The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XLVIII-4/W20-2025
Free and Open Source Software for Geospatial (FOSS4G) 2025 – Academic Track, 17–23 November 2025, Auckland, New Zealand

A Scalable Open-Source System for Impervious Land Mapping Using GRASS and the Python Ecosystem

Tomaž Žagar, Alen Mangafić

Geodetic Institute of Slovenia, tomaz.zagar@gis.si, alen.mangafic@gis.si

Keywords: Sentinel-2 time series, aerial imagery, GRASS, TorchGeo, EuroSAT, Random Forest.

Abstract

This paper presents a scalable open-source system for national-level impervious-surface mapping and monitoring that combines GRASS, HDF5 and Python-based machine-learning libraries. Continuous Sentinel-2 multispectral monitoring is paired with targeted very high resolution (VHR) orthophoto segmentation to efficiently detect and delineate land transformation. Sentinel-2 imagery is retrieved from the Microsoft Planetary Computer, organized into a GRASS space-time raster dataset and classified using TorchGeo models fine-tuned on the EuroSAT dataset. Persistent transitions toward impervious-related classes identify disturbance candidates, and these areas trigger semantic segmentation of corresponding orthophoto tiles using GPU-accelerated Random Forests in RAPIDS cuML. The resulting outputs are vectorized, enriched with registry data and disseminated through spatial database services.

Applied to Slovenia for the 2020–2025 period, the system detected 99 km² of new impervious surfaces across a 20,271 km² study area, corresponding to 0.5 % land transformation (0.1 % annually). These results demonstrate that integrating continuous multispectral time-series analysis with event-driven VHR segmentation provides an efficient, reproducible and high-detail approach for operational land-change detection and environmental assessment.

1. Introduction

Impervious-surface expansion is a persistent driver of environmental degradation, underscoring the need for systems that can observe land transformation continuously, process data efficiently, and deliver interpretable outputs at multiple spatial scales. Modern Earth-observation workflows must therefore unite two historically separate domains: scalable raster-storage formats designed for high-performance machine learning, and spatial data engines capable of precise geospatial processing.

Advances in array-based storage such as HDF5 and Zarr (Ambatipudi and Byna, 2022) have become central to large-scale machine-learning pipelines, offering chunking, compression, and lazy loading that allow models to ingest terabytes of multispectral or VHR imagery without exceeding system memory. At the same time, geospatial platforms such as GRASS provide mature spatial-temporal data models, topological operations, and algorithms essential for deriving meaningful land-change signals. For operational monitoring, both worlds must be accessible directly from Python, enabling seamless transitions between spatial analysis and GPU-accelerated learning.

Bridging these two infrastructures creates a unified workflow in which spatially explicit datasets are prepared, stored, processed, and analysed without format conversions or computational bottlenecks. This interoperability allows Sentinel-2 time series, orthophotos, labels, and machine-learning features to coexist in scalable formats while remaining tightly coupled with spatial algorithms. As a result, complex tasks—ranging from pixel-level classification to semantic segmentation and temporal change detection—can be executed reproducibly and at national scale.

This paper presents such a system: an open-source pipeline that blends continuous Sentinel-2 monitoring, event-driven orthophoto segmentation, GRASS-based spatio-temporal processing, and Python machine-learning frameworks to produce detailed, operational impervious-surface maps.

2. Materials and Methods

The system integrates scalable raster storage, spatial data management and Python-based machine learning into a unified environment for impervious-surface detection. GRASS is used for organizing and processing spatial and temporal datasets, while high-volume imagery and derived features are stored in chunked HDF5 arrays that support lazy loading and efficient retrieval during model training and inference. This connection between GRASS and HDF5 ensures that spatial algorithms and GPU-accelerated machine learning operate on the same data without intermediate conversion.

Figure 1 presents the diagram of the system. It summarizes the interaction between Sentinel-2 time-series processing, spatial data management, machine-learning components and high-resolution impervious mapping. The subsequent steps shown in Figure 1 include change detection in the time series, which highlights potential disturbance areas and triggers targeted semantic segmentation of the corresponding orthophoto tiles. The segmented outputs are postprocessed, compared with relevant registry data, and after manual inspection prepared for dissemination. Before dissemination, all outputs undergo an additional manual review step to ensure that geometries, class labels and detected change periods remain consistent with visual interpretation and registry datasets.

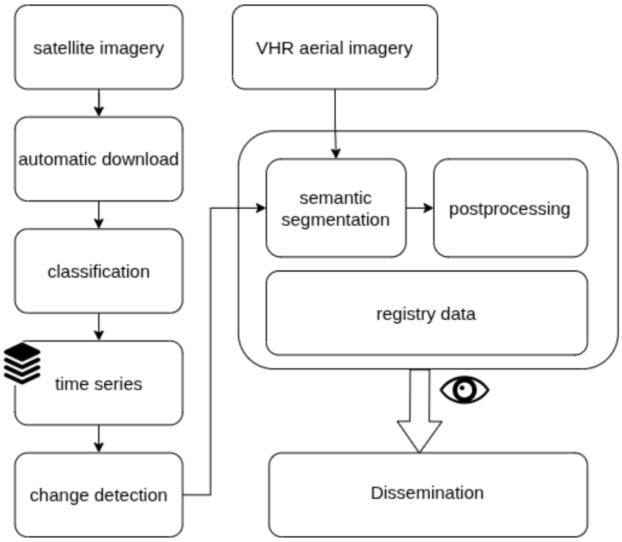

Figure 1. Overview of the impervious-surface mapping and monitoring system.

2.1 Data Sources and Management

The system ingests multiple raster inputs. These include Sentinel-2 multispectral bands from the Microsoft Planetary Computer, which provide continuous monitoring, and very high resolution orthophotos from the Mapping Authority of the Republic of Slovenia, containing near infrared, red, and green channels at 50 cm resolution. An example of these complementary inputs for the same area and acquisition date is shown in Figure 2. Orthophotos are used in regions where the Sentinel-2 time series indicates change.

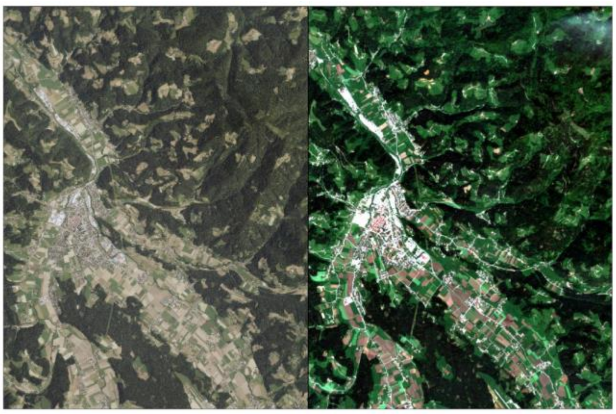

Figure 2. Orthophoto (left) and Sentinel-2 image (right) - Copernicus Sentinel data (2022) - of the same area acquired on the same day.

The system also incorporates reference impervious layers from earlier periods to generate labeled training data and establish baselines for change detection, as well as ancillary datasets such as agricultural and forest land boundaries to support comprehensive environmental assessments.

Sentinel-2 imagery is accessed through the Microsoft Planetary Computer (Microsoft Open Source et al., 2022), which provides a STAC-compliant data structure and consistent metadata, enabling automated retrieval across the full monitoring period from 1 January 2020 to 31 October 2025. Figure 3 shows the Sentinel-2 tiles covering Slovenia, illustrating the spatial subdivision used for organizing and downloading the imagery. We downloaded imagery with less than 5 % of clouds.

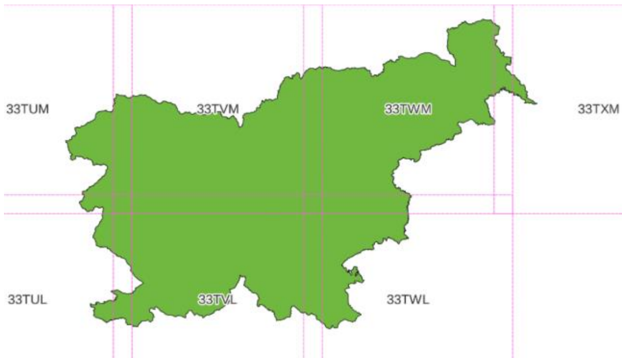

Figure 3. Sentinel-2 MGRS tiles covering Slovenia.

Aerial imagery for Slovenia is obtained from the national orthophoto programme (Surveying and Mapping Authority of the Republic of Slovenia, 2025), produced within the cyclical aerial surveying of Slovenia, in which only part of the country is photographed each year and full national coverage is achieved on a three-year cycle. The orthophotos used in this study include near-infrared, red, and green (NIR–R–G) channels and cover acquisition years 2020, 2021, 2022, 2023, and 2024, capturing sensor, regional, and seasonal differences relevant for robust impervious-surface mapping. Figure 4 shows the national aerial imagery used in the workflow.

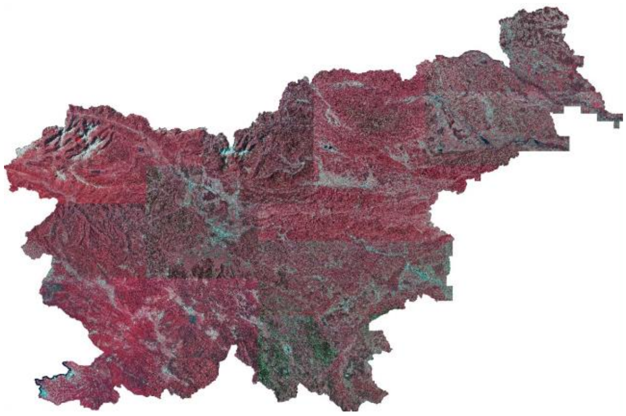

Figure 4. Aerial NIR orthophoto of Slovenia used in this study.

Orthophotos are preprocessed in GRASS and exported to HDF5. Image features and labels are stored as chunked arrays to enable efficient parallel processing and patch based classification. After classification, prediction maps are imported into GRASS as a space-time raster dataset. Results of time-series analyses are converted to vector layers, and loaded into a PostgreSQL/PostGIS database for downstream operations and integration.

2.2 Machine Learning Pipeline

2.2.1 Classification of Sentinel-2

Sentinel-2 land-cover classification is performed using TorchGeo, a geospatial deep-learning library developed as an extension of PyTorch. TorchGeo provides native support for spatially aware data structures, coordinate systems, georeferenced patch extraction and multi-band imagery, allowing models to operate directly on Earth-observation datasets without manual preprocessing steps typical of generic computer-vision frameworks. These capabilities make TorchGeo more suitable for remote sensing than pure PyTorch, especially when handling large multispectral archives, STAC indexing and heterogeneous footprints.

TorchGeo models rely on convolutional neural-network backbones, which function as feature extractors that learn hierarchical spatial and spectral patterns. A backbone is a pretrained architecture, such as ResNet, DenseNet or a Vision Transformer, initially trained on large image corpora to capture generalizable representations. Although multiple backbones are available, this system uses ResNet because it offers a strong balance between computational speed, classification accuracy and stability across diverse land-cover scenes.

Instead of training a model from scratch, the system initializes the ResNet backbone with pretrained weights from the SATLAS MSI dataset, a large collection of multispectral Sentinel-2 patches curated by Microsoft for remote-sensing model initialization. SATLAS MSI covers highly diverse geographical regions, atmospheric conditions and land-cover types, enabling the backbone to learn spectral–spatial abstractions closely aligned with Sentinel-2 data. This improves convergence, stabilizes early training phases and enhances generalization in regions with limited labeled samples.

Domain adaptation is necessary because EuroSAT and Sentinel-2 Level-2A products are not identical in format, radiometric scaling or metadata structure. To harmonize the inputs, each Sentinel-2 patch is standardized using z-score normalization. This method is preferred over range-based normalization techniques such as min–max scaling because it is independent of absolute reflectance minima and maxima, remains stable under illumination and seasonal variability, and avoids compressing low-variance spectral bands, which often contain informative content. Z-score normalization also matches the statistical assumptions embedded in convolutional feature distributions and thus benefits fine-tuning stability.

During training, the system applies data augmentation consisting of horizontal and vertical flips and 90-degree rotations. These transformations expand the effective training distribution, reduce overfitting and increase resilience to orientation-dependent land-use patterns, building geometries and acquisition variability. After fine-tuning, the model predicts ten land-cover classes and reaches 97% accuracy.

2.2.2 Semantic Segmentation of Orthophotos

High-resolution impervious mapping is performed using GPU-accelerated Random Forest models implemented in RAPIDS cuML. RAPIDS cuML is a machine-learning library that mirrors the scikit-learn API (Pedregosa et al., 2011), but executes computations on NVIDIA GPUs, providing substantial speedups for ensemble methods, nearest-neighbour searches and decision-tree operations. Its design philosophy is to offer drop-in replacements for classical algorithms while exploiting highly parallel GPU architectures. This results in significantly reduced training and inference times when compared with CPU-based implementations.

Random Forests in cuML are particularly suitable for orthophoto segmentation because decision-tree construction, feature-split evaluation and prediction over millions of pixels can be distributed efficiently across many CUDA cores. The similarity of the cuML interface to scikit-learn allows reproducible model configuration, while the GPU backend supports national-scale image segmentation without CPU bottlenecks. Training polygons representing bare soil and impervious surfaces, low vegetation and medium–high vegetation are exported from GRASS to HDF5 and assembled into tiles for GPU processing. Model development uses an 80/20 train–test split with three-fold stratified cross-validation.

2.2.3 Temporal Analysis

Classified Sentinel-2 mosaics are organized as a temporal stack within the GRASS space-time raster dataset, enabling aggregation and change detection. Monthly composites and six-month aggregates are produced following uniform STAC metadata conventions. Transitions into impervious-related classes are identified and validated using a persistence criterion that requires three consecutive temporal periods. These confirmed transitions guide the selection of orthophoto tiles for high-resolution segmentation.

Change detection is performed by comparing Sentinel-2 based land-cover prediction maps across years and by maintaining a Sentinel-2 time series in a GRASS space-time raster dataset. Classification maps are mosaicked by satellite (A, B, C) per date, generating 1532 processed scenes that consolidate into 421 unique dates, which are further aggregated into 69 monthly composites. Temporal summaries are constructed as time series of six-month aggregates, each of which is subsequently reclassified.

The TorchGeo model, fine-tuned on the EuroSAT dataset, provides land cover classes, and the temporal analysis focuses on transitions into highway, industrial and residential classes from other land cover types as indicators of areas with new impervious surfaces. In the reclassified time series, these three impervious-related classes are assigned a value of 1 (potential impervious), while all others—annual crop, forest, herbaceous vegetation, pasture, permanent crop, river and sea/lake—are assigned a value of 0. In practice, individual patches often contain mixtures of land-cover types, and each detected transition triggers a new classification step. To ensure reliability, classification stability is required for three consecutive temporal periods. This approach yields a conservative estimate, providing a credible lower bound for urban growth.

New Sentinel-2 scenes are downloaded automatically after publication, and GRASS spatio-temporal aggregation and change-detection tools are used to derive spatial differences that identify newly developed impervious areas. When new very high resolution orthophotos are available, the system automatically triggers segmentation only over tiles flagged by the Sentinel-2 disturbance signal.

2.3 Postprocessing and Dissemination

Postclassification filtering reduces noise. Cleaned rasters are vectorized and simplified within GRASS. Final geometries are stored in PostGIS, generalized, enriched with auxiliary land cover data, and disseminated as WFS services, enabling spatial queries, external integration, and comparison with land-use data for assessing impacts on soil and potential for restoration. Postprocessing also uses land-use data, and the resulting layers are manually verified before dissemination. The verified outputs are overlaid with the registry datasets of the land cadastre, planned land use and spatial planning acts of the Republic of Slovenia. Attributes assigned to the classified objects are the date of the VHR imagery, the period in which the change occurred, the land-cover transition detail (e.g. Forest to Residential), the area, and the distance from polygons of these registry datasets.

3. Results and Discussion

The system was applied to multi-year orthophoto and Sentinel-2 datasets. Continuous Sentinel-2 monitoring produced timely impervious updates. The HDF5-based chunking significantly reduces I/O bottlenecks when performing machine learning tasks. Targeted segmentation of orthophotos over Sentinel-2 disturbance areas improved efficiency while preserving detail. Change maps accurately identified new impervious areas. Vectorized outputs supported further spatial analysis, visualization, and comparative evaluation against agricultural and forested land-use data. WFS dissemination enabled direct use in client applications.

The analysis covered a study area of 20,271 km². Across the five-year period, we detected 99 km² of impervious-surface disturbances, corresponding to 0.5 % total land transformation (0.1 % annually). This annual growth rate lies within the typical 0.1–0.3 % urban-growth range reported for EU (Heuser, 2022).

Figure 5 illustrates an example of the temporal disturbance detection on Sentinel-2 data for a selected location. The two infrared images show the Sentinel-2 patch before and after the detected change, followed by an OpenStreetMap (OSM) image of the same area, where the new impervious features are not yet mapped. This example demonstrates also the potential use of the disturbance signal as support for OSM contributors.

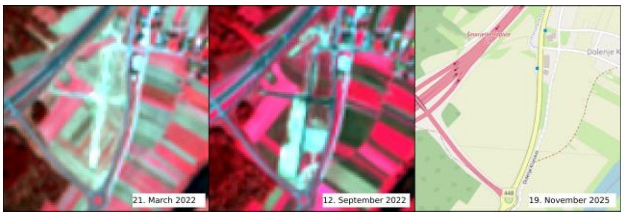

Figure 5. Sentinel-2 patch before and after change, with corresponding OSM view. Copernicus Sentinel data (2022).

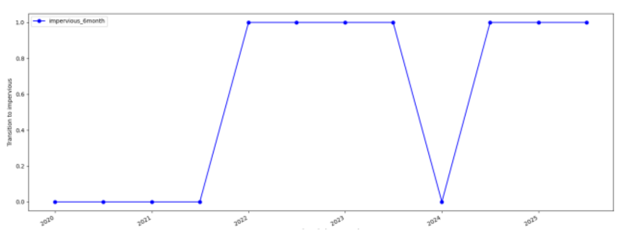

Figure 6. Land-cover transitions for the example shown in Figure 5, derived from the reclassified Sentinel-2 time-series aggregates.

The land-cover transitions for the same example are shown in Figure 6, where the reclassified time-series output highlights the detected change. In this example, the transition to impervious was confirmed because it persisted for at least three consecutive temporal periods, consistent with our conservative stability requirement. As visible in the time-series plot, a short-term transition back to a non-impervious class occurs later, but because it does not last for three consecutive periods it is not accepted as a valid change. The same rule applies to all subsequent detections—any newly observed transition must again persist for at least three consecutive periods before being classified as a confirmed change.

Figure 7 presents the corresponding orthophoto images for the same patch, with the left image showing the state before the change and the right image the state after the change.

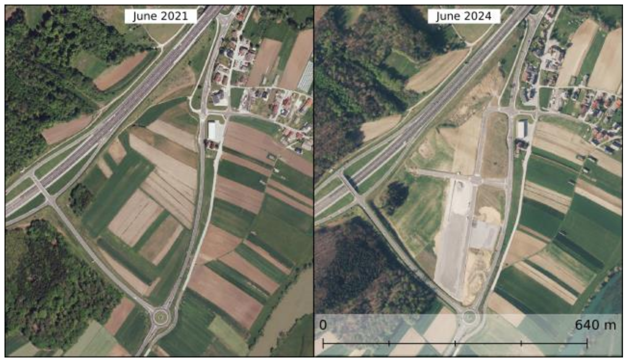

Figure 7. Orthophotos before (left) and after (right) the detected change.

Figure 8 shows the semantic-segmentation result for the 2024 orthophoto, and Figure 9 displays a selected detail where the resulting classes are clearly visible. Bare soil and impervious surfaces appear in dark tones, low vegetation in green, and medium-high vegetation in yellow. The unsegmented orthophoto imagery was acquired prior to 2024 and was used as the training dataset.

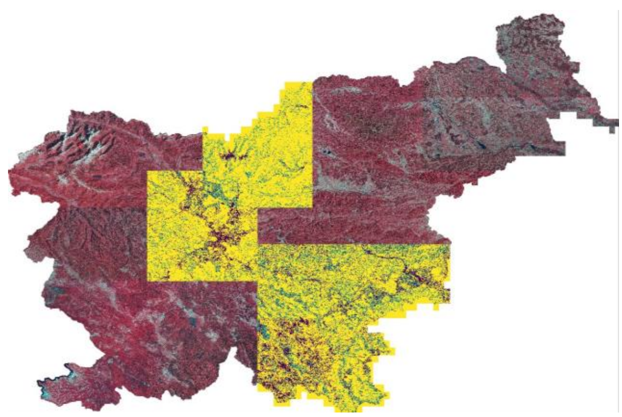

Figure 8. Result of the semantic segmentation of the 2024 orthophoto

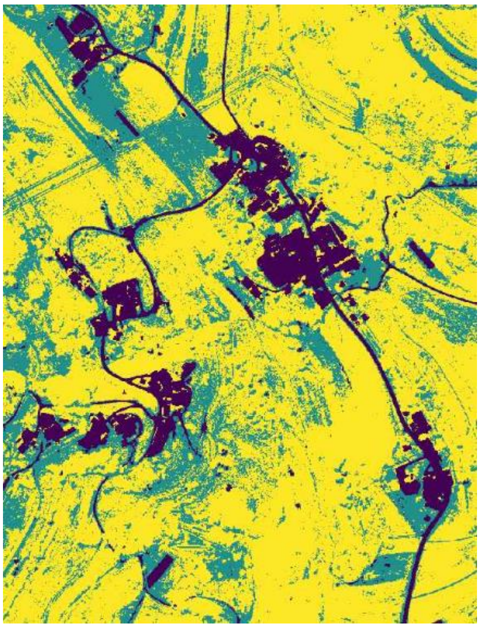

Figure 9. A selected detail of the semantically segmented orthophoto shown in Figure 8.

A second example of detected change is shown in Figure 10, together with instance-segmented polygons produced during orthophoto segmentation. The polygons closely follow the newly developed buildings and bare soil areas, showing that the segmented objects correspond well to the visible changes in the orthophoto.

The same example as the one shown in Figure 10 is presented also in Figure 11 demonstrating the use of the workflow's output after segmentation and overlaying the registry data as a WFS in QGIS.

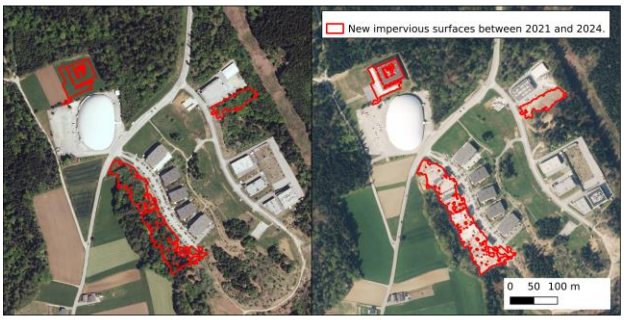

Figure 10. Polygons of new detected impervious surfaces.

One of the polygons over a newly developed impervious object is selected and its attributes are shown, demonstrating that the result of the workflow is an instance-segmented layer enriched with registry data attributes.

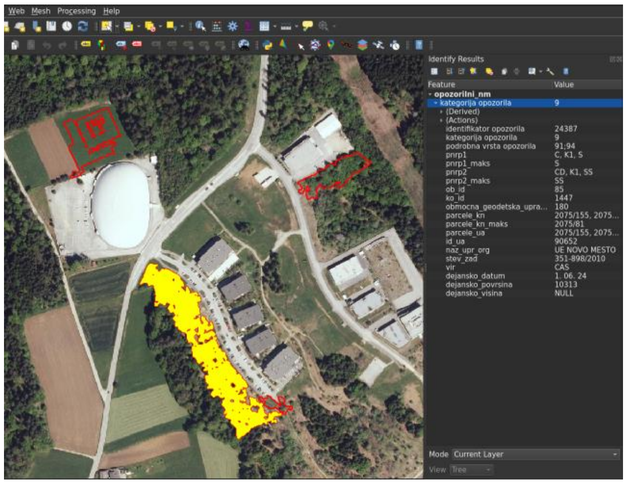

Figure 11. Selected segmented polygon with registry attributes in QGIS.

4. Conclusion

This study demonstrates a reproducible, open-source pipeline for high-resolution impervious land mapping with continuous Sentinel-2 monitoring. Integration of GRASS, HDF5, RAPIDS cuML and TorchGeo enables efficient classification, temporal analysis, event-driven segmentation of orthophotos, spatial output management, and environmental impact assessment. The architecture is adaptable to diverse data sources and modeling approaches, thus suitable for research, operational applications, and informed land management and restoration practices, with outputs shared through WFS services.

Future work will focus on extending the workflow in several directions. Semantic and instance segmentation of Sentinel-2 imagery, together with time series of such maps, will enable more detailed and concurrent analysis. Improvements to the aggregation of space-time datasets are planned, particularly the inclusion of meteorological information, for example: snow conditions that introduce noise, and per-patch cloud and noise removal, since the current approach discards entire scenes when cloud cover exceeds 5 %. Further developments will also explore semantic segmentation of VHR imagery with deep learning to support a larger set of classes, as well as sensor-agnostic models that combine satellite and VHR imagery and potentially additional sources.

Acknowledgements

This work was done within the scope of the targeted research project V2-24072: Development of a unified spatial physical change detection system with artificial intelligence,
financially supported by the Slovenian Research and Innovation Agency and the Surveying and Mapping Authority of the Republic of Slovenia.

References

Ambatipudi, S., Byna, S., 2022: A comparison of HDF5, Zarr, and netCDF4 in performing common I/O operations.

Breiman, L., 2001: Random forests. Mach. Learn., 45, 5–32. doi.org/10.1023/A:1010933404324.

Helber, P., Bischke, B., Dengel, A., Borth, D., 2017: EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification. IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens., 12, 2217–2226. doi.org/10.1109/JSTARS.2019.2918242.

Heuser, D.I., 2022: Soil governance in current European Union law and in the European Green Deal. Soil Secur., 6, 100053. doi.org/10.1016/j.soisec.2022.100053.

McFarland, M., Emanuele, R., Morris, D., Augspurger, T., 2022: Microsoft Planetary Computer. doi.org/10.5281/zenodo.7261897.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Dubourg, V., Passos, A., Brucher, M., Perrot, M., Duchesnay, E., 2011: Scikit-learn: Machine learning in Python. J. Mach. Learn. Res., 12, 2825-2830.

Raschka, S., Patterson, J., Nolet, C., 2020: Machine learning in Python: Main developments and technology trends in data science, machine learning, and artificial intelligence. Information, 11, 193. doi.org/10.3390/info11040193.

Stewart, A.J., Robinson, C., Corley, I.A., Ortiz, A., Ferres, J.M.L., Banerjee, A., 2022: TorchGeo: Deep learning with geospatial data. GIS: Proceedings of the ACM International Symposium on Advances in Geographic Information Systems. doi.org/10.1145/3557915.3560953.

Surveying and Mapping Authority of the Republic of Slovenia, 2025: Daljinsko zaznavanje – E-prostor. https://www.e-prostor.gov.si/podrocja/drzavni-topografski-sistem/daljinsko-zaznavanje/