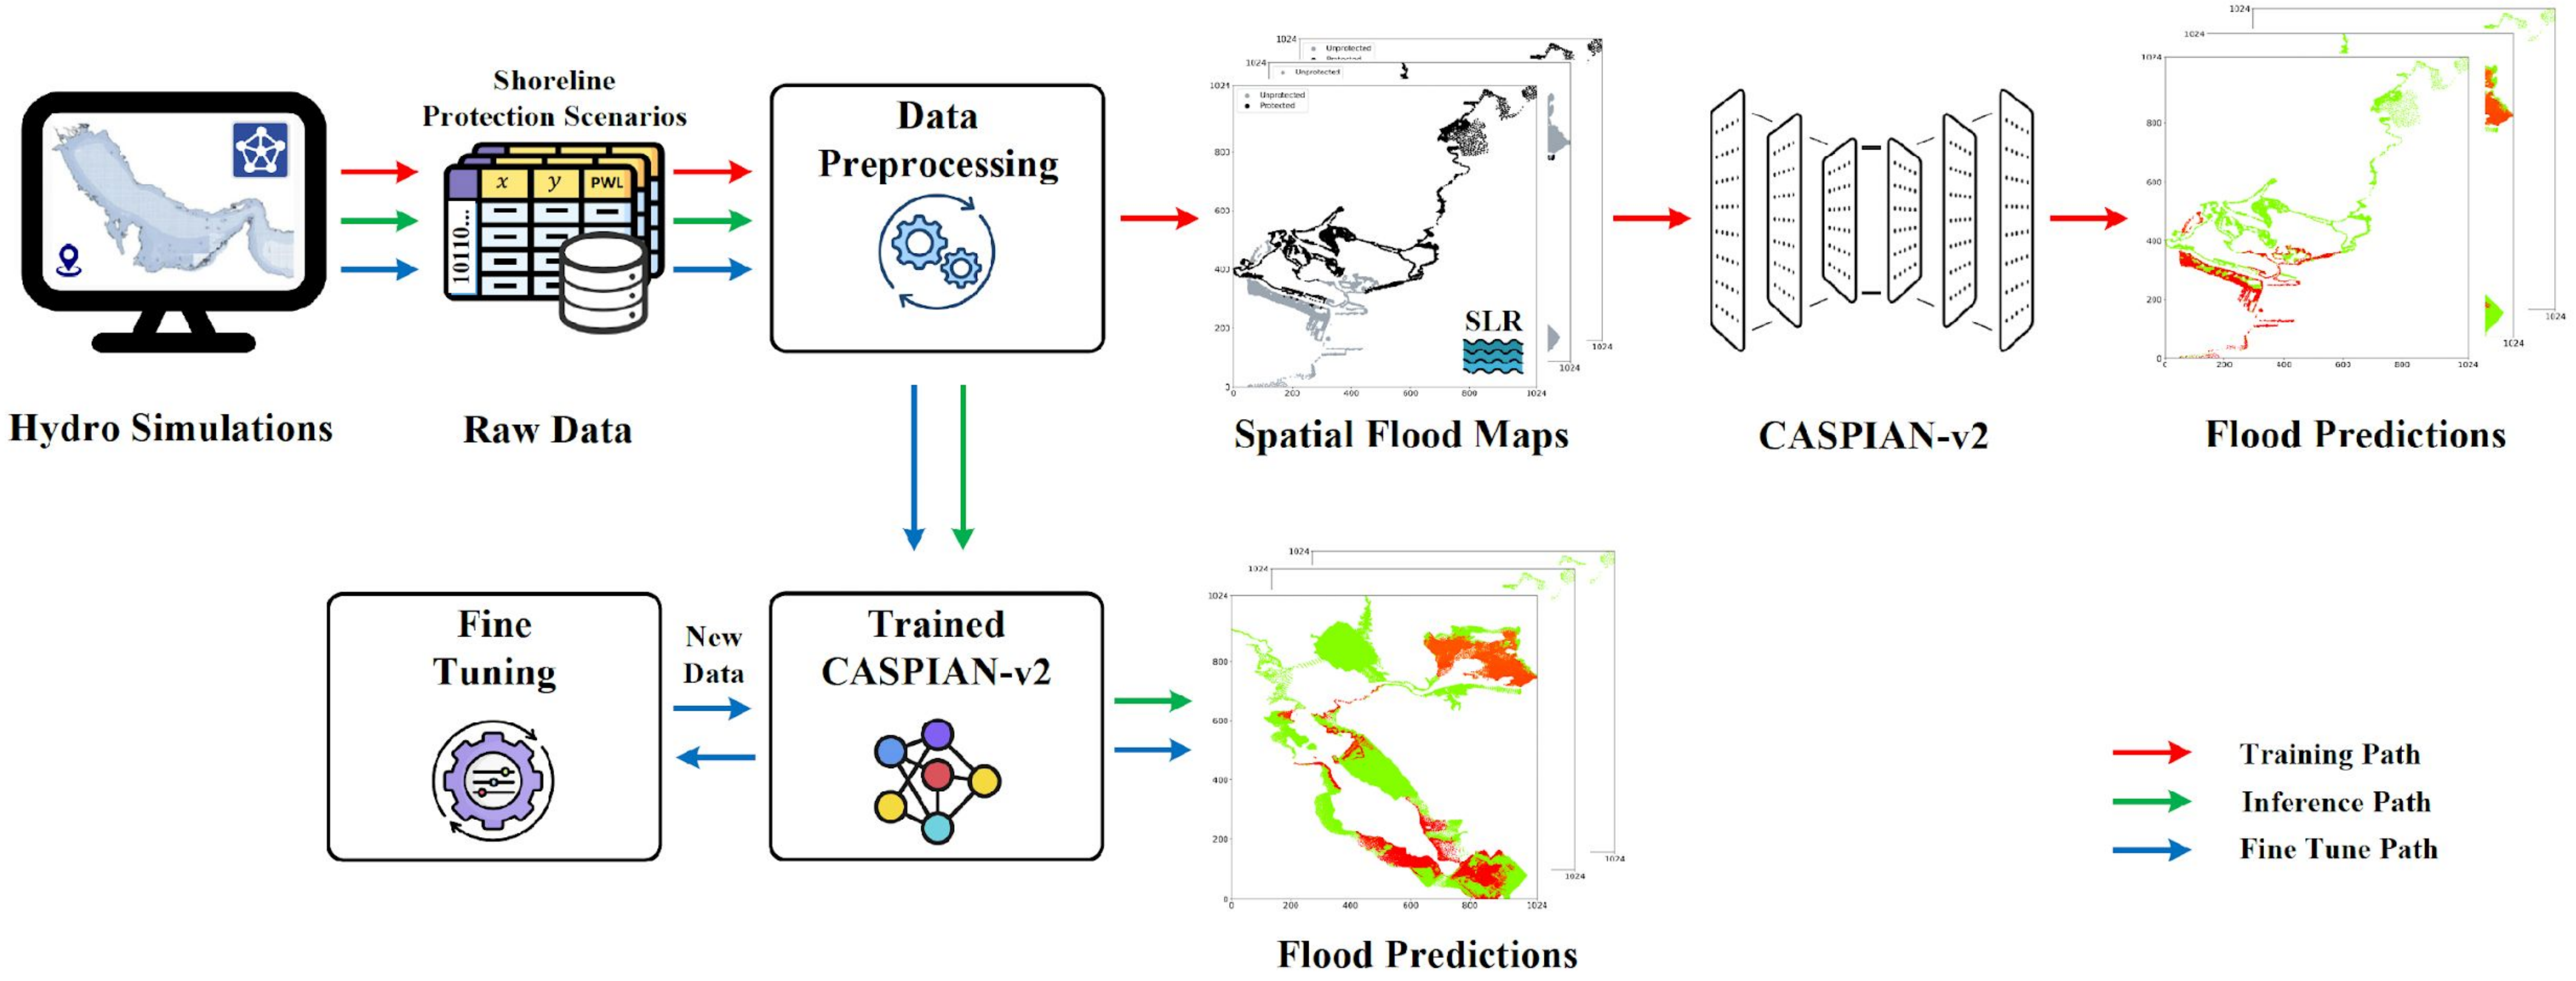

- Hydrodynamic simulations generate raw flood data under different SLR depths and protection scenarios.
- This data is transformed into 2D spatial maps and used to train the CASPIAN-v2 model.
- Once trained, CASPIAN-v2 rapidly predicts flood maps for new scenarios.
- With limited new data, the model can be fine-tuned to adapt to new regions or SLR levels.

CASPIAN-v2 MODEL

Encoder
- Feature extraction (FE) blocks with depthwise separable convs
- Downsample & preserve context via skip connections

Bottleneck
- ResNeXt + CBAM attention
- Emphasize critical spatial & channel features

Decoder
- Feature reduction (FR) blocks to upsample & fuse encoder features
- Integrate SLR to produce a high-resolution flood map

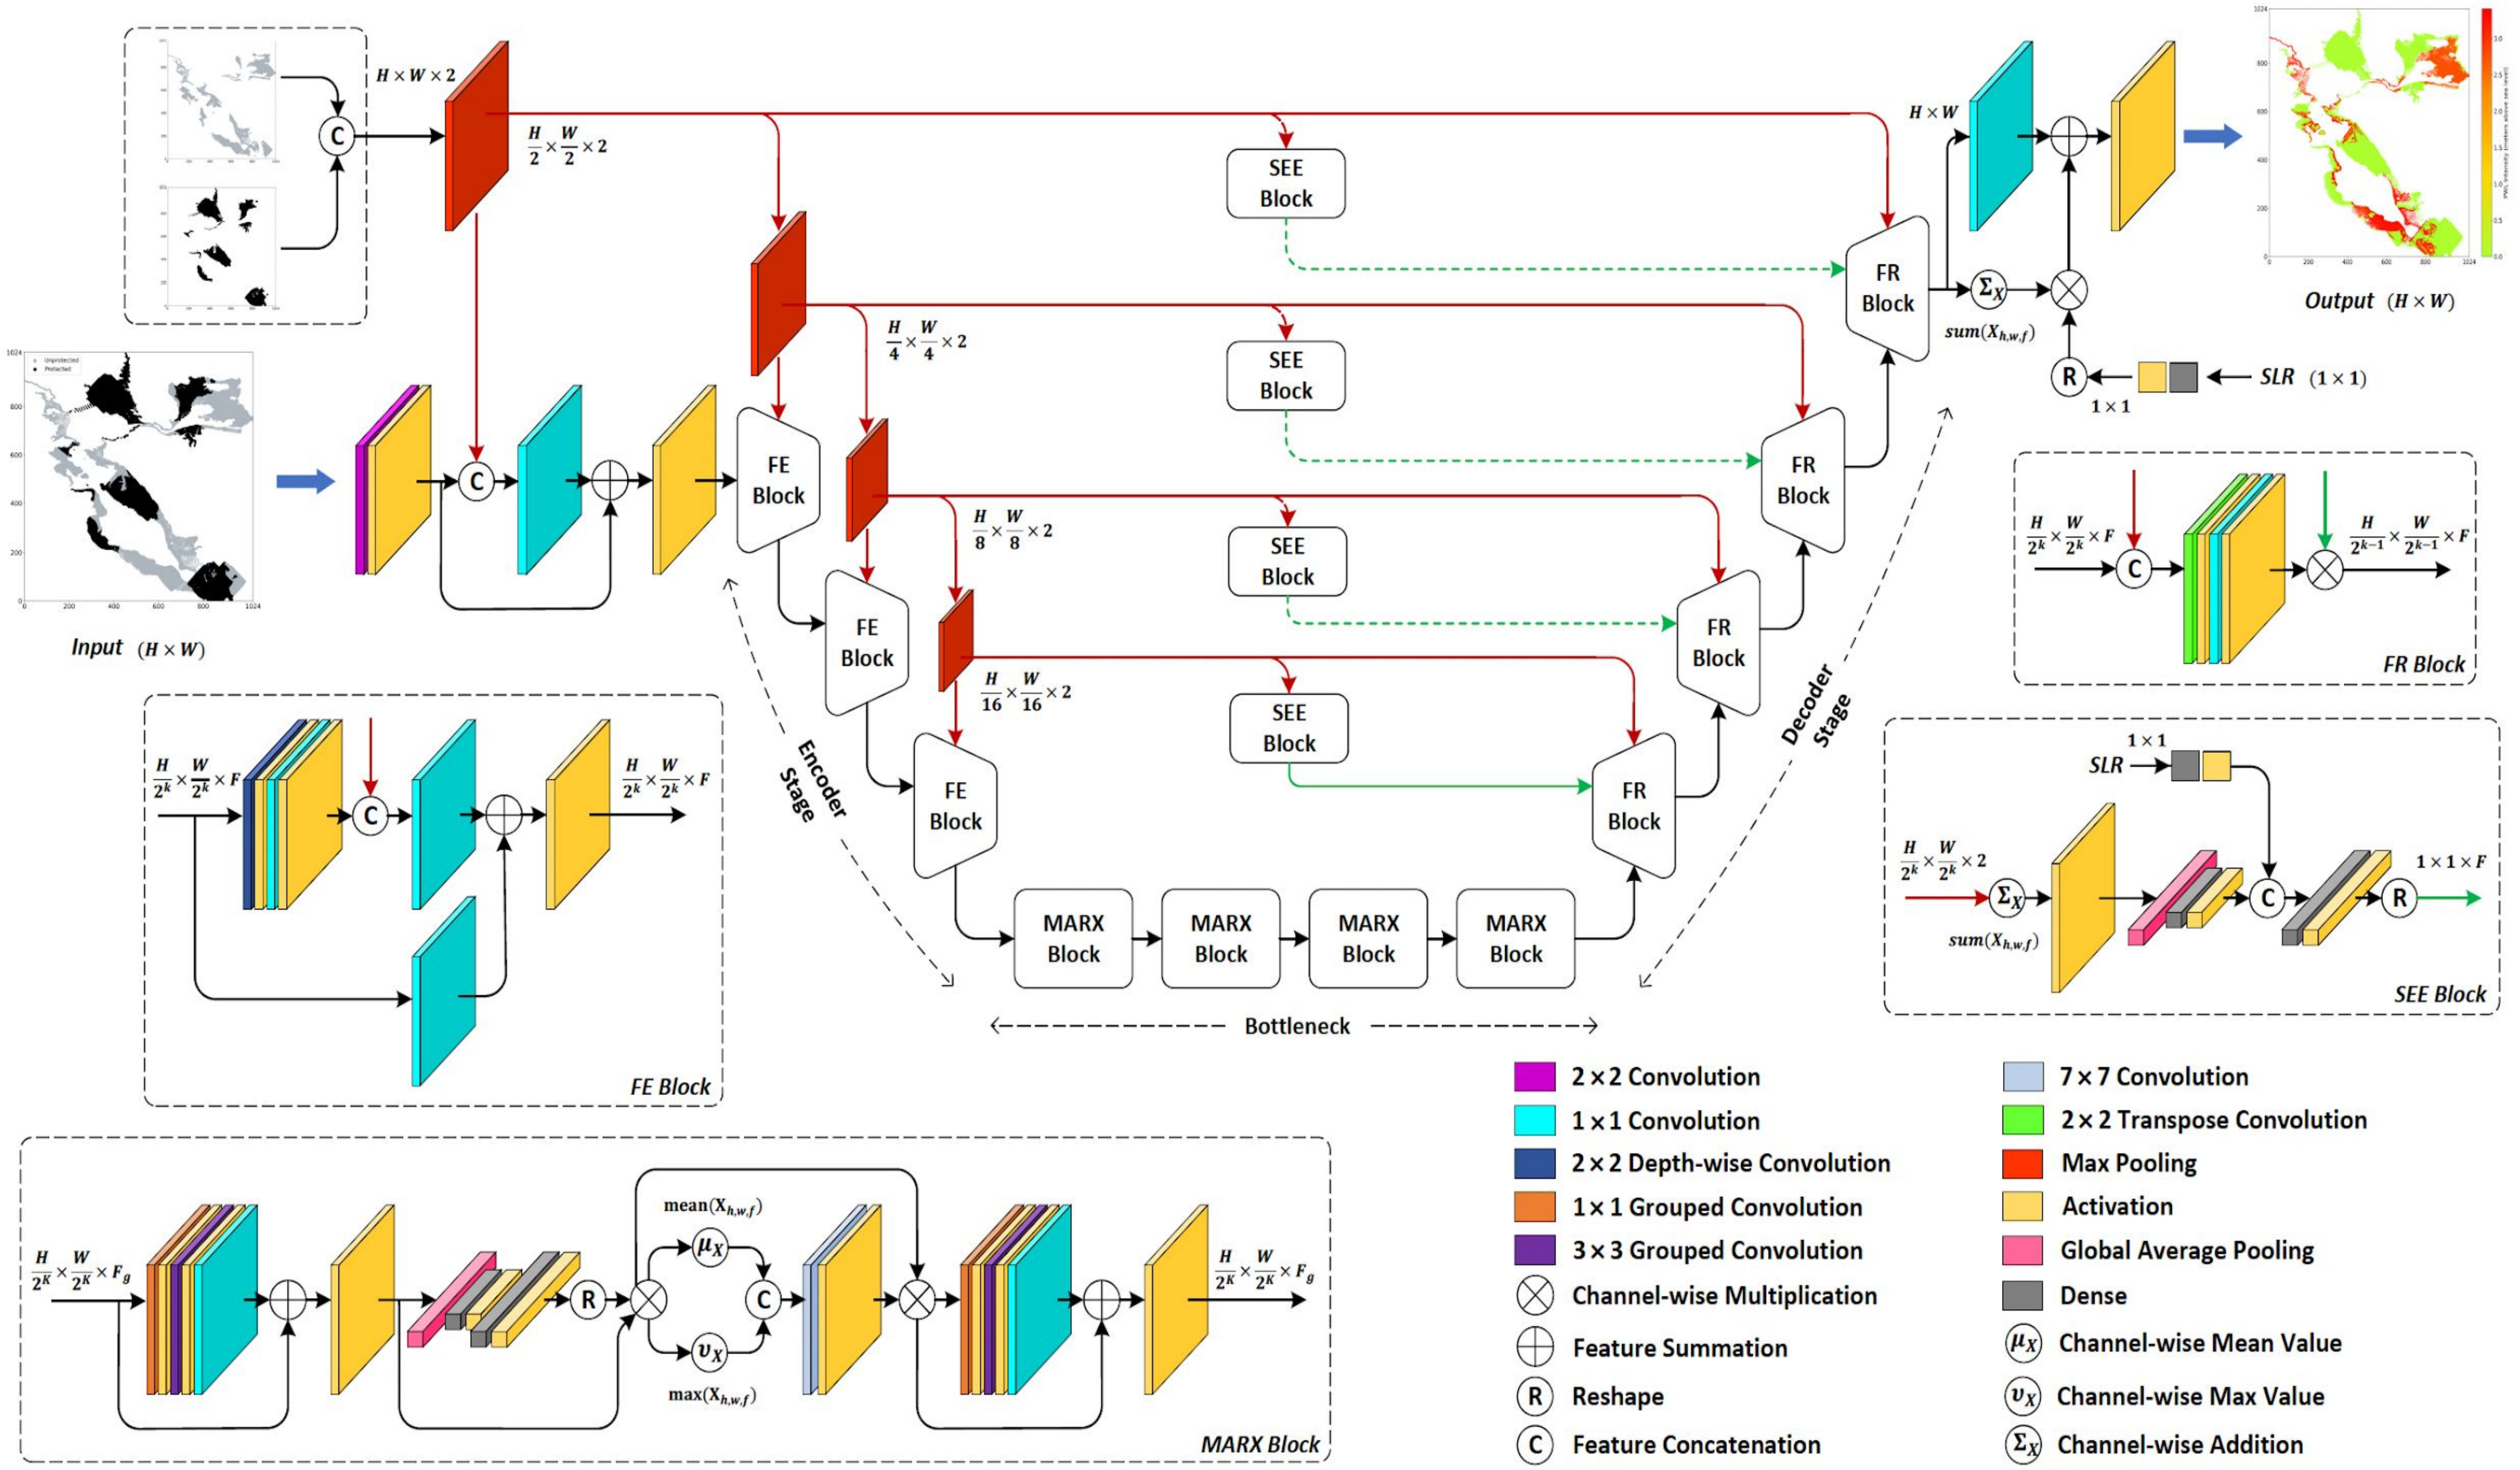

Evaluation Metrics

1. MAE
2. RMSE
3. RTAE
4. $\mathbf{R}^2$
5. Acc[0]
6. DSC
7. % of Small Errors (> 0.1m)
8. % of Big Errors (> 0.5 m)

QUANTITATIVE RESULTS

The top-performing result for each metric is highlighted in red, and the second-best is highlighted in blue

Type	Model	Prediction Accuracy								Computational Efficiency		
		MAE ↓	RMSE ↓	RTAE ↓	$\delta > 0.5$ ↓	$\delta > 0.1$ ↓	R^2 ↑	Acc[0] ↑	DSC ↑	Param↓	TT ↓	IT ↓
Simulator	AD Pipeline†	*Served as the ground truth*								-	-	71–73h
	SF Pipeline^o									-	-	3.5–6.0h
ML (1-D)	Naïve	1.53	3.54	1746.06	74.92%	80.11%	0.54	31.01%	0.38	-	62s	0.15s
	RF	0.54	0.73	264.95	36.77%	72.20%	0.79	34.19%	0.41	-	75s	0.18s
	Linear	0.12	0.19	64.98	7.87%	14.03%	0.94	59.28%	0.62	-	65s	0.16s
	XGBoost	0.25	0.24	164.16	16.27%	49.88%	0.93	44.10%	0.47	-	198s	0.21s
	SVR	0.20	0.24	72.31	9.24%	41.17%	0.92	45.46	0.48	-	79s	0.19s
	Lasso Poly	0.09	0.12	28.15	4.47%	15.04%	0.96	55.78%	0.64	-	72s	0.17s
	Kriging	0.10	0.24	39.90	5.22%	11.59%	0.94	62.88%	0.63	-	76s	0.18s
DL (1-D)	MLP	0.64	2.72	524.17	32.82%	41.94%	0.65	36.91%	0.43	0.01M	14h	5.03s
	CCT	0.90	2.32	843.54	48.08%	64.63%	0.66	34.01%	0.42	11.05M	18h	0.26s
DL (2-D)	Atten-Unet	0.10	0.37	11.82	3.14%	16.70%	0.91	95.26%	0.73	12.07M	46h	0.24s
	Atten-Unet*	0.10	0.36	11.65	3.31%	15.62%	0.92	94.99%	0.74	12.07M	47h	0.27s
	Swin-Unet	0.06	0.27	6.72	1.47%	12.94%	0.95	98.10%	0.80	8.29M	26h	0.24s
	CASPIAN	0.05	0.36	5.85	1.01%	4.79%	0.92	98.84%	0.82	0.36M	22h	0.22s
	Ours	0.04	0.30	6.54	0.89%	3.55%	0.93	99.39%	0.84	0.38M	22h	0.22s

* with pre-trained encoder on ImageNet [36].

† AD pipeline includes Delft3D + SWAN + post processing.

^{o} AD pipeline includes Delft3D + post processing.

Summary

- 51.65% AMAE reduction vs best ML baseline
- 19.96% AMAE improvement vs 2nd best DL model
- 0.38M parameters & 22 h training time
- 0.22 s inference per scenario vs up to 71–73 h for simulators

QUANTITATIVE RESULTS

Holdout Dataset

Region	AMAE ↓	ARMSE ↓	R^2 ↑	Acc[0] ↑
AD	0.0792	0.4871	0.9525	99.07%
SF	0.0317	0.2259	0.9694	99.64%
Combined	0.0512	0.3331	0.9625	99.41%

On holdout sets, CASPIAN-v2 sustains accuracy, demonstrating robust generalization.

Generalization

Dataset (SLR)	AMAE ↓	ARMSE ↓	R^2 ↑	Acc[0] ↑
SF (0.5 m)	0.0626	0.2996	0.9336	97.99%
SF (1.5 m)	0.1005	0.4565	0.9196	98.23%
AD Holdout (0.5 m)	0.0567	0.2274	0.9901	99.18%
SF Holdout (1.0 m)	0.0433	0.2318	0.9685	99.34%
Aggregate	0.0652	0.3040	0.9520	98.69%

Fine-tuning on limited new data preserves performance: CASPIAN-v2 generalizes across unseen SLR levels and regions with minimal loss in accuracy

Spatial Accuracy

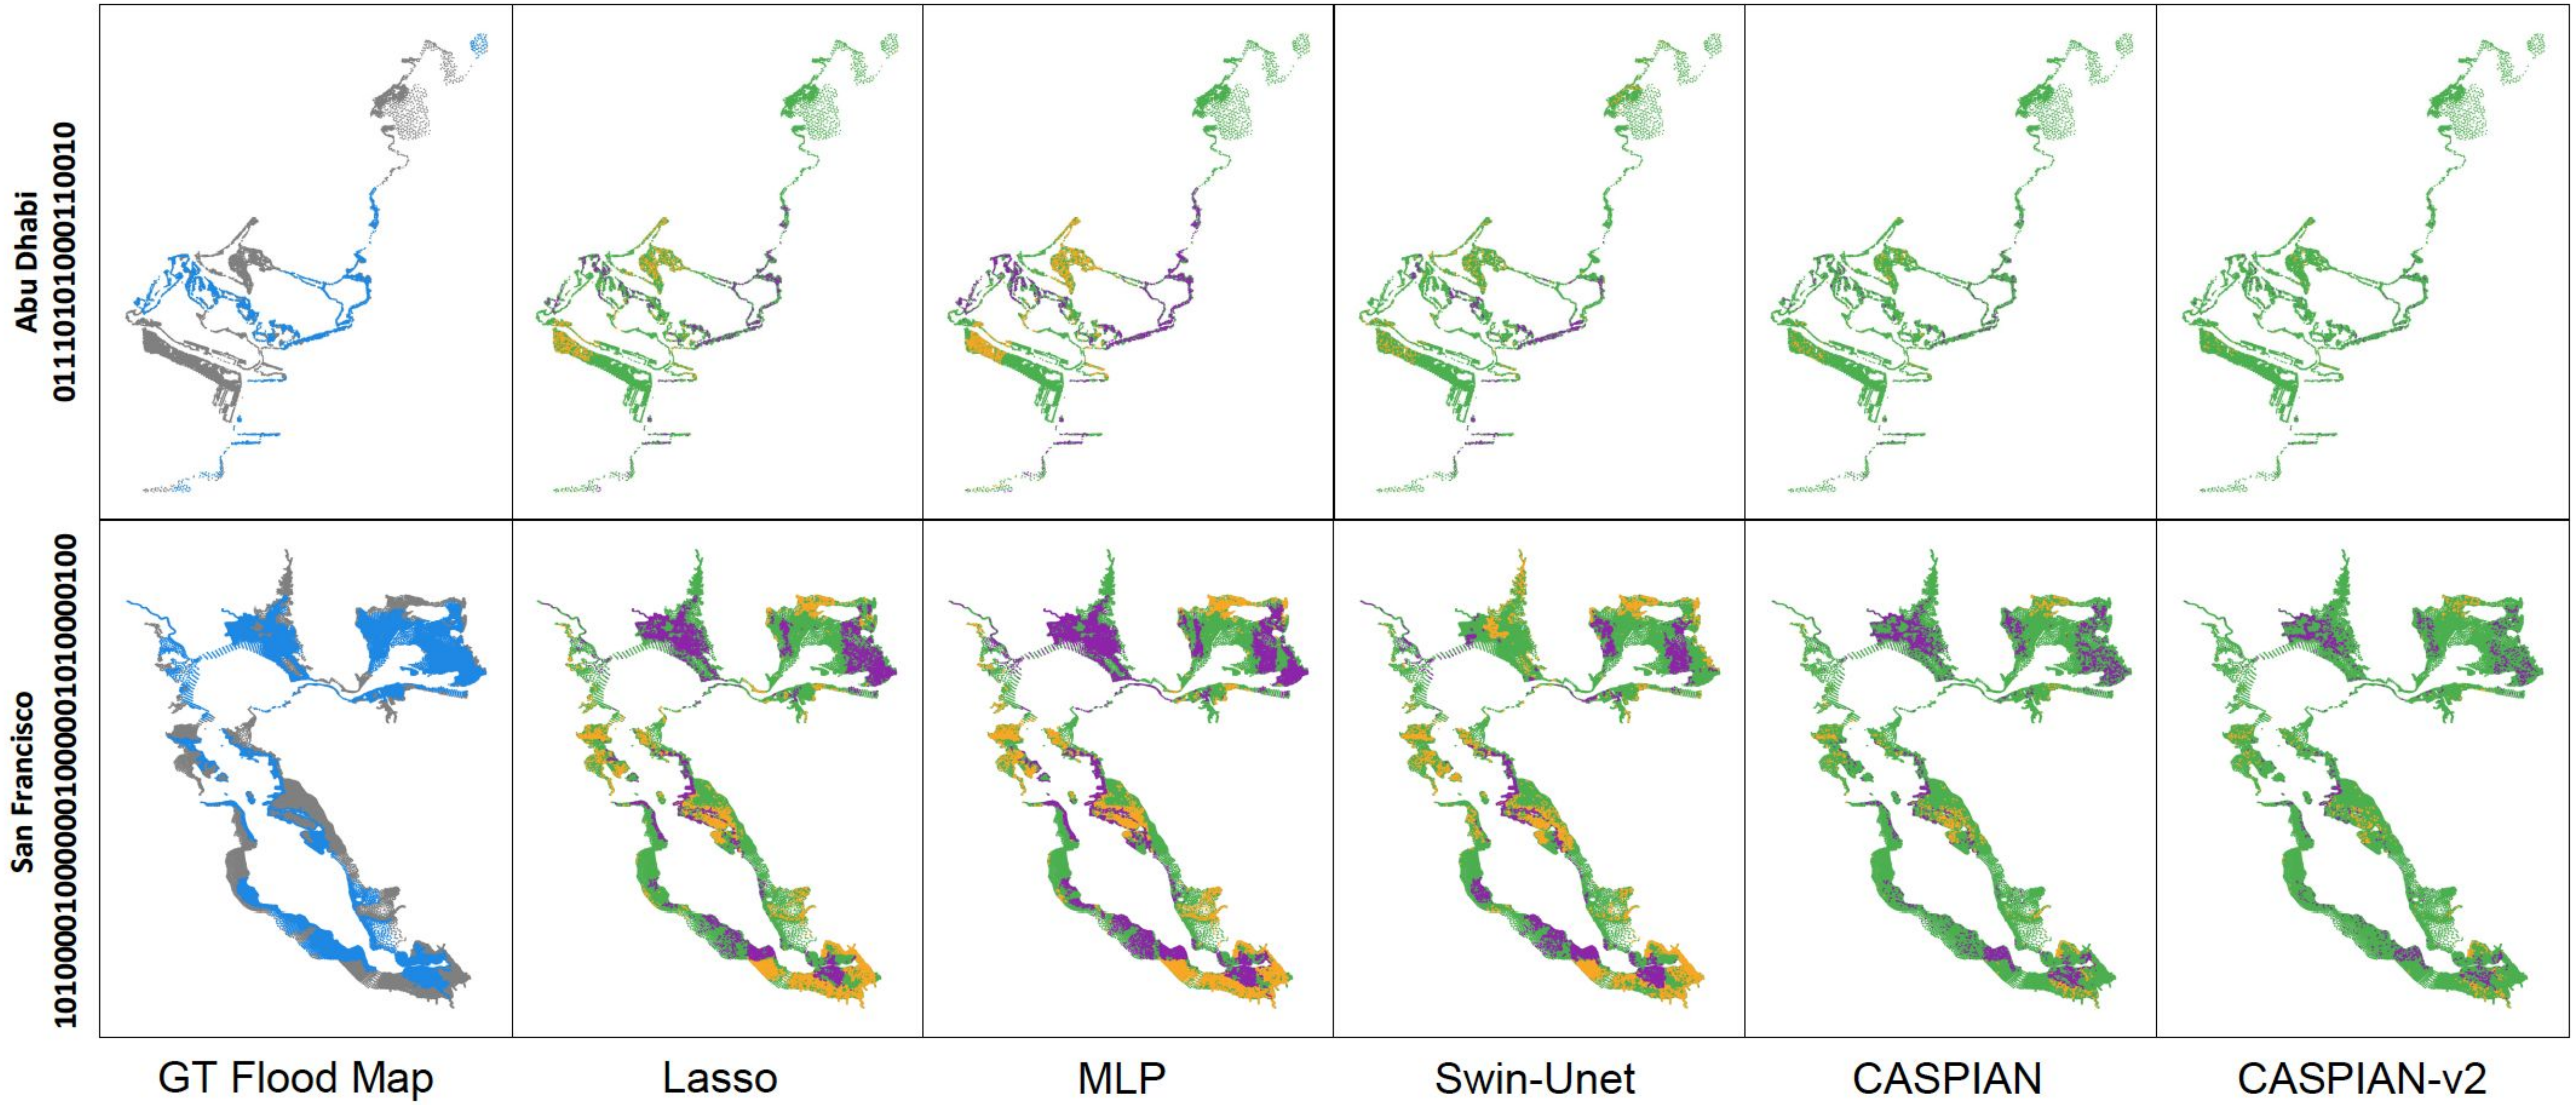

Visual comparison of spatial prediction accuracy for CASPIAN-v2 versus the top performing baseline model on a representative test case. Green indicates correctly predicted inundated areas (true positives), orange indicates over-prediction (false positives), and purple indicates underprediction (false negatives). CASPIAN-v2 demonstrates a larger matched area and more coherent flood boundaries.

MODEL INTERPRETABILITY

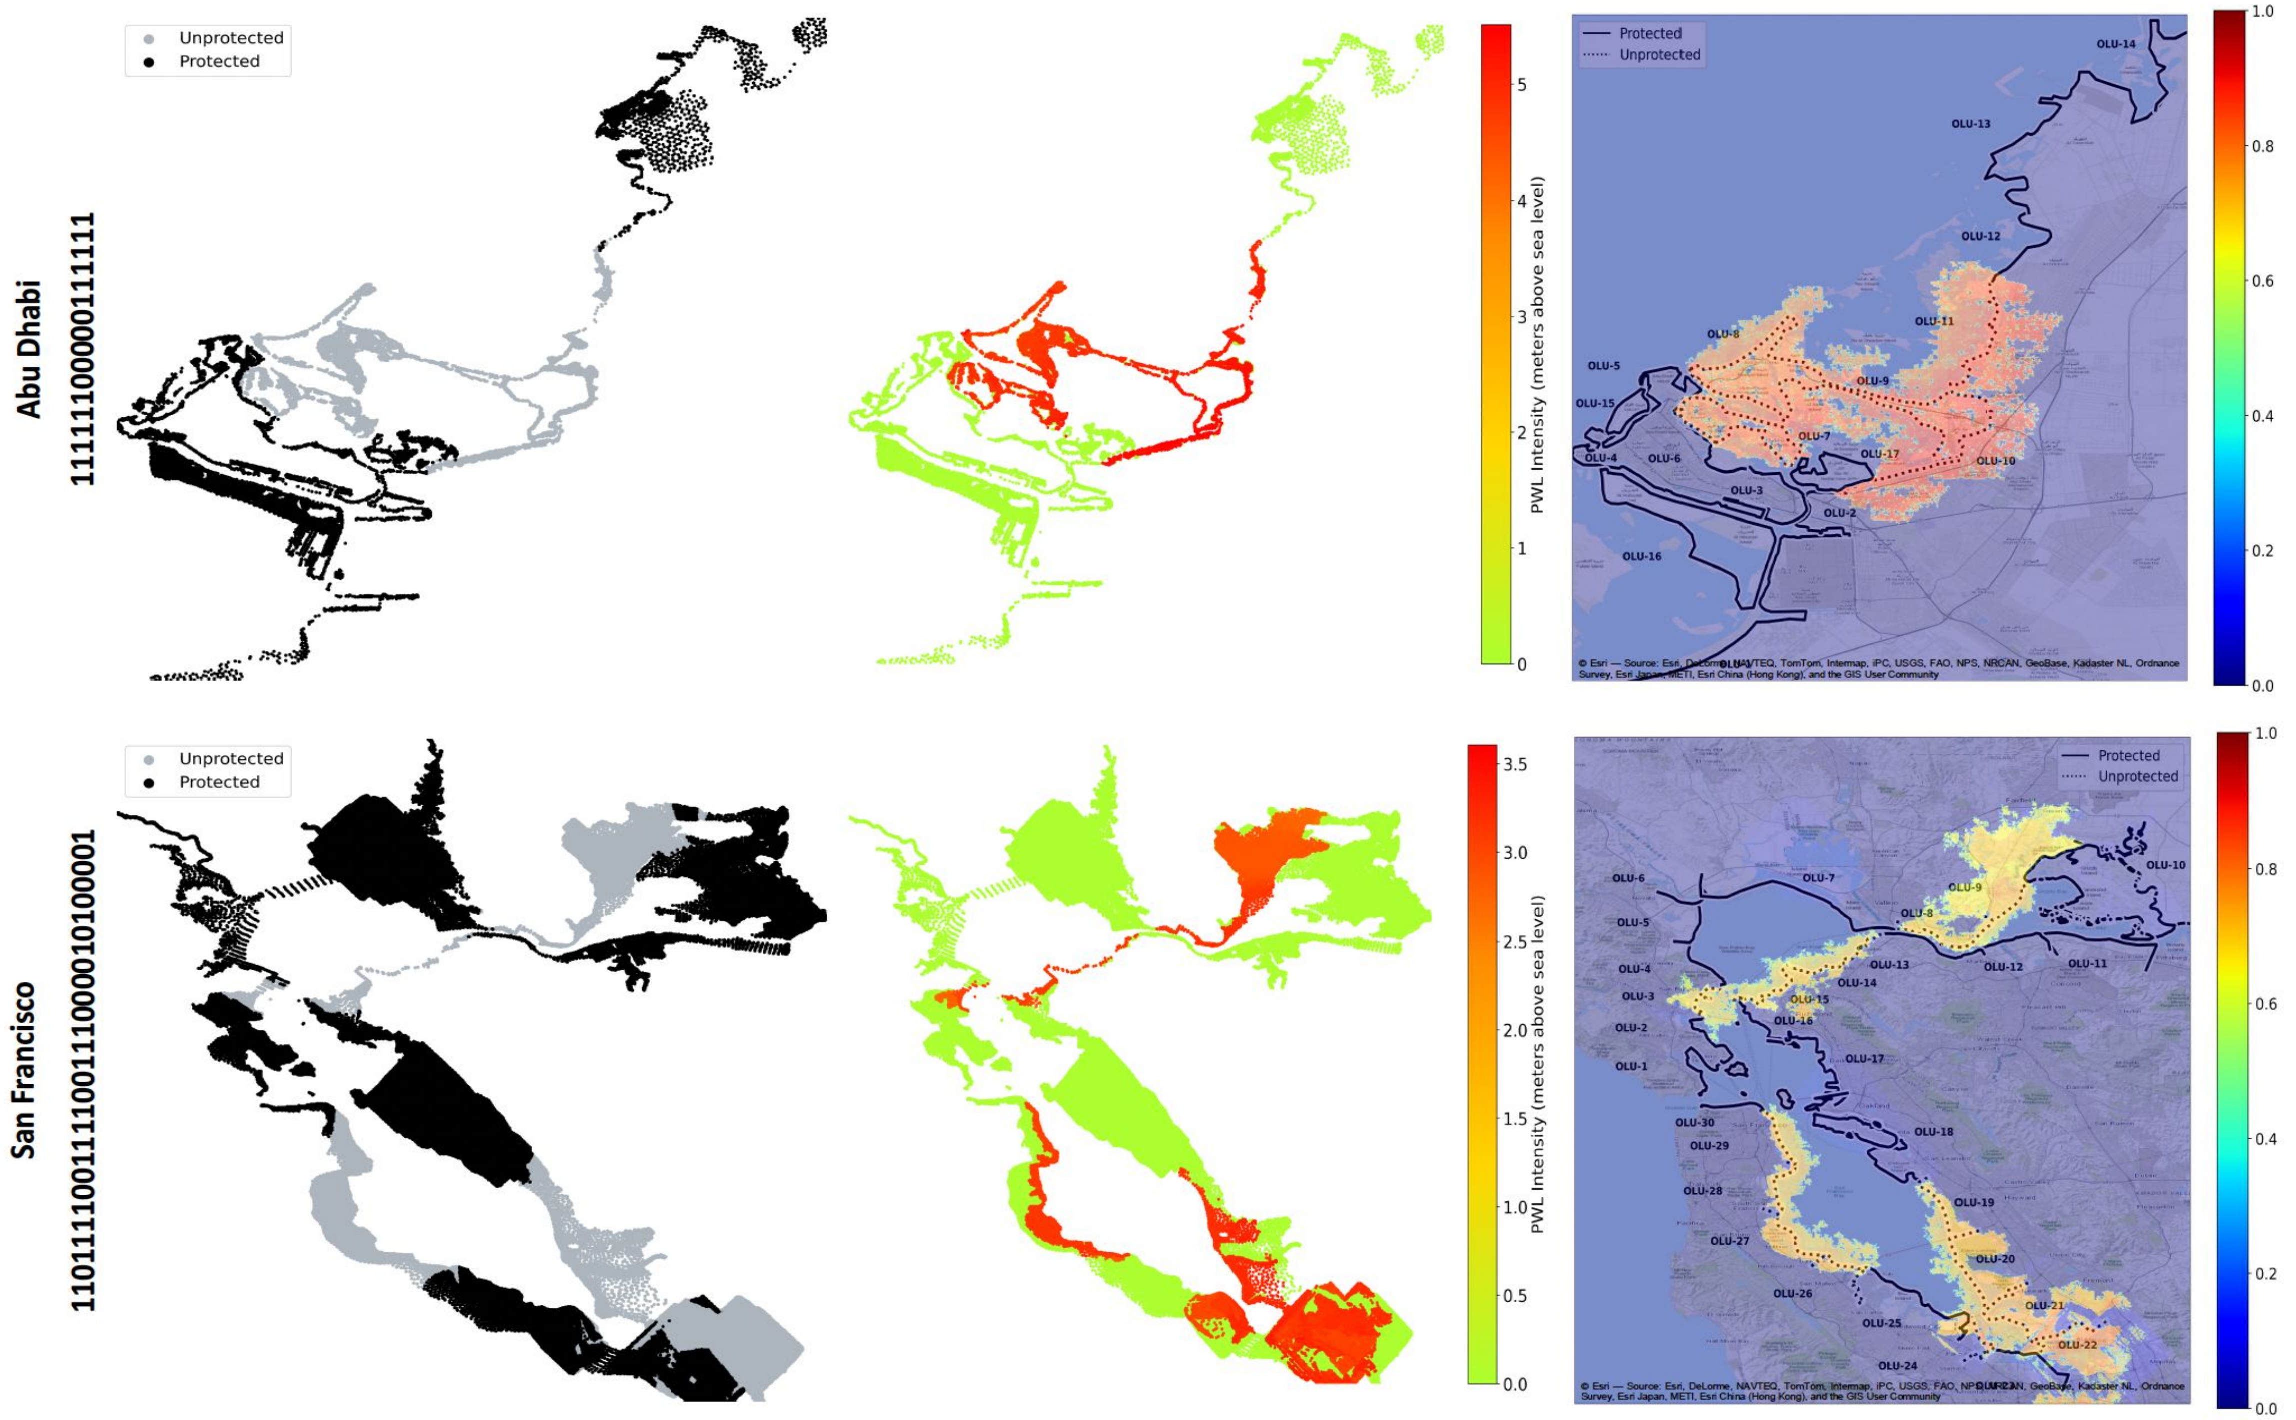

(a) (b) (c)

Visual comparison of CASPIAN-v2 inundation prediction and interpretability for AD (top) and SF (bottom): (a) Input maps, (b) Predicted inundation, (c) Grad-CAM visualizations highlighting model attention, which aligns with unprotected and vulnerable areas

UNCERTAINTY QUANTIFICATION

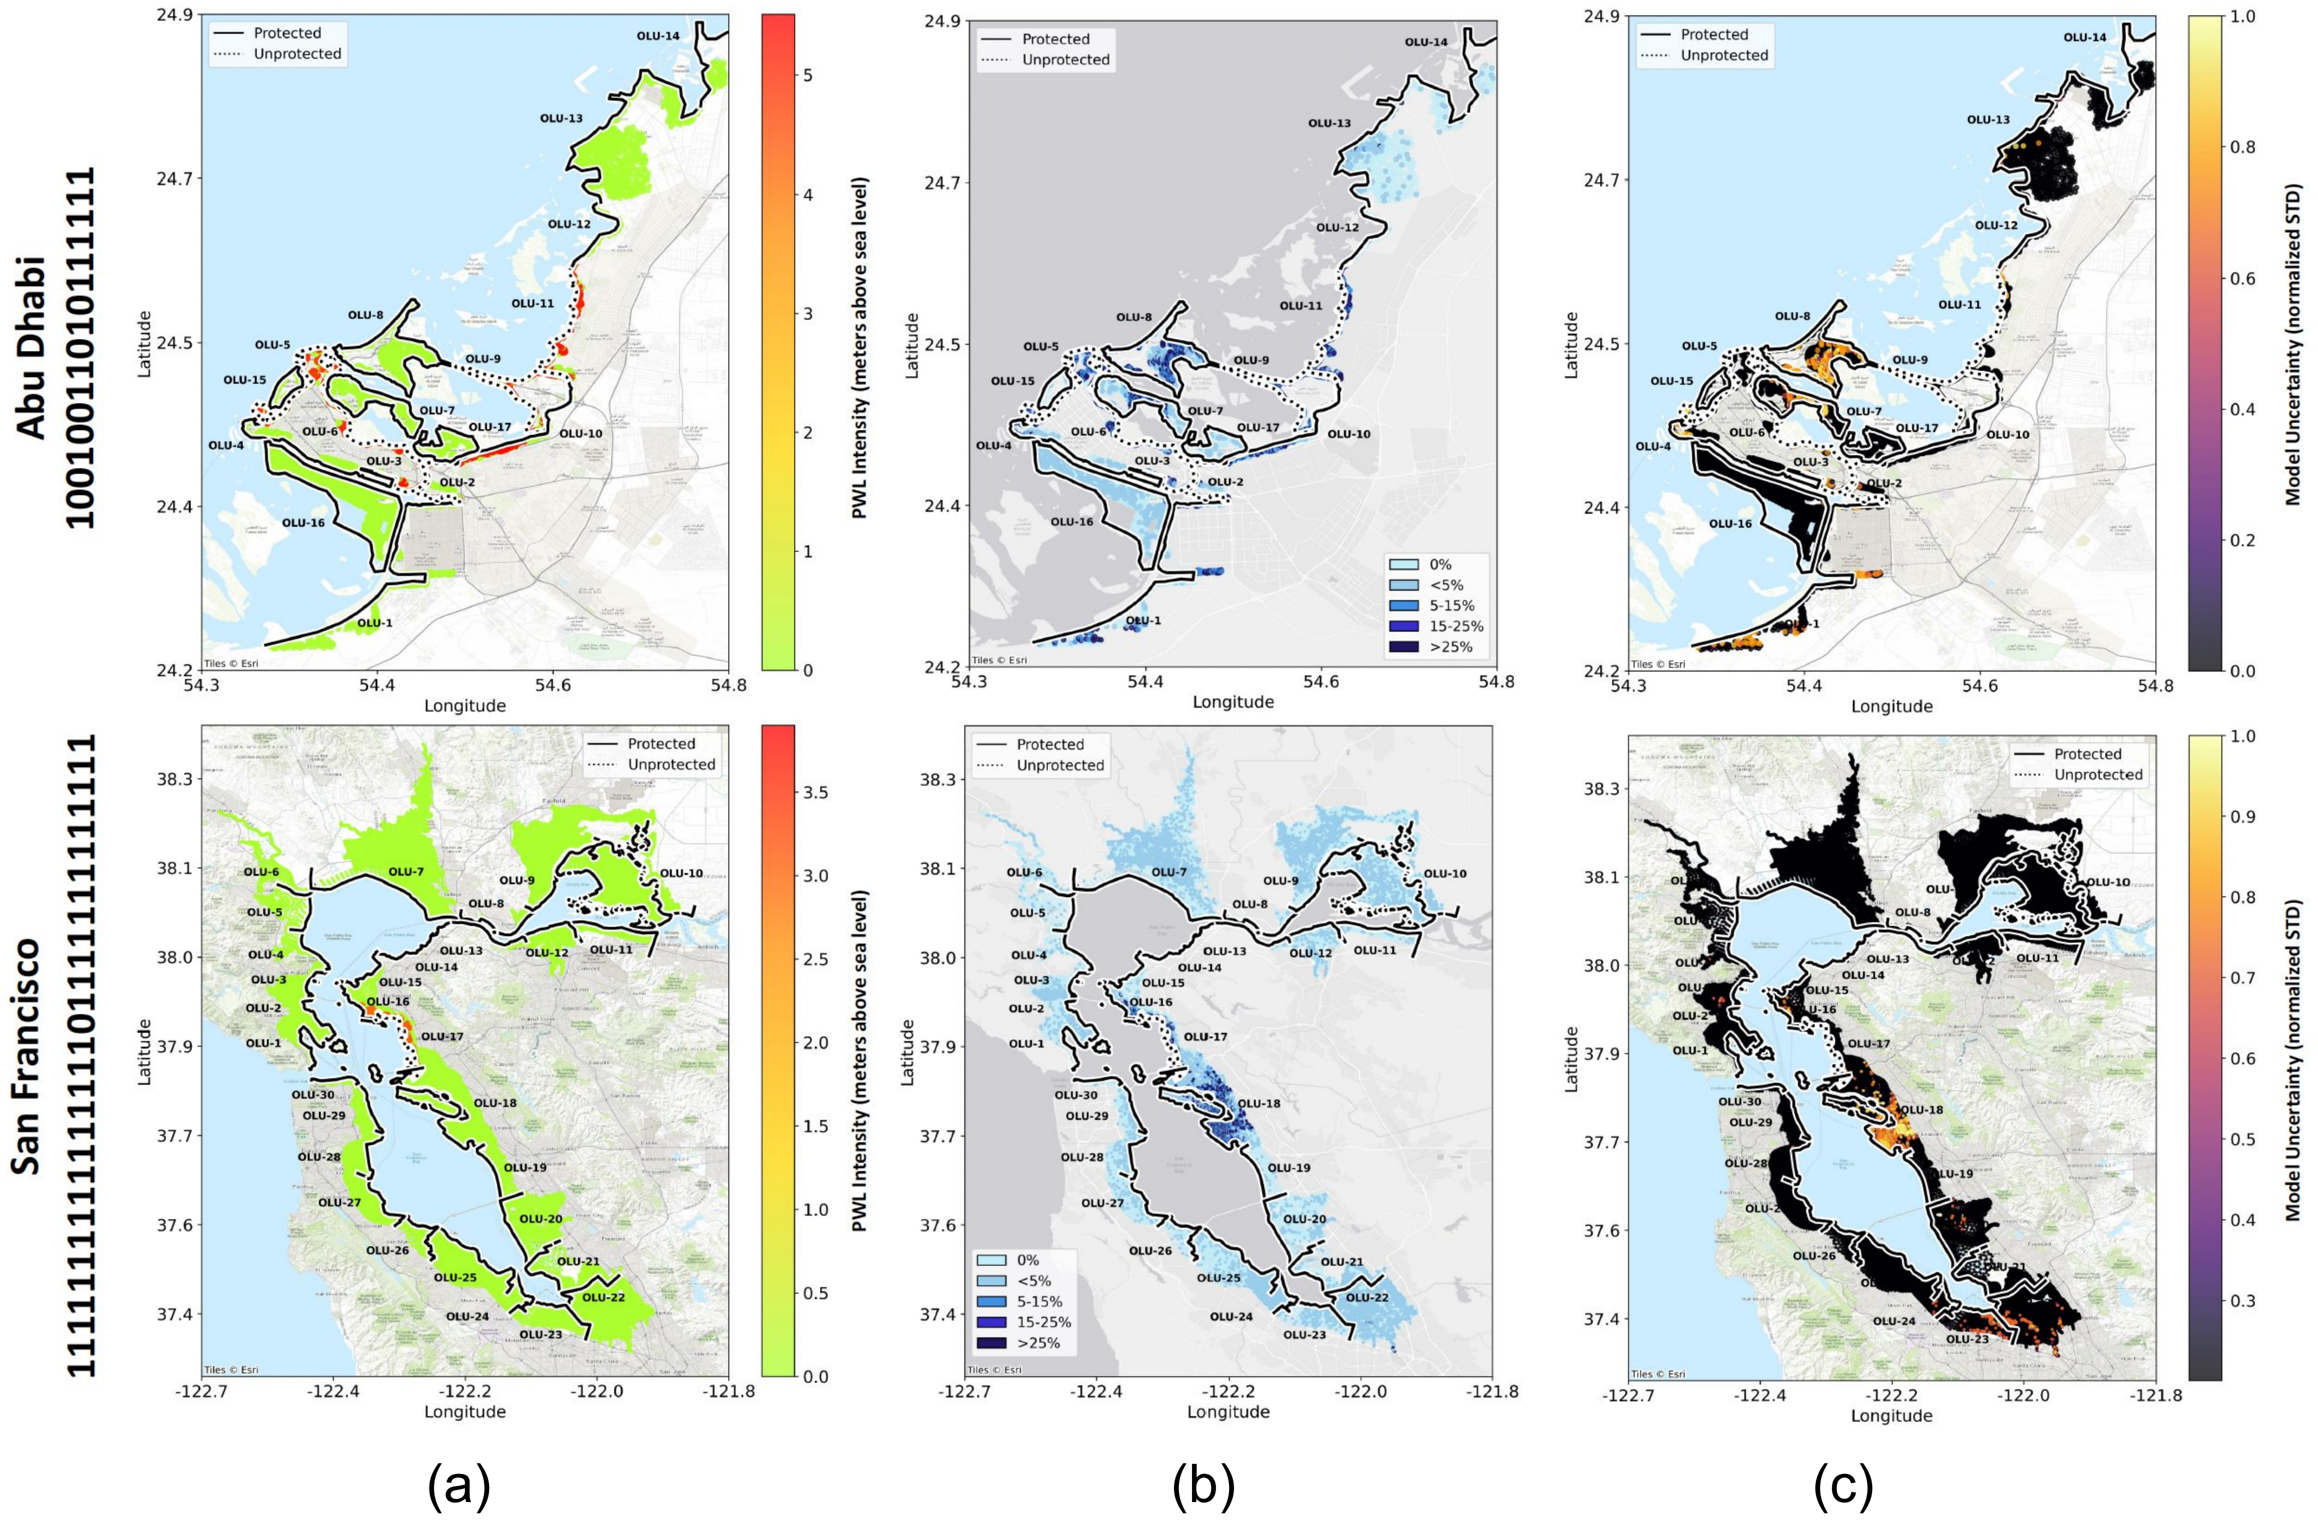

Predictive uncertainty maps for AD and SF scenarios. (a) Ground truth inundation. (b) Absolute error of the ensemble mean prediction. (c) Pixel-wise predictive uncertainty, where lighter colors indicate higher uncertainty and align with areas of higher error areas

Conclusions & Impact

Accuracy & Efficiency

- Achieves superior accuracy vs ML and DL baselines while remaining lightweight (0.38M parameters).
- Predictions for 72 scenarios in under 16s vs 115 days required for hydro simulations.

Generalization

- Generalizes across cities and SLR scenarios with minimal fine-tuning.
- Hybrid loss and SLR-enhanced encoding enable robust adaptation.

Policy & Future Work

- Open-source datasets & code empower researchers and planners.
- Future works: integrate dynamic adaptation strategies, expand to more coastal regions, and couple with socio-economic models.

CASPIAN-v2 provides a scalable, accurate and interpretable framework for coastal flood prediction. By enabling rapid exploration of adaptation options, it supports data-driven decision-making in the face of climate change.

THANK YOU!