Session 3: Emerging AI Trends in Earth Observation

Foundation Models, Self-Supervised Learning, and Explainable AI

Stylianos Kotsopoulos
EU-Philippines CoPhil Programme

Agenda & Pacing

- **Duration:** 2 hours (120 minutes)
- **Plan:**
  - 0–10: Session goals & context
  - 10–50: Foundation Models (GeoFMs)
  - 50–80: Self-Supervised Learning (SSL)
  - 80–110: Explainable AI (XAI)
  - 110–120: Discussion & next steps

Learning Objectives

- Define foundation models and list major GeoFMs (Prithvi, Clay, SatMAE, DOFA)
- Explain SSL and its value for unlabeled satellite data
- Apply XAI concepts (SHAP, LIME, Grad-CAM) to EO tasks
- Decide when to use FM/SSL/XAI in PH contexts

Foundation Models (GeoFMs)

What & Why

- Pre-trained on massive EO archives → transferable to many tasks
- Fine-tune with 100–500 labeled samples instead of 10,000+
- Global priors capture seasonal, cultural, and atmospheric variation

Key Characteristics

Capability	Traditional CNN	Foundation Model
Pre-training data	Task-specific (10–100k labeled)	Global EO archive (10^8+ patches)
Fine-tuning labels	5–10k	100–500
Sensors handled	Usually single modality	Optical + SAR + DEM (model dependent)
Transfer across regions	Limited	High (learned invariances)
Compute demand (fine-tune)	Multi-GPU days	Single GPU < 2 hours

Why It Matters for the Philippines

- **Label scarcity:** national agencies rarely hold 10k+ labeled samples
- **Disaster tempo:** need flood/drought products within hours, not weeks
- **Heterogeneous landscapes:** models must generalize from Metro Manila to BARMM
- **Budget efficiency:** reuse one model for crops, mangroves, settlements

Major GeoFM Examples

- **Prithvi (IBM–NASA, 2023):** temporal ViT on Harmonized Landsat-Sentinel; excels on floods & wildfires
- **Clay (Clay Foundation, 2024):** multimodal (S1/S2/DEM) transformer; strong for land cover and change
- **SatMAE (Microsoft, 2022):** masked autoencoding on Sentinel-2; few-shot scene classification
- **DOFA (2024):** “Dynamic One-For-All” generalist spanning optical + SAR tasks
- **Google AlphaEarth (2024):** planetary-scale transformer for climate & EO forecasting (temperature, precipitation, emissions)

How FM Works

```
flowchart TB
  A[Massive Archive] --> B[Self-Supervised Pre-training]
  B --> C[Foundation Model]
  C --> D[Fine-tune (100-500 labels)]
  D --> E[Downstream Task]
```

- Pre-train once (expensive), fine-tune many times (cheap)

Philippine Use Cases

- **Rapid flood response:** Prithvi + 200 PAGASA/NAMRIA labels → 15 min flood masks
- **National crop mapping:** Clay + regional DA samples → 12M ha crop inventory
- **Climate resilience planning:** AlphaEarth-driven forecasts ingested by PAGASA climate services
- **Mangrove health:** SatMAE + 400 DENR plots → quarterly coastal monitoring

Self-Supervised Learning

Why SSL for EO

- 99.99% of satellite data unlabeled; SSL learns from raw archives
- Reduces annotation costs and accelerates “cold-start” regions
- Produces representations reusable for segmentation, detection, forecasting

Common SSL Strategies

- **Masked autoencoding:** hide 60–90% of pixels/patches, reconstruct context (SatMAE, Prithvi)
- **Contrastive learning:** pull together positive pairs (same location/time) vs negatives (BYOL, MoCo, SimCLR)
- **Temporal pretext tasks:** order scenes, predict next time step, detect anomalies
- **Cross-modal alignment:** align Sentinel-1 ↔ Sentinel-2 or EO ↔ DEM for richer features

Building an SSL Pipeline

1. Collect unlabeled Sentinel/Landsat/Planet imagery for AOI
2. Apply cloud/shadow filtering; tile into patches or time series
3. Choose pretext objective (masked, contrastive, temporal)
4. Pre-train backbone (ViT, CNN, LSTM) on GPUs/TPUs
5. Fine-tune backbone on mission-specific labels

Explainable AI (XAI)

Why XAI

- Builds trust for operational agencies (DRRM, agriculture, enforcement)
- Reveals bias or drift (seasonal, regional, sensor)
- Supports governance and documentation requirements

Techniques at a Glance

Technique	Works With	Insight Provided
SHAP	Tree models, tabular DL	Feature contribution per prediction
LIME	Any classifier/regressor	Local surrogate explaining single sample
Grad-CAM	CNN-based models	Heatmap of salient pixels/patches
Integrated Gradients	DNNs (incl. transformers)	Attribution along input path

Applying XAI in EO

- **Flood mapping:** Grad-CAM over segmentation masks to show river overflow focus
- **Crop yield models:** SHAP to rank NDVI, rainfall, soil moisture drivers
- **Illegal fishing detection:** LIME on vessel trajectories to explain alerts
- Pair XAI dashboards with validation from local analysts (LGUs, NAMRIA)

PH Integration & Roadmap

1. **Assess:** audit data readiness, choose priority missions (floods, crops, mangroves)
2. **Fine-tune:** adopt FM/SSL backbone, label 100–500 local samples
3. **Deploy:** package API/dashboard (DIMER, AIPI, Space+)
4. **Scale:** monitor drift with XAI, refresh labels, expand to new regions

Discussion & Q&A

- Where can FM/SSL/XAI help your agency most?
- Data readiness and labeling strategies
- Risks, governance, and explainability