Geo-AI for Environmental Monitoring

Dr. Shilpi Yadav
Department of Geography, Sophia College (Autonomous), Ajmer Raj.

***Abstract*— GeoAI – the fusion of geospatial information systems (GIS) and artificial intelligence (AI) – is revolutionizing environmental monitoring. By integrating advanced machine learning (ML) and deep learning (DL) methods with diverse spatial data (satellites, drones, sensors, etc.), GeoAI enables automated analysis of ecosystems at unprecedented scale and resolution. In air and water quality, multi-source fusion (satellite +ground) with spatio-temporal neural networks yields highly accurate forecasts of pollutants. For land cover and forestry, deep CNNs and segmentation models (e.g. U-Net, Deep Lab) process Landsat/Sentinel time series to map annual habitat change. In biodiversity, initiatives like NOAA's GAIA use very-high-resolution satellite imagery plus cloud-based ML to detect whales across oceans. GeoAI also advances disaster risk (floods, wildfires, storms) and urban analysis (heat islands) by extracting features (e.g. thermal indices from ECOSTRESS) and applying hybrid CNN+RNN/Transformer models.**

GeoAI workflows involve many data sources and steps: acquiring multi-spectral images (Sentinel-2, Landsat, MODIS, Planet, etc.), LiDAR point clouds, SAR data (Sentinel-1) and in-situ/IoT sensor readings; preprocessing (georeferencing, calibration, cloud masking) and feature engineering (vegetation indices like NDVI, water indices, spatial filters); constructing spatial–temporal models (3D CNNs, ConvLSTM, spatio-temporal GNNs, Transformer-based nets); and deployment (cloud platforms, mobile/edge devices, GIS integration). Open-source libraries (GDAL/Rasterio, PyTorch/TensorFlow) and tools (Google Earth Engine, QGIS, ArcGIS Pro) are used widely. Key challenges remain: heterogeneous data quality (cloud cover, noise), scale (petabytes of imagery), and biases (uneven station coverage) can limit model generalization. Model interpretability and regulatory compliance (e.g. EU AI Act's transparency requirements) demand explainable AI methods. Privacy and legal issues arise when using high-resolution imagery or mobile data. High-performance computing costs also constrain deployments. Best practices – data augmentation, cross-region validation, XAI auditing, privacy-preserving federated learning – are recommended, and future work is pushing toward large foundation models for geoscience, multimodal sensor fusion, and operational decision-support systems.

This paper provides a comprehensive survey of recent GeoAI methods and applications for environmental monitoring. We define GeoAI's scope, review AI/GIS techniques, summarize domain applications (air, water, land, biodiversity, disasters, climate, urban heat, agriculture), describe key datasets and sensors, outline common pipelines (with flowcharts and tables), and discuss tools, case studies, and open data. Major references are drawn from the last five years, including seminal journals and official reports.

***Index Terms*— Artificial Inteligence (AI); Remote Sensing; Geographical Information System (GIS); Environmental Monitoring.**

I. INTRODUCTION TO GEOAI AND ITS SCOPE

Geospatial Artificial Intelligence (GeoAI) merges AI algorithms with geospatial data to solve Earth-system problems. In essence, GeoAI "integrates advanced AI techniques with diverse geospatial data and methodologies". This synergy enables automated analysis of complex environmental processes and patterns that are otherwise too large or nuanced for traditional methods. For example, deep learning on satellite images can map deforestation or urban growth at continental scales; graph-based models can link sensor networks; and AI-driven analytics can fuse multi-source observations (imaging, meteorology, socioeconomics) into climate models.

GeoAI has emerged as a vital field because Earth observation datasets have grown explosively (billions of images, sensor networks, IoT) and compute power has soared (GPUs, cloud platforms). TorchGeo notes that "petabytes of freely available satellite imagery" and powerful GPUs are bringing AI to bear on climate and disaster challenges. Likewise, Ma et al. (2024) observe that new satellite missions (Landsat-9, Sentinel-2/1, GEDI LiDAR, PlanetScope) and ubiquitous UAVs have vastly expanded data availability. These rich data, combined with DL breakthroughs (e.g. convolutional and attention networks), allow GeoAI systems to learn fine-grained environmental features directly from raw data.

GeoAI: Definitions and Key Concepts

"GeoAI" has no single fixed definition, but broadly refers to applying AI/ML to geospatial problems. Song et al. (2025) describe GeoAI as "reshaping our understanding of Earth and urban systems by integrating advanced AI techniques with diverse geospatial data". Key elements include: spatial data (often remote sensing imagery, maps, GPS coordinates, etc.), AI methods (statistical ML, DL, etc.), and GIS concepts (projections, spatial analysis, vector/raster). GeoAI systems often combine remote sensing (satellite or aerial imaging) with spatial analysis and modern ML. It goes beyond classical GIS mapping by adding learning-based pattern recognition and prediction.

GeoAI spans both supervised and unsupervised AI. For example, land-cover maps can be generated by supervised DL segmentation; while unsupervised clustering or self-supervised techniques may discover anomalies (e.g. new fire zones) without labels. Emerging approaches include transfer learning (leveraging models trained on one region/sensor for another), active learning (iteratively querying unlabeled points for labeling by experts), and federated learning (training on decentralized data, e.g. satellites and IoT). Explainable AI (XAI) is gaining attention to interpret GeoAI decisions.

The geographical scope of GeoAI is global, but many applications focus on specific scales or regions. Crucially, GeoAI interlinks with Earth system science and policy: its outputs feed into climate models, conservation planning, and smart-city dashboards. By automating tasks like feature extraction and anomaly detection, GeoAI can accelerate environmental monitoring (e.g. rapid deforestation alerts) and support decision-making (e.g. urban heat mitigation strategies).

II. KEY AI AND GIS TECHNIQUES

GeoAI leverages the latest in AI/ML algorithms alongside traditional GIS methods. We summarize the main components:

- Machine Learning and Deep Learning: GeoAI uses both classical ML (Random Forests, SVMs, boosting) and DL (neural nets). Convolutional Neural Networks (CNNs) are standard for image data (land cover segmentation, object detection). Recurrent Neural Networks (RNNs) including LSTM/GRU are applied to temporal sequences of observations (e.g. time-series of satellite images or sensor readings). More recently, Transformers and self-attention architectures have been adapted to handle spatial data and long-range dependencies. Hybrid models combine these: e.g. CNN-LSTM or CNN-Transformer to capture both spatial and temporal patterns. Graph Neural Networks (GNNs) are a rising class of GeoAI models that operate on graph-structured data – for instance, modeling networks of environmental sensors or road/river systems. Statistical ML techniques (e.g. kriging, spatial regression) remain used, especially for geostatistical interpolation, but deep learning dominates cutting-edge GeoAI.
- Transfer Learning: Many tasks suffer from limited labeled data. Transfer learning – fine-tuning pre-trained models (often on generic image datasets or different geographies) – is widely used in remote sensing. Ma et al. (2024) note that deep nets often have millions of parameters, so source-domain training on large data (e.g. ImageNet or global satellite datasets) followed by task-specific tuning helps overcome scarce labels. Transfer approaches also address domain shift: differences in sensor types, seasons, or geography that can otherwise degrade model performance.
- Active Learning: To reduce annotation cost, GeoAI incorporates active learning: models select the most informative unlabeled samples for human labeling. For example, one study showed that an active learning CNN trained on just 20% of satellite image labels (chosen via uncertainty sampling) achieved 93.2% accuracy, saving ~80% of labeling effort compared to full annotation. This strategy is especially useful in environmental monitoring, where labels (e.g. precise species counts or pollutant measurements) are expensive. Active learning helps iteratively refine models while minimizing expert time.

- Federated Learning: Many geospatial datasets are held by different agencies or countries, raising privacy and sovereignty concerns. Federated learning allows training a global model using local data without sharing raw data. Moreno-Álvarez et al. (2024) review Federated Learning for remote sensing, noting it's an emerging privacy-preserving approach: collaborating entities train local models and aggregate them, enabling shared model building (e.g. for land cover classification) without centralizing sensitive data. Federated GeoAI can combine satellite data from multiple nations or distributed sensor networks while keeping each owner's data private.
- Explainable AI (XAI): Environmental decisions often require trust and transparency. Methods like saliency maps, concept activation vectors, or inherently interpretable models are being adapted to GeoAI. A recent review highlights the gap between raw performance and explainability: although ML in remote sensing has surged, interpretability lags behind. XAI helps validate model outputs (e.g. ensure a flood map isn't driven by spurious features) and aligns with regulations like the EU AI Act (which mandates transparency and fairness).
- GIS and Spatial Analysis: On the GIS side, GeoAI builds upon techniques in remote sensing, spatial statistics, and geostatistics. Remote sensing provides raster imagery (optical multispectral, thermal, synthetic aperture radar) and derived products. GIS operations include reprojection, mosaicking, and clipping to areas of interest. Spatial analysis (e.g. kernel density estimation, Moran's I, hotspot analysis) is often used before or after ML. For example, environmental variables may be interpolated by kriging, or elevation and proximity features computed in GIS and fed to ML models. Spectral indices (e.g. NDVI for vegetation, NDWI for water) are classic hand-crafted features from remote imagery that serve as inputs or benchmarks. Geostatistics provides tools (variogram modeling, point pattern analysis) that complement learning algorithms. In summary, GeoAI often uses ML to enhance or replace manual GIS workflows, but still relies on GIS foundations for organizing and interpreting the spatial data.

III. ENVIRONMENTAL DOMAINS AND CASE STUDIES

GeoAI has been applied across virtually all environmental monitoring domains. Below we highlight key examples and references in each area.

Air Quality

Air pollution monitoring uses satellites (for NO_2, CO_2, aerosols), ground stations, and models. Recent GeoAI work fuses multi-source data in hybrid DL frameworks. For example, Kalaiselvi et al. (2026) developed MAST-Net, a CNN+LSTM with attention, to predict urban concentrations of $PM_{2.5}$, PM_{10}, NO_2 and O_3. It ingests Sentinel-5P (TROPOMI), MODIS and Landsat-8 imagery along with meteorological and in-situ data. By dynamically selecting features and including uncertainty estimation, MAST-Net achieved 23–31% lower RMSE than traditional methods (correlation ~0.92 across pollutants). Other studies similarly report that deep spatio-temporal models outperform classical ML (e.g. ARIMA, Random Forest) by capturing complex pollutant dependencies. Key to accuracy is data fusion: combining satellite aerosol and trace-gas retrievals with ground sensors and meteorology. Anggraini et al. (2024) used multi-station time series + satellite imagery to create a global Air Quality Index system with improved coverage. In practice, GeoAI systems use convolutional layers to extract spatial pollutant patterns from images and recurrent layers for temporal trends. Transfer learning from global datasets (e.g. using pre-trained CNNs on large satellite collections) helps in regions with sparse stations. Citizen IoT networks also play a role: projects like AirQo deploy low-cost air sensors and use AI models on their data for local mapping.

Case Study: In India, a deep learning pipeline combined Sentinel-5P NO_2 maps, MODIS aerosol optical depth, ground $PM_{2.5}$ and weather features to predict daily pollution. The model provided city-wide forecasts 24–48h ahead, informing health advisories. In China, CNN-based mapping using TROPOMI has enabled near-real-time NO_2 hotspot detection in megacities. The table below summarizes common data in air monitoring.

Data/Platform	Spatial Resolution	Notes
Sentinel-5P (TROPOMI)	~3.5×7 km	Daily global NO_2, O_3, CO, CH_4, aerosol maps
MODIS (Terra/Aqua)	250m – 1km	Aerosol optical depth, fire anomalies, UV
Landsat-8/9	30m (visible bands)	Captures urban land cover for emission sources
Ground stations (EPA etc)	Point sensors	$PM_{2.5}$, PM_{10}, NO_2, O_3 measures
IoT networks (e.g. AirQo)	100–1000m	Dense, low-cost sensors for local AQ mapping

IV. WATER QUALITY

Monitoring lakes, rivers, and coasts involves estimating chlorophyll, turbidity, nutrient levels, etc. GeoAI approaches use optical and sometimes radar imagery. For instance, Kandasamy et al. (2025) proposed a hybrid deep learning system for chlorophyll-a in the Ganges River. They applied NASA's QAA-v5 algorithm to Sentinel-3 data to derive reflectance parameters, then used a pre-trained CatBoost regressor (ensemble tree model) to predict chlorophyll in real time ($R^2 \approx 0.985$). For forecasting trends, they then used time-series models (N-BEATS, CatBoost TS) for future predictions. This pipeline satellite retrieval + ML outperformed static models.

Water quality GeoAI often leverages spectral indices: NDWI and its variants for water presence, turbidity algorithms (NIR/Red ratios), or specialized bio-optical inversion models. ML regressors (Random Forest, XGBoost) and CNNs are trained on in-situ samples to map continuous quality metrics. For example, in reservoirs, deep nets applied to Sentinel-2 have extracted algal bloom extents from RGB bands. UAV/drone imagery also aids by giving high-res surface observations (e.g. plant/harmful algal bloom segmentation).

Case Study: In Qingcaosha Reservoir (China), researchers combined Landsat-8 spectral indices with ground measurements of BOD, COD, and WQI (Water Quality Index). A stacked autoencoder DL model ("VAW-Net") provided accurate water quality classification, detecting pollution changes more sensitively than classical classifiers. Similarly, global water quality frameworks (like the European Water Framework Directive) are beginning to incorporate satellite +AI alerts for deteriorating conditions.

Data/Platform	Spatial Resolution	Use-cases
Sentinel-2 MSI	10–60 m	Chlorophyll, turbidity, algal bloom mapping
Landsat (LM sensors)	30 m	Water index (NDWI), sediment transport
Sentinel-3 OLCI	300 m	Coastal water color (Chl-a, CDOM)
UAV Hyperspectral	~1 cm	Local blooms, pollutant plumes
In-situ probes	Point data	Validation (pH, turbidity, nutrients)
Sea/swim buoys	~<1 km	Continuous water level/salinity sensors

V. LAND USE AND LAND COVER (LULC) CHANGE

Mapping LULC and its changes (urbanization, crop rotation, deforestation) is a classical GIS task now enhanced by DL. Recent efforts aim for annual, high-resolution maps using deep neural networks on the long Landsat/Sentinel archives. For example, the USGS Land Cover AI project (LCAMS) uses multiple CNN-based models on Landsat time series to produce annual 30-m national maps of thematic cover and impervious surfaces. Their pipeline has "spatial classification" (per-scene CNN), "spatiotemporal refinement" networks, and regression for impervious cover over time. These models integrate knowledge from past NLCD products and LCMAP, exploiting temporal consistency (harmonic analysis) to improve accuracy.

Worldwide, many groups apply U-Net, SegNet or DeepLabV3 on high-res images. For instance, CNNs trained on Sentinel-2 have achieved ~80–90% accuracy in land cover categories (forest, water, urban, etc.) in Europe and Asia. Time-series models (3D-CNN or CNN+LSTM) detect change by learning temporal patterns of reflectance; they pinpoint deforestation or new constructions with multi-year data. Foundation datasets like ESA's WorldCover use supervised CNNs to release global LULC maps at 10-m resolution annually. Pitfalls include mixed pixels and class imbalance (e.g. urban areas are rarer), so techniques like class weighting and data augmentation are common.

Case Study: A recent MDPI study used high-resolution PlanetScope and Sentinel imagery with a LinkNet-ResNet34 model for African savanna land cover. The network, pre-trained on ImageNet then fine-tuned on labeled patches, achieved 87% overall accuracy in distinguishing cropland, forest, grassland, and settlements. The resulting map revealed post-2015 cropland expansion trends. Similarly, Google Earth Engine users have produced global decadal crop maps using ML on fused Landsat/Sentinel stacks.

VI. DEFORESTATION AND FOREST MONITORING

Forest loss tracking is one of the earliest GeoAI use-cases. Automated segmentation algorithms on satellite data detect tree-cover change faster than manual surveying. Wieland et al. (2019) and others note that high-resolution (0.5–5m) data from Planet, Maxar, or UAVs, combined with CNN segmentation (U-Net, SegNet, DeepLab) can identify clear-cuts or burn scars. For example, one systematic review found the majority of deep learning deforestation studies use CNNs trained on labeled patches from aerial/satellite imagery. U-Net is especially prevalent (shown 25+ times in the literature) because of its accuracy in pixel-level classification.

In practice, a GeoAI deforestation workflow might take bi-annual or annual SAR or optical images and run a change-detection network. Multi-temporal CNNs or Siamese nets (comparing imagery at two dates) highlight lost canopy. The DeepForest tool (a PyTorch library) uses ResNet and Efficient Net for tree detection in aerial imagery. Governments and NGOs use these methods to generate near-real-time alerts; e.g. Brazil's INPE uses DL on Planet data to flag Amazon deforestation hotspots. Active learning is used to refine maps: the model queries experts on ambiguous patches, rapidly improving map accuracy.

Case Study: LiDAR also plays a role: by capturing 3D forest structure, it can help GeoAI distinguish old-growth from new regrowth. A recent approach fused Sentinel-1 SAR and GEDI LiDAR with a hybrid CNN to forecast above-ground biomass changes, enabling early detection of logging. Overall, GeoAI has accelerated forest monitoring from monthly or annual mapping to near-daily alert systems. For example, the Global Land Analysis & Discovery (GLAD) lab and Planet Data have used nightly datasets to detect deforestation alerts in minutes.

VII. BIODIVERSITY AND WILDLIFE MONITORING

Monitoring species and habitats is challenging due to scale and access. GeoAI contributes by using remote sensing and automated imagery analysis for biodiversity proxies. NOAA's GAIA project is a flagship example: it integrates very-high-resolution (sub-meter) satellite imagery with cloud-based ML to detect whales in the ocean. GAIA's system provides a platform for experts to annotate whales in images, and trains detectors (CNNs) that can scan vast ocean areas for endangered species. The NOAA site notes GAIA uses AI and geospatial analysis to locate North Atlantic right whales, with plans to expand to other species.

Satellite image-based biodiversity monitoring extends beyond whales. For example, AI models on aerial images can count individual trees or classify species in a plot. In tropical forests, CNNs trained on multispectral drone images identify species groups and detect habitat fragmentation. Camera-trap and acoustic networks also benefit: ML (including CNNs for images or spectrograms) automatically detect animal presence from vast passive-monitoring data. Another thread uses habitat mapping: linking land cover change to species loss using geospatial data in joint AI models.

Figure: GeoAI uses very-high-resolution satellite imagery for species monitoring. (Image credit: NASA/NOAA)

Figure: Satellite whale monitoring (NOAA GAIA). Very-high-resolution satellite imagery (e.g. this NOAA photo of Earth from space) is now used to track wildlife. NOAA's GAIA project combines VHR optical images with AI in the cloud to detect whales. The model and platform enable conservationists to survey remote oceans efficiently.

VIII. DISASTER RISK AND EARLY WARNING

Natural hazards (floods, fires, storms, landslides) are prime GeoAI targets for damage assessment and early warning.

- Flood mapping: CNN-based segmentation of SAR or optical images is used to delineate floodwater. For example, DL flood mapping on Sentinel-1 SAR images can achieve >90% pixel accuracy, enabling rapid flood extent maps during hurricanes.

- Fire detection/prediction: ML models fuse vegetation fuel maps (from LiDAR or NDVI) with weather data. Recent reviews classify wildfire-risk models into image segmentation (spot active fires), time-series forecasting (daily fire probability), and spatio-temporal prediction. Xu et al. (2025) note that features include fuel loads, topography, climate, and that explainability is important. One study used a CNN-LSTM hybrid on MODIS and climate data to forecast regional wildfire risk, outperforming static indices.
- Storm impact and landfall: GeoAI analyses satellite imagery (e.g. GOES radar, Sentinel) to estimate hurricane intensity and track by combining convolutional features with physics-based models.
- Landslides: Terrain and precipitation data fed into random forests or neural nets predict slope failures; graph models can capture river- or road-networks linked to slides.
- Earthquakes: While primarily measured by seismic networks, satellite InSAR plus AI is used to map surface deformation patterns.

Case Study: The NASA-ISRO SAR (NISAR) mission's data, when available, will fuel global deep learning flood models. Early experiments on Sentinel-1 show CNNs classifying inundated pixels. Similarly, CNNs trained on Landsat images identified burned areas in California wildfires with >95% accuracy, aiding rapid damage assessment. These systems often include uncertainty estimation (e.g. Monte Carlo dropout) to flag unreliable areas and are integrated into disaster dashboards. Xu et al. emphasize that multi-modal models (fusing imagery, topography, climate) and coupling with numerical weather prediction are promising directions.

Climate Change Impacts

GeoAI supports monitoring climate-driven changes. Cryosphere: AI analyzes glacier imagery to track retreat. Geo youth efforts highlight combining ML with multi-sensor data to detect glacial changes and emerging proglacial lakes in near-real time. For instance, one recent method detects glacier boundaries in Sentinel-2 images using a U-Net, revealing rapid ice loss. A GEO press release notes that GeoAI and satellite EO are "transforming how we monitor glacier change" in remote regions.

- Sea-level rise and coasts: Coastal inundation and erosion are mapped by segmenting shorelines in multi-temporal imagery. The USGS shows that processing Landsat time series can now extract high-resolution shoreline changes, matching field surveys. This effectively shifts coastal monitoring "from a data-scarce to a data-abundant era".
- Ecosystem shifts: Habitat transformations (e.g. forest dieback due to drought) are tracked by time-series anomaly detection in NDVI or by training models to recognize climate stress signatures. Carbon flux: DL models estimate forest biomass from LiDAR and SAR to monitor carbon sequestration trends. Permafrost and land-carbon: Machine learning on thermal-infrared satellite data can identify thaw lakes or permafrost melt patterns.

Case Study: A deep learning model applied to Landsat images identified newly exposed land from retreating Himalayan glaciers, generating alerts for downstream communities. Such tools are increasingly used in climate adaptation: e.g. Colombian Andes region employed CNNs to map glacier lakes for flood risk. Overall, GeoAI turns the "petabytes" of climate-affected imagery into timely information for resilience planning.

Urban Heat Islands

Monitoring urban heat islands (UHI) relies on thermal imagery. NASA's ECOSTRESS (on the ISS) provides ~70m LST (land surface temp) data. A NASA tutorial describes processing ECOSTRESS data in Earth Engine and applying a machine-learning downscaling to 10m maps. In practice, GeoAI uses random forests or CNNs to fuse thermal bands with high-res imagery (Landsat, Sentinel-2) to produce street-level heat maps. These models identify hotspots (e.g. parking lots, rooftops) and predict human exposure risk. The Earthdata training notes that ECOSTRESS LST has been used with RF downscaling to pinpoint heat-vulnerable communities.

Urban planners use these maps for mitigation (e.g. adding green roofs). For example, in Houston the 10m thermal map (downscaled via ML) revealed neighborhoods with >5°C higher nighttime temperatures, guiding tree-planting programs. CNNs trained on aerial night-time images (like VIIRS data) are also used to infer UHI intensity in absence of direct thermal imagery.

Agriculture and Precision Farming

Precision agriculture uses GeoAI to optimize inputs and forecast yields. The state-of-the-art is reviewed by Muhammad et al. (2026): they note that deep learning models consistently outperform traditional methods for crop yield estimation, thanks to learning hierarchical spatial features. Remote sensing (multispectral satellite, UAV, drones) provides data on crop health (NDVI, LAI, etc.) and phenology; weather and soil sensors provide context. DL architectures (CNNs, sometimes with LSTM for temporal yield trends) map these inputs to crop yield or stress detection.

Examples include CNNs that estimate grain yield from Sentinel-2 time series in wheat and rice fields, achieving $R^2>0.9$ with minimal ground sampling. Other cases use U-Net on drone images to detect plant diseases (like rust or blight) with >85% precision. Transfer learning is common: e.g. a model pre-trained on global cropland imagery is fine-tuned on a local farm's fields. Active learning is applied by letting farmers label a few worst-predicted areas.

Case Study: Aman et al.'s review shows that DL models (CNNs, RNNs) have revolutionized yield forecasting: e.g., a ConvLSTM model for maize fields in China reduced mean absolute error by 25% over a traditional regression. Another example: the Wheat Yield Prediction Challenge achieved top results using ensemble DL models trained on multi-year Landsat data. Open-source tools like RasterVision and TorchGeo are increasingly used to build these pipelines (Table 2).

Data/Platform	Spatial Resolution	Use-cases
Sentinel-2 MSI	10–20 m (optical)	Crop health indices (NDVI, NDRE), crop type
Landsat-8/9 OLI	30 m	Large-scale yield and field mapping
PlanetScope	~3 m	Field-level monitoring (e.g. vegetable farms)
UAV (multispectral)	cm-scale	Disease/weed detection, phenotyping
Soil/Weather IoT	point, hourly	Inputs for growth models (humidity, temp)
Yield/Pheno Data	field-scale	Ground truth for training and validation

IX. DATA AND SENSORS

GeoAI relies on diverse remote sensing datasets and sensors. Key sources include:

- Satellite Imagery:
- Sentinel-2 (Copernicus) – 10–20 m multispectral (13 bands), 5-day revisit. Widely used for vegetation, land cover, coastal water quality.
- Landsat 4–9 (USGS/NASA) – 30 m multispectral (11 bands), 16-day revisit. Longest-running land imaging record; foundation for many ML models.
- MODIS (Aqua/Terra) – 250m–1km, daily global. Used for time-series climate signals (e.g. fire, aerosols).
- Commercial VHR: e.g. PlanetScope (~3m), Maxar (~0.3–1m), WorldView, etc. These provide very high resolution for local studies (e.g. tree crown counting) but often require purchase.
- Thermal: ECOSTRESS (70m), VIIRS (375m, night-light), Landsat TIRS (100m) for LST/UHI.
- SAR (Radar): Sentinel-1 (10 m, C-band, 6-day), ALOS PALSAR, etc. Radar penetrates clouds; crucial for all-weather mapping (floods, deforestation). SAR time series enable oil spill and flood detection with ML.
- LiDAR: Airborne laser scanning (point clouds) and spaceborne (GEDI, ICESat-2) provide 3D structure. ML on LiDAR is used for forest biomass, terrain modeling, and now combining with optical imagery for more accurate feature extraction.
- Aerial and UAV Imagery: High-resolution orthophotos (e.g. US NAIP – 1m RGB/IR) and drones (cm-scale) are used for fine-grained tasks (e.g. tree species, crop stress). These require special GIS libraries but offer detailed views.
- In-situ Sensors: Fixed networks measure environment at point locations (e.g. air quality stations, water monitoring buoys, meteorological stations). These data are often used as labels or fused inputs. Low-cost IoT sensor networks (e.g. community air sensors, soil moisture probes) are rapidly expanding; projects like AirQo deploy LoRaWAN sensor nodes and use AI for city-scale pollution maps.
- Citizen and Mobile Data: Geo-tagged photos, smartphone apps (e.g. iNaturalist observations), and crowd-sensed data can complement satellite

data for on-the-ground monitoring, though quality control is needed.

Ma et al. (2024) summarize these sources: for example, Sentinel-1/2 and Landsat are ubiquitous, SAR complements optical (cloud-penetration for forests, floods), LiDAR (GEDI) adds elevation/canopy structure, and UAVs are increasingly used for precision farming. They note that advances in sensors mean more modalities (hyperspectral imagery, smartphone imagery) entering GeoAI.

Data preprocessing is critical: satellite images are often level-2 (atmospherically corrected) but still require co-registration and mosaicking. Specialized libraries like GDAL/Rasterio handle file formats (GeoTIFF, NetCDF, Shapefile). Common preprocessing steps include atmospheric correction (e.g. using ACOLITE, SMAC), cloud/cloud-shadow masking (e.g. FMask), and reprojection to a common coordinate system. After spatial alignment, features such as vegetation indices are computed (e.g. NDVI = (NIR−RED)/(NIR+RED)). Multi-temporal stacks are assembled (e.g. seasonal composites or time series) and, if needed, downsampled or aggregated to match ground data. For ML, raster data are often tiled into smaller patches or chips due to memory limits, with augmentation (rotations, flips, spectral jitter) applied to improve generalization.

X. MODELING AND ANALYSIS

GeoAI modeling involves spatial, temporal, and spatio-temporal architectures, training strategies, and rigorous evaluation. Key components include:

Model Architectures

- Convolutional Neural Networks (CNNs): The backbone of GeoAI. CNNs (and fully convolutional variants like U-Net, FCN) excel at spatial feature extraction from imagery. They capture textures, shapes, and local context via learned filters. In practice, pre-trained CNN encoders (ResNet, EfficientNet) are fine-tuned on geospatial tasks. CNNs are used for per-pixel classification (land cover mapping) and object detection (e.g. tree or building detection with YOLO/RetinaNet). Their strength is handling Euclidean grid data, but they cannot directly model irregular graph data or long-range context without many layers.
- Recurrent Neural Networks (RNNs/LSTM): Designed for sequences, RNNs process ordered data. In GeoAI, LSTM or GRU layers capture temporal patterns from time series of sensor data or multi-date imagery. For example, a 1D RNN can forecast future pollutant levels from past readings, or an LSTM can process a sequence of NDVI images to predict yield. RNNs are powerful for temporal correlation but by themselves have limited spatial perception.
- Transformers: The attention-based architecture now enters GeoAI. Vision Transformers (ViT) and spatio-temporal transformers handle global context by relating all pixels/patches. For instance, a transformer can learn that a change in forest pixels far apart may be correlated. Transformers have been adapted to remote sensing (e.g. a Conv-attention hybrid on Sentinel data) and are promising for fusing multimodal data (satellite + text + sensors). They require large data and compute, but offer flexibility in modeling complex spatio-temporal dependencies.
- Graph Neural Networks (GNNs): GNNs operate on graph-structured data and non-Euclidean spatial relations. For example, you can construct a graph where nodes are river gauges or weather stations and edges represent hydrological flows or proximity. GNNs can propagate information across such networks to model diffusion of pollutants or heat. Recent work surveys show GNNs applied to climate, disaster response, air quality, and agriculture in EO. GNNs' strength is in modeling relationships beyond regular grids (e.g. social-ecological networks, road networks), but they require explicit graph construction and can be computationally heavy on large graphs.
- Hybrid Models: Many GeoAI systems combine architectures to capture different aspects. Common hybrids include CNN+LSTM (for spatio-temporal data), CNN+Transformer, or multi-branch models fusing imagery and tabular data. For instance, an encoder CNN extracts spatial features from an image, then an LSTM ingests that plus time stamps. Transformer-based fusion may attend jointly over spatial patches and temporal tokens. Ensembles of models (averaging multiple networks) are also used to improve robustness.

Table 1 compares these model types in GeoAI contexts:

Model / Architecture	Typical Use-Cases	Strengths	Limitations
CNN (U-Net, ResNet)	Image classification, semantic segmentation (e.g. LULC, deforestation)	Excellent at spatial feature learning (textures, shapes); well-supported in frameworks	Requires large training sets; less explicit temporal modelling
RNN / LSTM	Time-series forecasting (e.g. air quality, crop yield over time)	Captures temporal dependencies; handles sequence data	Struggles with long-range spatial correlations; prone to vanishing gradients
Transformer	Multimodal fusion, long-range context (e.g. global climate trends)	Models' global relationships via attention; flexible to modalities	Data-hungry, high compute; may ignore local details if not combined with CNN
GNN	Graph-structured data (e.g. sensor networks, road maps, river systems)	Encodes non-Euclidean relations; integrates heterogeneous data nodes	Needs explicit graph; scalability issues on very large graphs; less common code support
Hybrid CNN+RNN	Spatio-temporal modelling (e.g. satellite image sequences)	Leverages strengths of both spatial and temporal models	More complex training; risk of overfitting

Table 1: Comparison of common GeoAI model architectures, their use-cases, strengths, and limitations.

XI. TRAINING STRATEGIES

- Data Augmentation: Rotations, flips, spectral shifts, and noise are applied to remote sensing images to increase training diversity. Time-series can be augmented by synthetic temporal shifts. Augmentation helps models generalize across seasons and sensors.
- Transfer Learning (TL): As noted, TL is crucial in GeoAI to cope with limited labeled data. Models pretrained on one area (or even non-EO data like ImageNet) are fine-tuned on target tasks. Domain adaptation techniques (e.g. aligning feature distributions) further improve TL across regions or sensor types.
- Active Learning: Selecting informative samples (e.g. edge cases, high-uncertainty areas) for human labeling reduces annotation effort. This is used in iterative mapping: the model flags pixels it is most unsure about, an expert labels them, and the model is retrained.
- Federated and Distributed Training: For privacy or large-scale, models may be trained across multiple nodes. For example, separate servers could train local models on country-specific data, then securely aggregate weights (Federated Averaging) to form a global model. This approach is still experimental in GeoAI but is gaining traction.
- Explainability and Regularization: During training or post-analysis, techniques like SHAP values or Grad-CAM are used to explain model decisions. Regularization (dropout, weight decay) and Bayesian DL (e.g. MC-dropout) provide uncertainty estimates, useful for risk-sensitive environmental applications.

Evaluation Metrics

Model evaluation depends on task type. Classification/segmentation metrics include accuracy, precision, recall, F1-score, and Intersection-over-Union (IoU) for pixel-wise tasks. For object detection, average precision (AP) is used. Regression metrics (for continuous outputs like pollutant concentration) use RMSE, MAE, and R^2. For example, in the air quality study, RMSE and correlation coefficients (0.91–0.94) quantified performance gains. In crop yield forecasting, MAE or percentage error are typical. Geospatial problems often require spatially aware validation: e.g. hold-out regions or time-splits to avoid overfitting to one area. Uncertainty metrics (prediction intervals, entropy) are also reported in advanced studies to measure confidence.

Deployment

GeoAI systems are deployed on various platforms:

- Cloud Platforms: Google Earth Engine (GEE) is a widely-used cloud environment with petabytes of data and built-in ML tools. Although not open for citation here, it is often the backbone for large-scale processing. Major GIS vendors (Esri) now offer cloud notebooks (ArcGIS Online with Python) that integrate DL libraries.
- Edge Devices: For field applications (UAVs, sensor nodes), lightweight models are deployed on edge hardware (NVIDIA Jetson, Raspberry Pi

with Coral TPU). For example, a UAV could run an onboard CNN for real-time wildfire detection. However, edge deployment must consider model size and energy constraints.

- GIS Integration: Many organizations integrate AI into GIS workflows. ArcGIS Pro provides deep learning tools (building extraction, classification) via ArcPy or ModelBuilder, using local or cloud training. QGIS (open-source) uses plugins (e.g. Orfeo Toolbox) to call ML models. Rasterio/GDAL pipelines can feed predictions back into GIS rasters. APIs and web services (REST, OGC WPS) allow GeoAI models to be called from spatial databases or GIS clients.
- APIs and Mobile Apps: Some GeoAI models are exposed as web APIs. For instance, a FloodMapping API might accept coordinates and dates to return inundation maps. Mobile apps (e.g. CAJAL for crop advisors) can request satellite-based analytics (e.g. NDVI time series) for farmers.

XII. OPEN-SOURCE TOOLS, LIBRARIES, AND DATASETS

Numerous tools support GeoAI development. We highlight some widely used ones (see Table 2):

- Google Earth Engine (GEE): A cloud platform with scalable access to satellite imagery and geospatial analysis tools. Users can script analysis and some ML tasks in JavaScript or Python, leveraging GEE's infrastructure (though recent ML integration is more manual). (See e.g. [34†L77-L85] for an example of multi-sensor fusion that could run on such platforms.)
- QGIS: Free, open-source GIS desktop. With Python plugins (e.g. processing tools, semi-automated classification) and integration with GDAL/Rasterio, it can run ML models via Python scripts or external calls. It also supports plug-ins for remote sensing analysis (e.g. Orfeo ToolBox).
- ArcGIS Pro: Commercial GIS with built-in deep learning tools. It has tools for object detection and classification using CNNs (e.g. Detect Objects Using Deep Learning), and supports training models in integrated notebooks. ArcGIS integrates Python libraries (ArcPy, ArcGIS API for Python) for custom ML pipelines.
- GDAL/Rasterio: Core geospatial libraries for reading/writing raster data. They handle coordinate transforms, reprojection, and low-level I/O of GeoTIFF, JPEG2000, etc. Rasterio (Python) and GDAL (C/C++) are foundations for most custom GeoAI scripts.
- PyTorch and TensorFlow: Leading ML frameworks. They support GPU acceleration and have large ecosystems. PyTorch, for instance, is used in many research projects and has a rich library of vision models. TensorFlow also offers geospatial data handling via tfio (TensorFlow I/O) and ecosystem tools.
- TorchGeo: A PyTorch domain library for geospatial ML. It provides datasets, transforms, and pre-trained models specific to EO (e.g. VHR maps, multispectral data). TorchGeo simplifies loading large geospatial datasets and common workflows in PyTorch.
- RasterVision: A higher-level Python framework (by Azavea) for remote sensing ML. It orchestrates data preprocessing (tiling, augmenting), training (integrates with PyTorch/TensorFlow), and evaluation, specifically tailored for geospatial imagery.
- Other libraries: Tools like scikit-learn (ML library), Xarray (spatio-temporal data arrays), EarthPy (Raspberry Earth Engine tools), and SpaceNet/TorchVision/VIP (geospatial modeling) are also used. We note each has strengths: e.g. GDAL for data, QGIS/ArcGIS for geospatial workflows, PyTorch/TensorFlow/TorchGeo for modeling, and Google Earth Engine for large-scale data access.

Table 2: Open-source tools and libraries for GeoAI.

Tool / Library	Category	Description (and link)
Google Earth Engine	Cloud platform	Large-scale geospatial analysis (free sat. data); Python API for analysis and ML (e.g. image classification). earthengine.google.com
QGIS	GIS desktop	Open-source GIS with raster/vector support, Python scripting; integrates GDAL for remote sensing tasks. qgis.org

Tool / Library	Category	Description (and link)
ArcGIS Pro	GIS commercial	GIS platform with deep learning tools (object detection, segmentation), Python/API integration. esri.com
GDAL / Rasterio	Geospatial I/O	Library for reading/writing geospatial raster/vector formats; handles projections and metadata. gdal.org
PyTorch	ML framework	Deep learning library; supports CNN/RNN/Transformer; GPU-accelerated; large user base. pytorch.org
TensorFlow	ML framework	Deep learning library from Google; wide adoption; has TF Data and Keras APIs. tensorflow.org
TorchGeo	Geo ML library	PyTorch extension with geo-datasets, transforms, pre-trained models. github.com/microsoft/torchgeo
RasterVision	Geo ML framework	Orchestrates data pipeline for remote sensing (tiling, augmentation, training). rastervision.io
scikit-learn	ML library	Classic ML (RF, SVM, k-NN) for smaller problems or baseline models. scikit-learn.org
GDAL	Geospatial I/O	Core translator library for geospatial data formats (GeoTIFF, Shapefile).

Datasets

Key publicly-accessible datasets include:

- Landsat Archive: >50 years of 30m multispectral imagery, globally available via USGS EarthExplorer or Cloud. See USGS Landsat pages.
- Copernicus Sentinel: Sentinel-2 MSI (10m multispectral) and Sentinel-1 SAR (10m radar), free through ESA portals (Copernicus Open Hub) or via GEE.
- MODIS Products: e.g. MOD13A2 (16-day 1km NDVI) or MOD14 (fire hotspots). Accessible through NASA's LP DAAC.
- Global DEMs: Shuttle Radar Topography Mission (30m SRTM), ALOS World 3D (30m), Copernicus DEM (30m), used for elevation features.
- Climate and Weather: ERA5 reanalysis (0.25° hourly), WorldClim (bioclimatic layers), PRISM (gridded weather). These are often downloaded via APIs or packaged datasets.
- OpenStreetMap (OSM): Worldwide vector data (roads, buildings, land use). Although not imagery, OSM features are often fused with satellite data for GeoAI tasks.
- In-situ Networks: e.g. OpenAQ (global air quality measurements), NOAA Buoys (marine/ocean data), USGS stream gauges.
- Crowdsourced Datasets: e.g. iNaturalist for species observations, [TreeSnap] for urban tree data.
- Benchmark Remote Sensing Datasets: Popular ML benchmarks (SpaceNet for building footprints, DeepGlobe for land cover, IARPA functional Map). These provide labeled data to train/validate GeoAI models.
- Specialized Datasets: e.g. NASA's GEDI LiDAR data, NOAA weather radar archives, ESA CCI climate datasets, etc.

Each dataset usually has open-access links. For example, Landsat data can be accessed via GEE or USGS, Sentinel data via Copernicus. The table above and references provide starting points for obtaining these resources.

XIII. CHALLENGES AND CONSIDERATIONS

While GeoAI offers great promise, it faces significant challenges:

- Data Quality and Heterogeneity: Remote sensing data have noise (sensor artifacts, atmospheric effects) and missing observations (clouds in optical images). Lidar and SAR require careful calibration. In-situ sensors can have biases or outages. Diverse data modalities (multispectral, hyperspectral, LiDAR) complicate integration. Poor-quality labels (e.g. misclassified training pixels) can mislead models. Rigorous preprocessing and uncertainty filtering are needed.
- Scale and Computation: Processing satellite imagery (often terabytes or petabytes) demands heavy computation. Training DL models on global datasets requires GPUs or TPUs on cloud clusters, which can be costly. Real-time applications (e.g. disaster alerts) need near-edge inference or distributed computation (like GEE

parallelism). Managing scale also includes handling differing spatial resolutions; fusing 10m data with 1km data involves resampling and preserving detail.

- Model Bias and Transferability: AI models may overfit to training regions or sensors and perform poorly elsewhere. For example, a forest classifier trained in Europe may fail in tropical Africa. Uneven data coverage (more data from developed countries) can bias global models. Transfer learning mitigates this, but careful cross-validation (e.g. leaving-out regions) is essential.
- Interpretability and Trust: As Höhl et al. note, understanding "the inner workings of these models" is critical for EO tasks. Policymakers and stakeholders require explanations of AI-driven conclusions. For instance, if a model flags an area as flooded, visualizing which pixels and features drove that decision builds trust. Explainability is not just technical but also regulatory: the upcoming EU AI Act mandates transparency and accountability in AI systems. GeoAI must align with such frameworks (safety, privacy, transparency) to be socially responsible.
- Privacy and Legal Issues: Very-high-resolution imagery can inadvertently capture private activities (e.g. cars in driveways). Similarly, geotagged data from mobile devices raise privacy concerns. GeoAI systems must respect legal constraints on data use (some satellite companies restrict use of VHR data; GDPR rules apply to personal data). Federated approaches and anonymization help address privacy, while clear governance of data licenses is needed.
- Environmental and Ethical Biases: GeoAI models may embed biases (e.g. focusing on urban, neglecting rural). Ecological data are often sparse, risking underrepresentation of certain ecosystems. Ethically, automated monitoring could have negative consequences (e.g. surveillance concerns in protected areas). Community engagement is recommended to ensure technology serves environmental justice.
- Validation and Ground Truth: Many environmental variables (biodiversity, ecosystem health) lack precise ground truth. Validating models with sparse field data is challenging. Benchmark datasets and standardized metrics are fewer in GeoAI than in other AI fields. Collaborative programs (e.g. GEO's Africa initiative) aim to build evaluation datasets.

Despite these challenges, the literature shows strategies to mitigate them: data-driven calibration (e.g. removing outliers in coastal mapping), ensemble modelling for robustness, and continuous learning (models updated with new data). The increasing focus on XAI, federated methods, and international standards reflects the community's response to ethical and regulatory pressures.

XIV. BEST PRACTICES AND FUTURE DIRECTIONS

Based on the surveyed literature, we recommend:

- Robust Preprocessing: Use state-of-the-art correction and masking (e.g. FMask for clouds, BRDF normalization) to improve input quality. Align and resample disparate datasets carefully (common CRS). Compute informative indices (NDVI, NDWI, NBR, etc.) as auxiliary features.
- Data Augmentation and Regularization: Given limited labeled data, aggressively augment training images. Consider synthetic data (e.g. GAN-generated patches for rare classes). Regularize models to prevent overfitting to one region/sensor.
- Transfer and Multi-Task Learning: Pretrain models on large generic geospatial datasets and fine-tune on local tasks. Use multi-task learning (e.g. concurrently predict land cover and elevation) to leverage shared features and reduce labeled data needs.
- Active and Federated Learning: Engage domain experts in active learning loops to improve labels. For collaborative projects (e.g. global climate monitoring), employ federated learning so that countries can jointly train models without data sharing.
- Explainability and Uncertainty: Incorporate XAI methods to interpret model decisions, especially for high-stakes outputs (e.g. evacuation alerts). Report predictive uncertainties using ensembles or Bayesian neural nets so users know confidence levels.
- Edge-Cloud Hybrid Deployment: Strategically balance edge and cloud: use edge (drones, IoT) for

rapid local tasks, and cloud (GEE, AWS) for large-scale analysis. Containerize models (Docker) and use APIs for interoperability.

- Open Data and Code: Embrace digital public goods – publish models, code, and datasets with clear licenses. Platforms like GEE, OpenAerialMap, and community notebooks (e.g. Binder) foster reproducibility. As one GEO author notes, open satellite data and cloud tools "expand access and participation" in GeoAI.
- Interdisciplinary Collaboration: GeoAI lies at the intersection of remote sensing, ecology, and AI. Future progress depends on collaboration: integrating indigenous and local knowledge, aligning with policy needs, and building user-centric tools. For instance, community mapping combined with AI analysis was shown to help tailor tech to local climate needs.

Future Research

The field is moving toward:

- Foundation Models for Earth Science: Large self-supervised models trained on massive geospatial corpora (images, sensor streams) that can be fine-tuned for various tasks. Like GPT or BERT but for remote sensing data, enabling few-shot learning on new problems.
- Multimodal GeoAI: Integrating text (scientific reports, social media posts), imagery, and sensor data. For example, combining satellite change detection with news/social media to verify disaster impacts.
- Physics-Informed AI: Embedding physical constraints (e.g. conservation laws, topography) into models, so that predictions (e.g. air pollution spread) respect known processes.
- Real-Time and Edge Intelligence: Deploying lightweight GeoAI on satellites or unmanned aerial vehicles for on-board analysis and downlinking only alerts or summaries.
- Global Digital Twins: Developing comprehensive digital twins of the Earth that integrate GeoAI predictions with climate and ecosystem models for real-time simulation and planning.

In conclusion, GeoAI is a rapidly evolving field that brings cutting-edge AI to bear on environmental challenges. Recent advances (summarized in this paper) demonstrate its potential, while ongoing research and community efforts address the remaining gaps. With continued innovation and interdisciplinary collaboration, GeoAI will become an indispensable component of environmental monitoring and action.

REFERENCES

[1] M. A. M. Aman *et al.*, "A review of remote sensing-based crop yield estimation: Machine learning techniques and environmental, algorithmic, and hardware limitations," *Frontiers in Plant Science*, vol. 17, Art. no. 1742689, 2026, doi: 10.3389/fpls.2026.1742689.

[2] E. Bainomugisha *et al.*, "AI-driven environmental sensor networks and digital platforms for urban air pollution monitoring and modelling," *Societal Impacts*, vol. 3, Art. no. 100044, 2024, doi: 10.1016/j.socimp.2024.100044.

[3] R. Fleckenstein *et al.*, "A framework for integrating spatiotemporal deep learning methods with Landsat for annual land cover and impervious surface mapping," *Remote Sensing of Environment*, vol. 303, Art. no. 115347, 2026, doi: 10.1016/j.rse.2026.115347.

[4] Höhl *et al.*, "Opening the black-box: A systematic review on explainable AI in remote sensing," *arXiv preprint arXiv:2402.13791*, 2024, doi: 10.48550/arXiv.2402.13791.

[5] L. Kandasamy *et al.*, "Enhanced remote sensing and deep learning aided water quality detection in the Ganges River, India supporting monitoring of aquatic environments," *Results in Engineering*, vol. 25, Art. no. 103604, 2025, doi: 10.1016/j.rineng.2024.103604.

[6] M. Koch, "World Day for Glaciers: How GeoAI and satellite data are transforming cryosphere monitoring in climate extreme environments," *Group on Earth Observations (GEO) Blog*, Mar. 20, 2026. [Online]. Available: https://earthobservations.org/about-us/news/world-day-for-glaciers-how-geoai-and-satellite-data-are-transforming-cryosphere-monitoring-in-climate-extreme-environments

[7] Y. Ma, S. Chen, S. Ermon, and D. B. Lobell, "Transfer learning in environmental remote sensing," *Remote Sensing of Environment*, vol.

301, Art. no. 113924, 2024, doi: 10.1016/j.rse.2023.113924.

[8] Md Jelas, M. A. M. A. Zulkifley, and M. A. M. Abdullah, "Deforestation detection using deep learning-based semantic segmentation techniques: A systematic review," *Frontiers in Forests and Global Change*, vol. 7, Art. no. 1300060, 2024, doi: 10.3389/ffgc.2024.1300060.

[9] S. Miao *et al.*, "Utilizing active learning and attention-CNN to classify vegetation based on UAV multispectral data," *Scientific Reports*, vol. 14, Art. no. 31061, 2024, doi: 10.1038/s41598-024-82248-3.

[10] S. Moreno-Álvarez *et al.*, "Federated learning meets remote sensing," *Expert Systems with Applications*, vol. 255, Art. no. 124583, 2024, doi: 10.1016/j.eswa.2024.124583.

[11] NASA/JPL, "Welcome to ECOSTRESS," 2026. [Online]. Available: https://ecostress.jpl.nasa.gov

[12] Y. Song, F. Biljecki, G. Camps-Valls, and P. Atkinson, "GeoAI: Beyond mapping earth and cities through explainability, adaptability, and sustainability," *iScience*, vol. 29, no. 1, Art. no. 114407, 2026, doi: 10.1016/j.isci.2025.114407.

[13] U.S. Geological Survey, "EarthExplorer," [Online]. Available: https://earthexplorer.usgs.gov/

[14] European Space Agency, "Copernicus Open Access Hub (SciHub)," [Online]. Available: https://scihub.copernicus.eu/

[15] NASA, "MODIS (Moderate Resolution Imaging Spectroradiometer)," [Online]. Available: https://modis.gsfc.nasa.gov/

[16] Google, "Google Earth Engine," [Online]. Available: https://earthengine.googleQGIS Development Team, "QGIS Geographic Information System," [Online]. Available: https://qgis.org/

[17] [Esri, "ArcGIS Pro," [Online]. Available: https://www.esri.com/en-us/arcgis/products/arcgis-pro/overview

[18] GDAL/OGR contributors, "GDAL," and S. Gillies *et al.*, "Rasterio," [Online]. Available: https://gdal.org/, https://rasterio.readthedocs.io/

[19] PyTorch Contributors, "PyTorch," [Online]. Available: https://pytorch.org/

[20] TensorFlow Team, "TensorFlow," [Online]. Available: https://tensorflow.org/

[21] Microsoft, "TorchGeo," [Online]. Available: https://github.com/microsoft/torchgeo/

[22] Azavea, "Raster Vision," [Online]. Available: https://github.com/azavea/raster-vision