SAR Object Detection with Self-Supervised Pretraining and Curriculum-Aware Sampling

Yasin Almalioglu, Andrzej Kucik, Geoffrey French, Dafni Antotsiou, Alexander Adam, Cedric Archambeau

Helsing, London, UK   {first.last}@helsing.ai

Motivation & Contributions

- **Task:** Detect tiny objects (e.g., vehicles) in SAR satellite images.

- **Challenges:** Small size, severe class imbalance, noise, annotation scarcity.

- **Solution:** **TRANSAR** vision transformer:
  - Self-supervised pretraining (Masked Image Modeling, MIM)
  - Curriculum- and feedback-aware adaptive sampling scheduler
  - Auxiliary segmentation to enhance small object recovery
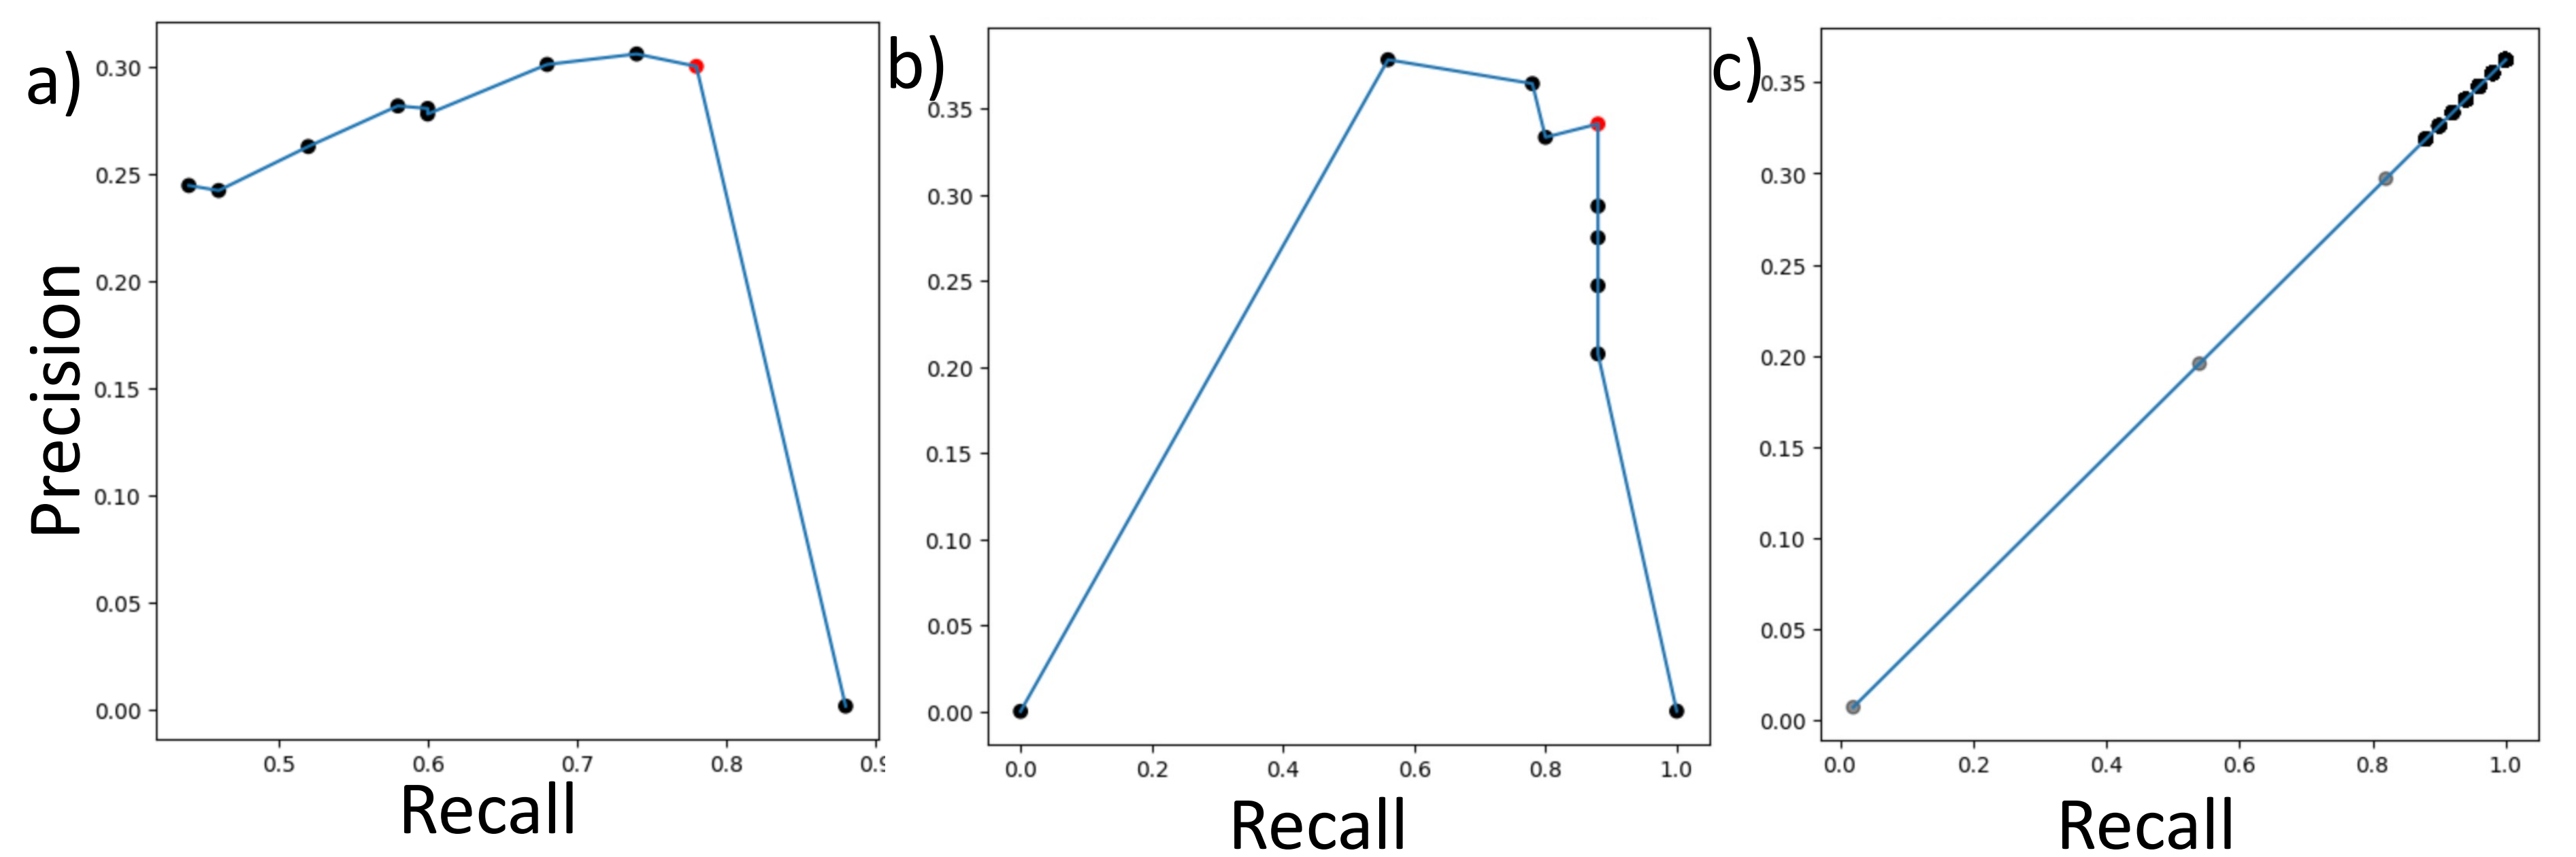
Adaptive Sampling Scheduler
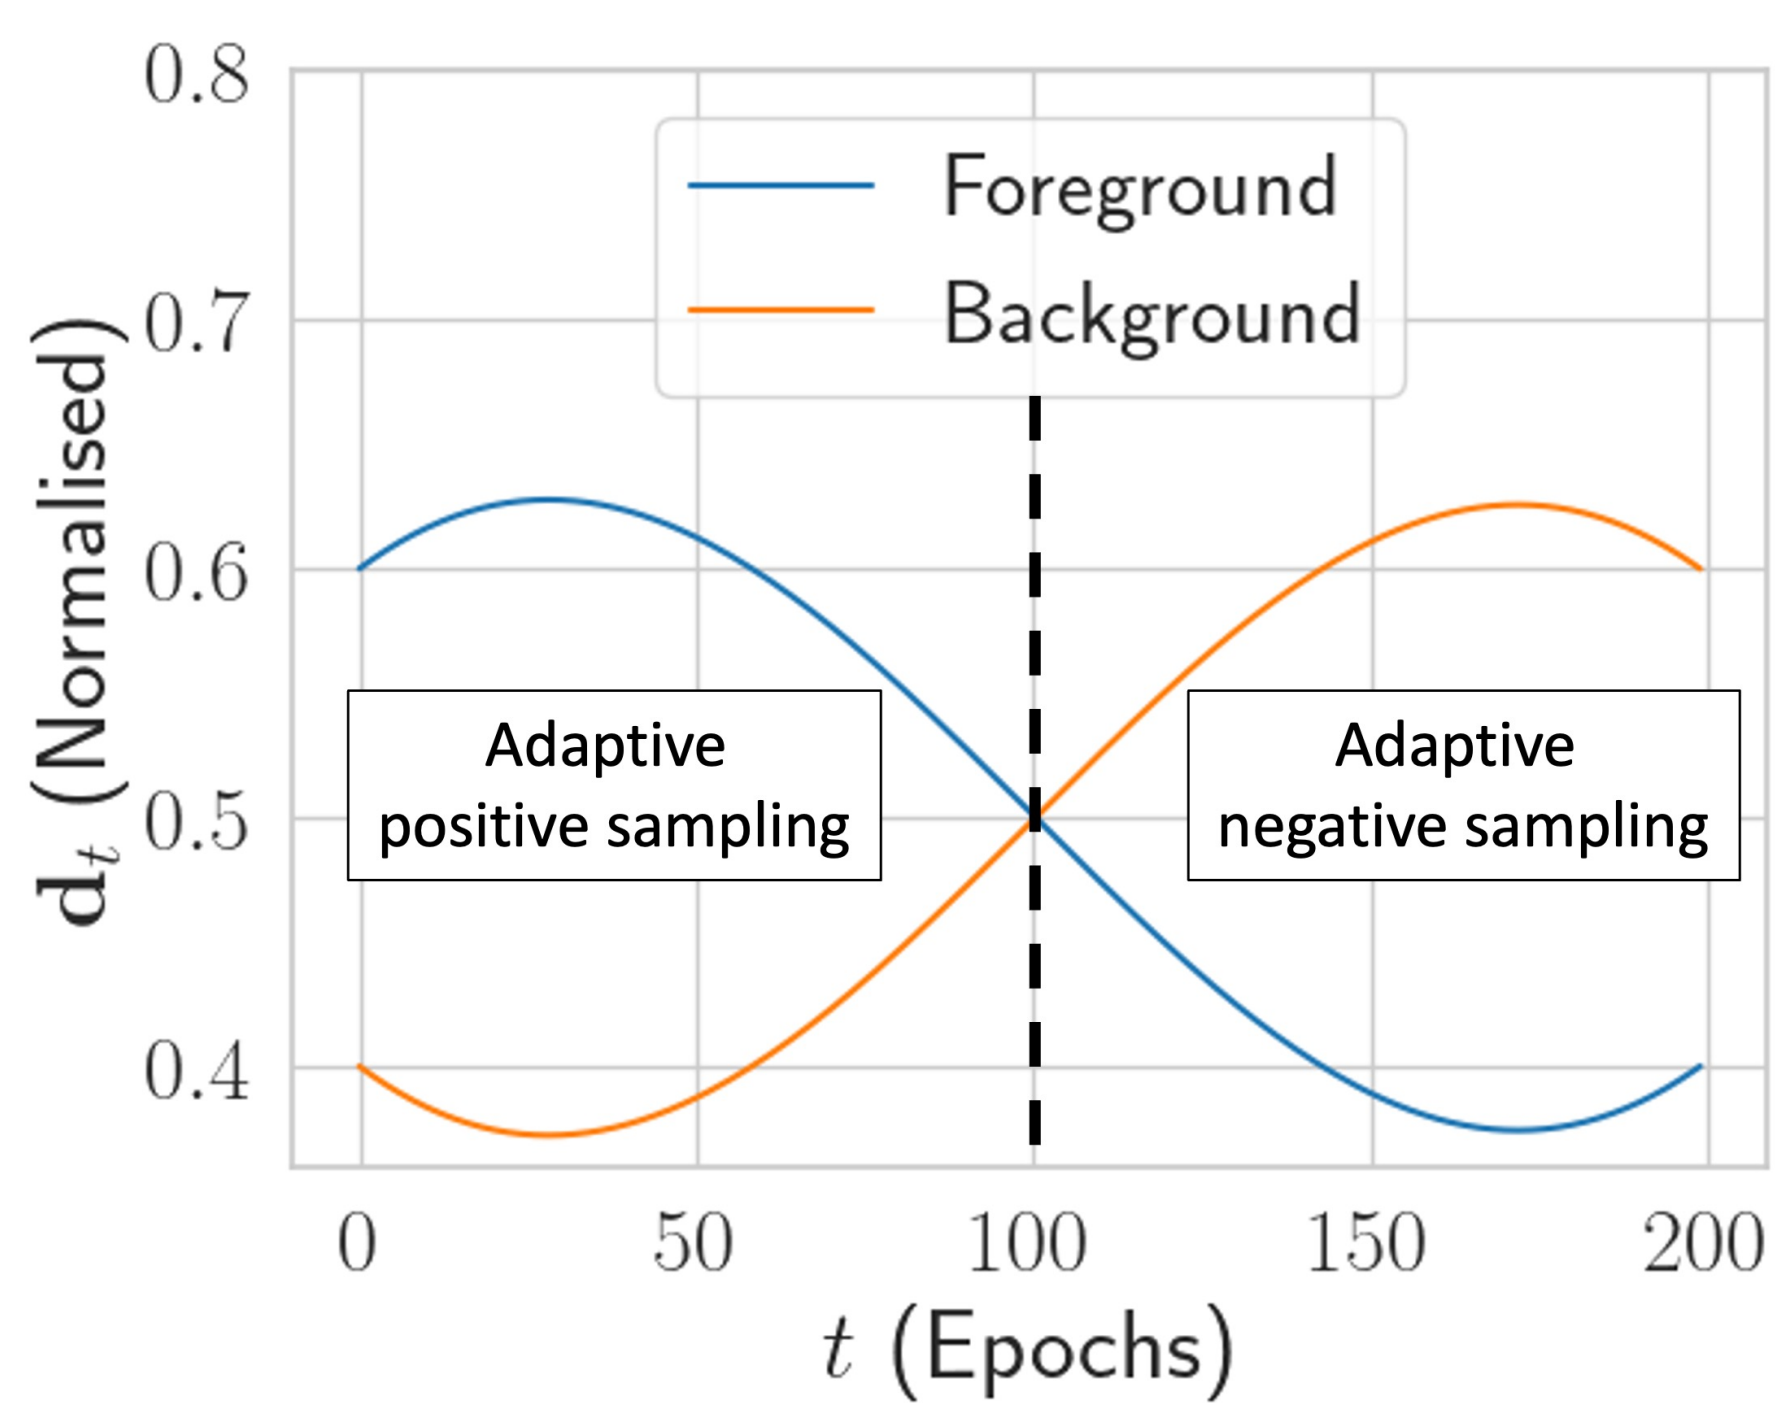
- Class imbalance tackled via curriculum-based scheduler.

- Gradually shifts sampling from real distribution → balanced.

- Real-time adjustment using model feedback (e.g., F1 score).

Pipeline Overview
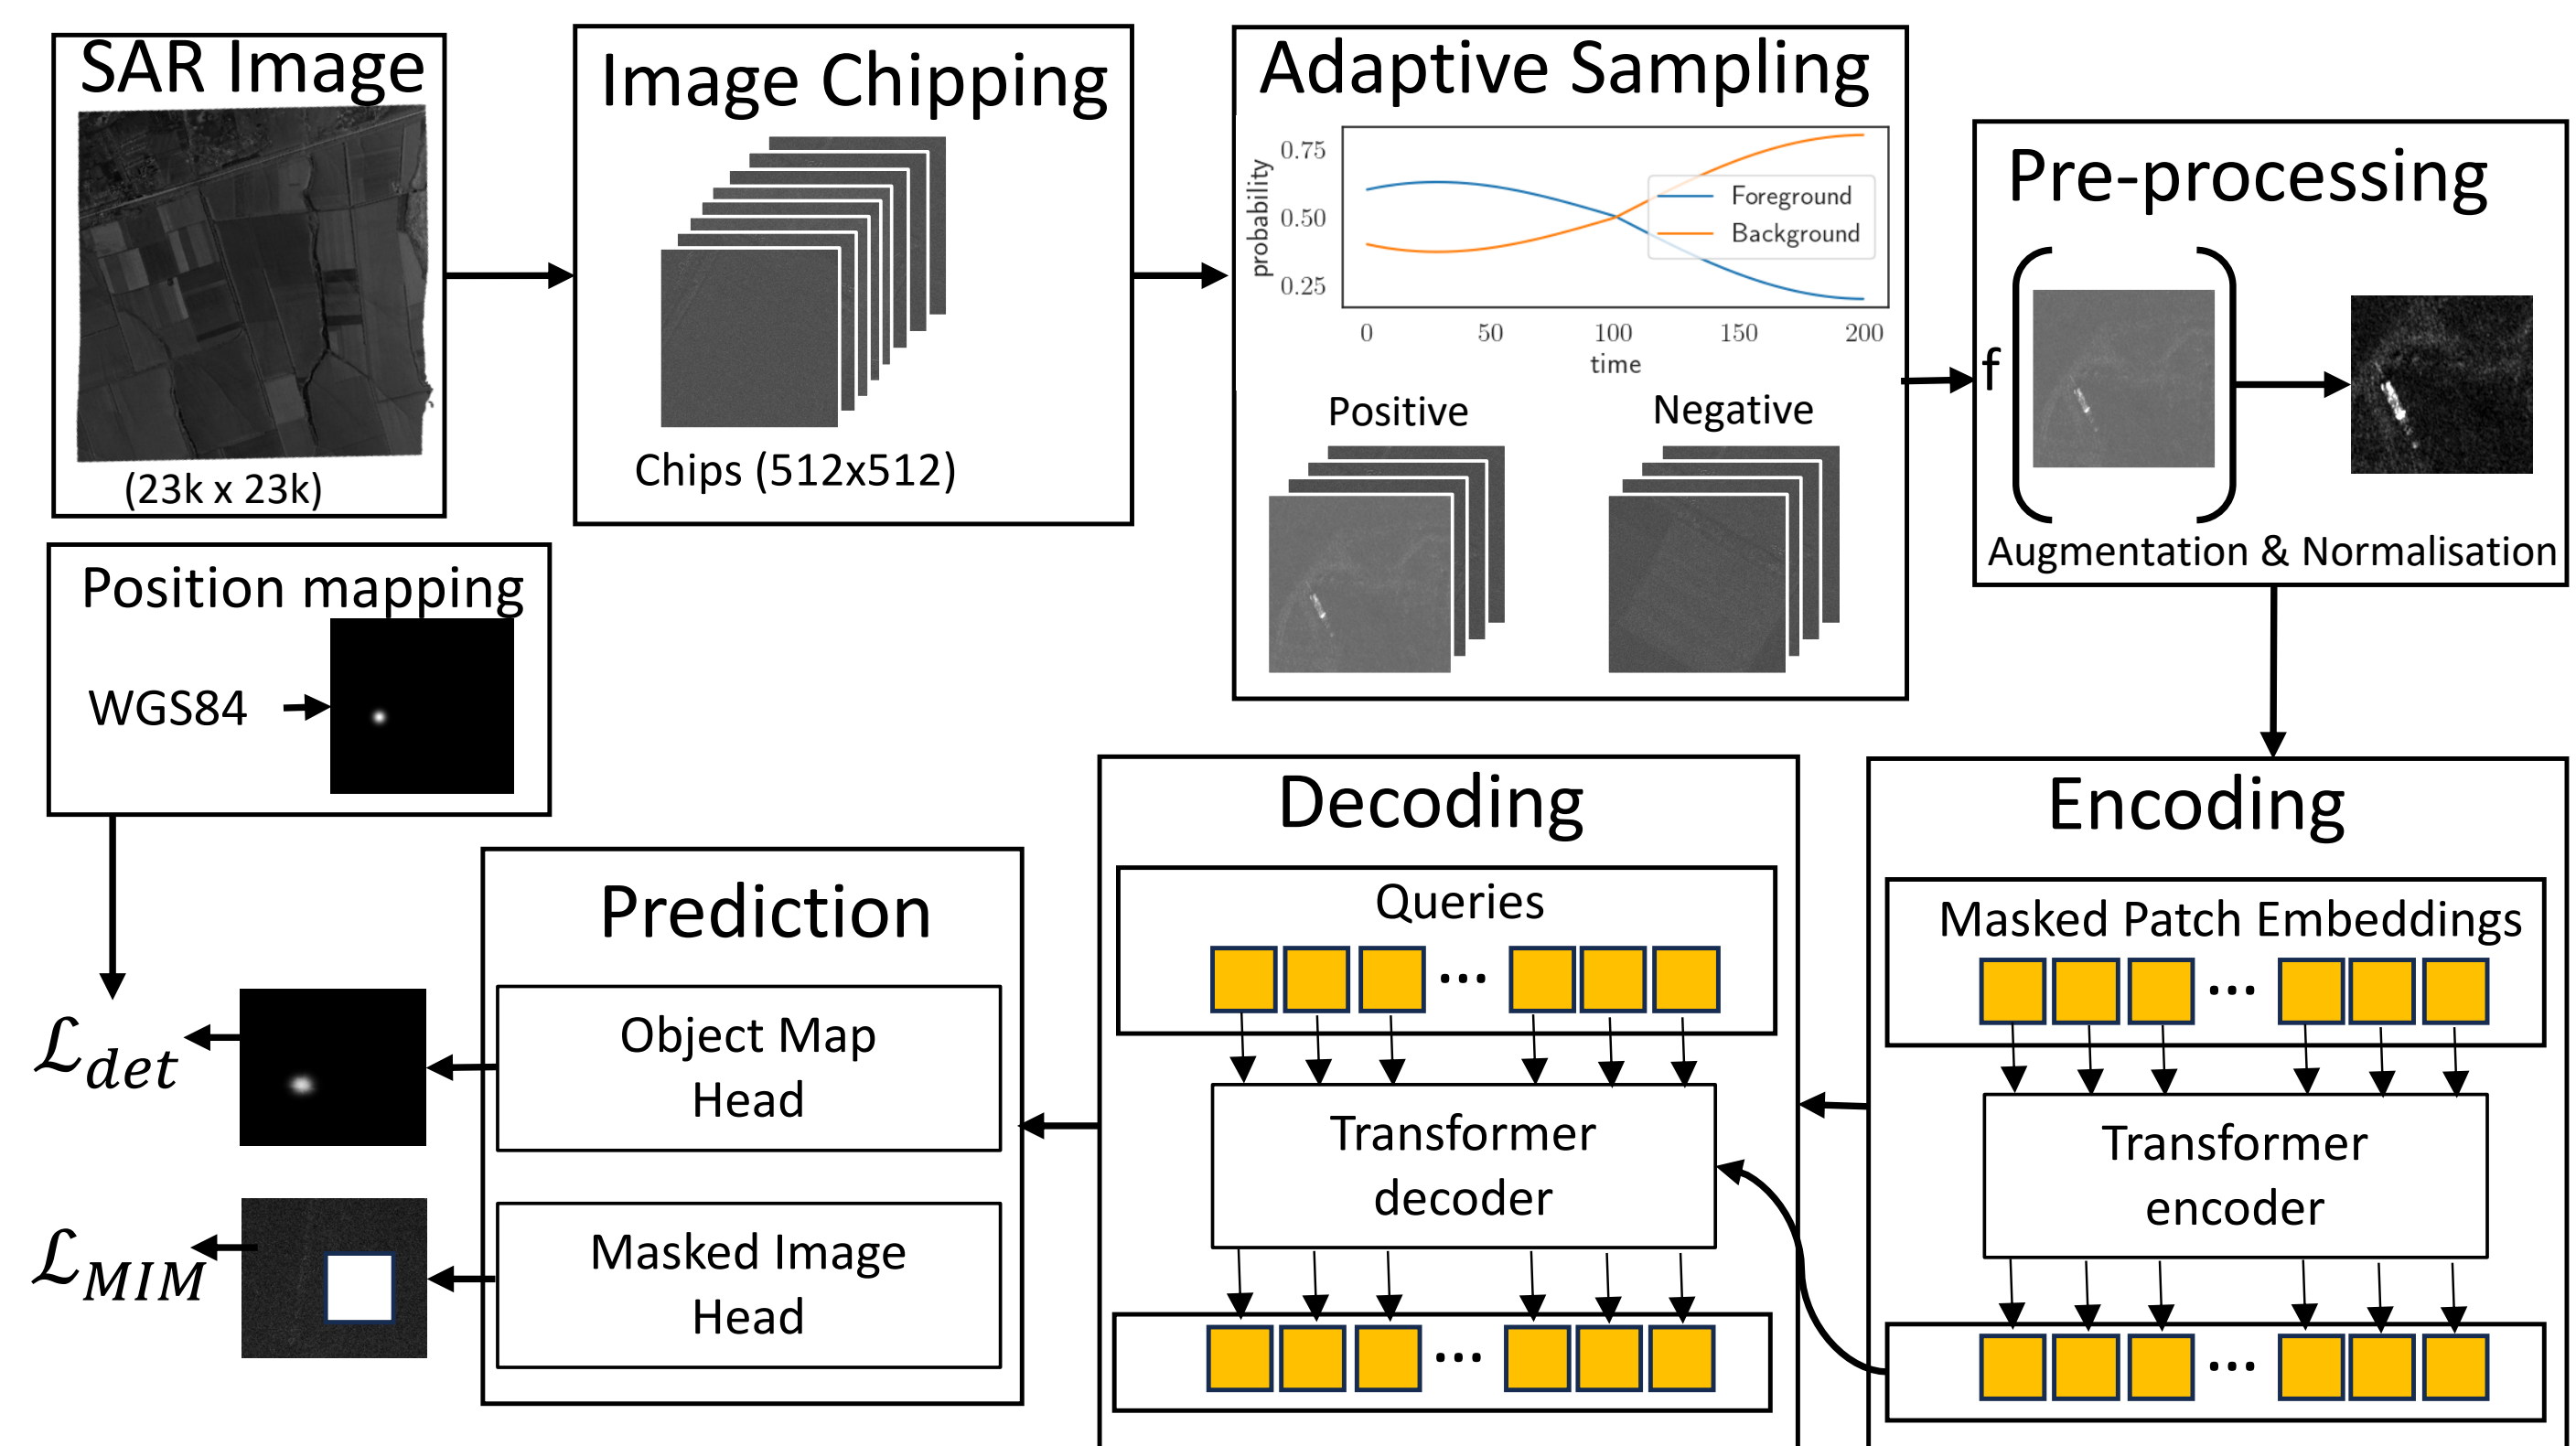
- **SSL Pretraining:** Learns structure from unlabelled SAR using block-wise masking.

- **Finetuning:** Detection head with adaptive sampling and auxiliary segmentation.

- **Output:** Probability heatmaps → peak detection → object locations.

Ablation

*Key findings:* Adaptive sampling and mask size crucial; optimal hyperparameters from ablation studies.

Benchmarking & Sensitivity

Detection Scores

	mAP	AP50	AP75	F1	Prec.	Recall
TRANSAR-large	**66.8**	**68.9**	**85.4**	**79.2**	77.9	80.5
ViT-MIM (SAR)	52.9	55.0	58.8	58.5	55.3	62.2
SegFormer (RGB)	35.8	36.4	40.9	40.6	38.6	42.8
UNet-SENet (SAR)	44.5	45.2	47.4	50.6	46.1	56.1

Ablation Table

Setting	F1 (%)
No Adaptive Sampling	61.6
Linear Scheduler	69.7
Cosine Scheduler	**79.2**
Block-mask size 8	**79.2**
Block-mask size 16	75.4
Block-mask size 32	70.6

TRANSAR outperforms prior baselines; ablations reveal crucial sampling and mask-size choices.

Qualitative Results
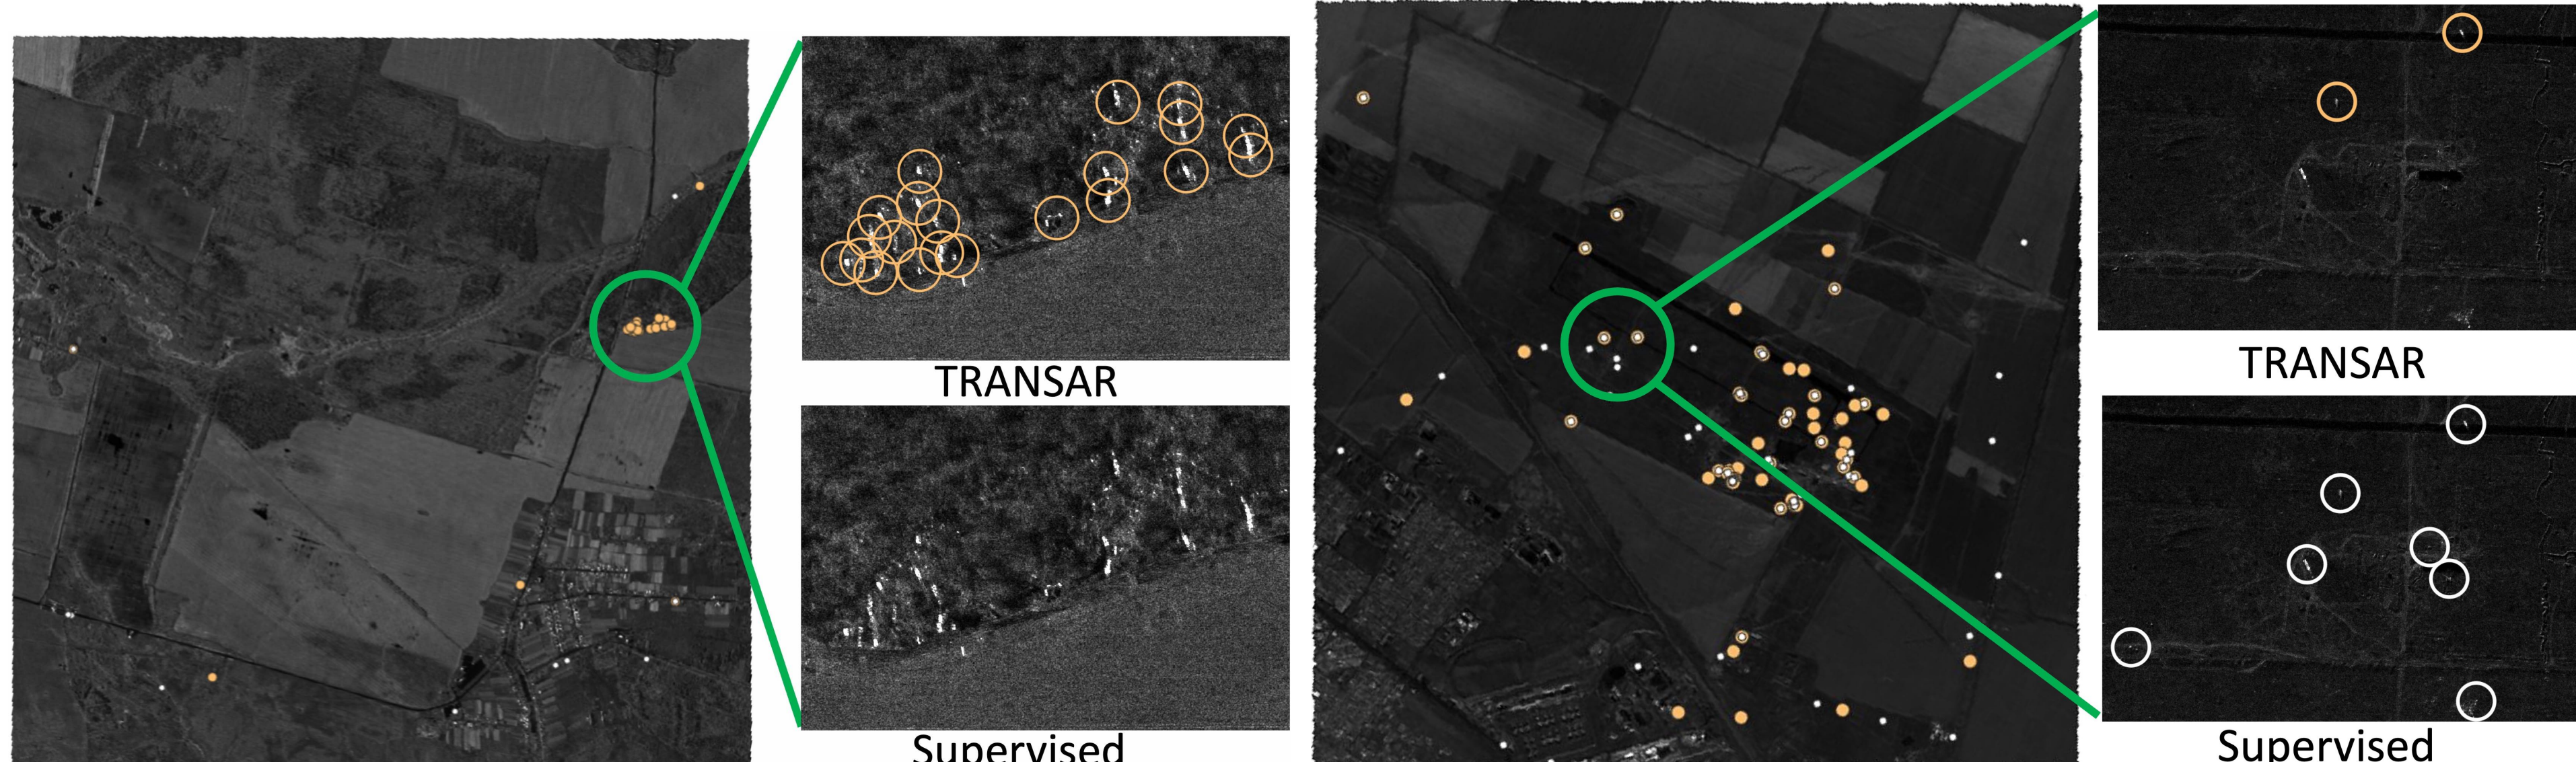
Fine-grained detections, robustness to false positives, higher precision (urban), parity in clear rural scenarios.

Takeaways & Outlook

- **Self-supervised ViTs** enable strong SAR detection even with few/zero labels.

- **Adaptive sampling scheduler** addresses data imbalance and maximizes F1.

- Outperforms prior SSL and supervised models on all key metrics.

- **Disaster relief, remote monitoring, urban analytics**—real-world impact.

- More annotated datasets and benchmarks needed for large-scale generalisation.