

---

# Artificial intelligence to advance Earth observation: a perspective

Devis Tuia<sup>1,\*</sup>, Konrad Schindler<sup>2</sup>, Begüm Demir<sup>3,4</sup>, Xiao Xiang Zhu<sup>5</sup>, Mrinalini Kochupillai<sup>5</sup>, Sašo Džeroski<sup>6</sup>, Jan N. van Rijn<sup>7</sup>, Holger H. Hoos<sup>7,13</sup>, Fabio Del Frate<sup>8</sup>, Mihai Datcu<sup>9,10</sup>, Volker Markl<sup>3,4</sup>, Bertrand Le Saux<sup>11</sup>, Rochelle Schneider<sup>11</sup>, Gustau Camps-Valls<sup>12</sup>

<sup>1</sup> Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

<sup>2</sup> ETH Zurich, Switzerland

<sup>3</sup> Technische Universität Berlin, Germany

<sup>4</sup> BIFOLD - Berlin Institute for the Foundations of Learning and Data, Germany

<sup>5</sup> Technische Universität München, Germany

<sup>6</sup> Jozef Stefan Institute, Ljubljana, Slovenia

<sup>7</sup> Leiden University, the Netherlands

<sup>13</sup> RWTH Aachen University, Germany

<sup>8</sup> University of Rome “Tor Vergata”, Italy

<sup>9</sup> German Aerospace Center (DLR), Germany

<sup>10</sup> University Politehnica of Bucharest (UPB), Romania

<sup>11</sup> European Space Agency (ESA) Φ-lab, Italy

<sup>12</sup> Universitat de València, Spain

\* Corresponding author: [devis.tuia@epfl.ch](mailto:devis.tuia@epfl.ch).

This article is a preprint of a paper accepted at IEEE Geoscience and Remote Sensing Magazine (DOI: 10.1109/MGRS.2024.3425961). Official published version available at: <https://ieeexplore.ieee.org/document/10669817>.

---

# Contents

<b>Introduction</b>	<b>4</b>
<b>I Modelling — Machine learning, computer vision and processing</b>	<b>7</b>
<b>1 Machine learning and computer vision for Earth observation</b>	<b>7</b>
1.1 Specific challenges of Earth observation . . . . .	7
1.2 General problem settings relevant for Earth observation . . . . .	8
1.3 AutoML for Earth Observation . . . . .	10
1.4 Beyond computer vision . . . . .	11
<b>2 Advanced processing and computing</b>	<b>12</b>
2.1 Earth observation ecosystems – new system architectures . . . . .	12
2.2 Federated learning . . . . .	13
2.3 AI interoperability . . . . .	14
2.4 On-board data processing . . . . .	14
<b>II Understanding — Physics-machine learning interplay, causality and ontologies</b>	<b>16</b>
<b>3 Knowledge-based AI and Earth observation</b>	<b>16</b>
3.1 Decision modelling and decision support with Earth observation . . . . .	16
3.2 Knowledge-based machine learning and Earth observation . . . . .	16
3.3 Ontologies for open science in machine learning and Earth observation . . . . .	17
<b>4 Explainable AI and causal inference</b>	<b>18</b>
4.1 Explainable AI . . . . .	18
4.2 Causal inference . . . . .	19
<b>5 Physics-aware machine learning</b>	<b>20</b>
5.1 Inverting physical models with machine learning . . . . .	21
5.2 Hybrid machine learning: interactions of transfer models and ML . . . . .	21
5.3 Discovering equations from data . . . . .	22
<b>III Communicating — Machine-user interaction, trustworthiness &amp; ethics</b>	<b>24</b>
<b>6 User-centric Earth observation</b>	<b>24</b>
6.1 Decreasing the user load with active learning . . . . .	24

---

6.2	Language-based interaction . . . . .	25
6.3	Enhanced visualisation . . . . .	26
<b>7</b>	<b>Earth observation and society: the growing relevance of ethics</b>	<b>28</b>
7.1	AI and Data Ethics . . . . .	28
7.1.1	AI Ethics . . . . .	28
7.1.2	Data Ethics . . . . .	30
7.2	Ethical Opportunities for Sustainability: Environment and Health . . . . .	31
	<b>Conclusions</b>	<b>31</b>

---

## Abstract

Earth observation is increasingly used for mapping and monitoring processes occurring at the surface of Earth. Data acquired by satellites nowadays allow us to have a global view, consistent in time, of the state of our forests, oceans, and growing urban areas. However, such a wealth of data has little value without appropriate processing chains able to convert the pixel values to information useful for decision makers.

Recently, machine learning has seen fast advances – especially thanks to the rise of deep learning methodologies – and is increasingly deployed in Earth observation image processing systems. The ever-growing models from computer vision and natural language processing have inspired developments in remote sensing, and new approaches are constantly proposed in the field. However, despite their impressive results, the ever-growing mass of approaches and solutions makes it complicated to have a holistic overview and to know the most promising approaches from the field.

In this paper, we aim to fill this knowledge gap and propose to review the thriving ecosystem focusing on developing AI models for Earth observation, its recent trends, and sketch potential pathways for future advances.

## Introduction

Earth observation (EO) is a prime instrument for monitoring land and ocean processes, studying the dynamics at work, and taking the pulse of our planet. A large variety of sensor data (active / passive / of many resolutions) are nowadays accessible to researchers, agencies, and the general public. However, a final barrier remains the need for technology to convert the enormous quantities of raw EO data generated into the valuable information necessary for making decisions and taking concrete action. The ability to extract meaningful information from raw EO data is an essential prerequisite to achieving significant impact, be it in terms of monitoring or documenting (e.g. progress towards the United Nations' sustainable development goals), predicting and issuing timely warnings (e.g. related to future natural disasters [1] and need for emergency evacuations), or projecting the effects of human actions and natural processes on nature and society [2–4].

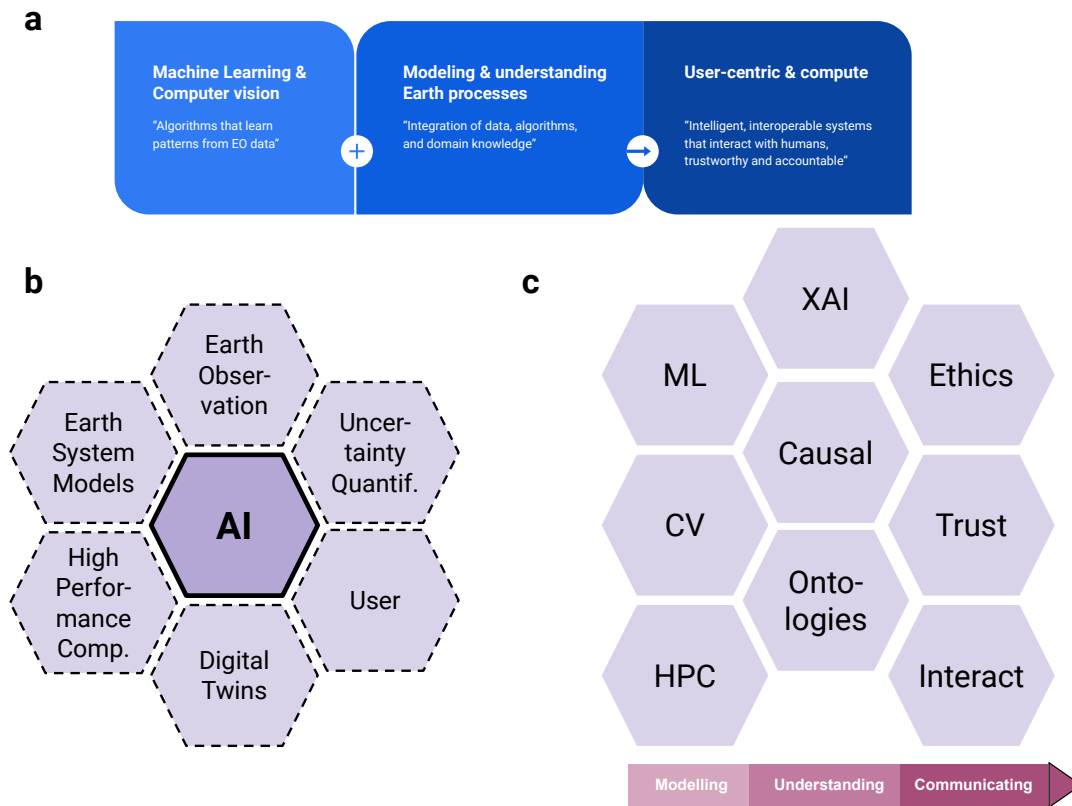
From the beginning of EO research, computer science and signal processing techniques have been important for extracting meaningful information from raw data. We have observed rapid technological advances in these fields over the last few years, mainly owing to the increasing use of artificial intelligence (AI) methods – and in particular, techniques from machine learning (ML) in the context of Earth observation. Notably, the use of deep learning (DL) and other ML techniques has had a broad and significant impact on EO and remote sensing [5, 6], where they are used across entire processing chains, from data compression and transmission to image recognition and predictions of environmental variables (land cover, land use, biomass, etc.). Similarly, ML techniques are increasingly widely used in many areas of environmental science [7], with potential key roles to play to support decision making in areas such as food security, environmental health or biodiversity monitoring.

However, while classical processing tasks, such as land cover classification, seem to have reached a

---

certain degree of maturity, partially due to the better integration of EO and ML methods, substantial work remains to render these approaches truly useful for users and society. In a recent article, an agenda for future technical achievements required in this context has been proposed [8], which triggered this effort from AI4EO scholars and the European Space Agency  $\Phi$ -Lab. Accordingly, with this overview article, we aim to explore and discuss in more detail the development and use of AI methods in the field of EO, with an emphasis on ML techniques, and to inspire the EO community, thus to realise the transformative advances made possible by these techniques.

More specifically, this article provides a comprehensive overview of the essential scientific tools and approaches that inform and support the transition from raw EO data to usable EO-based information. The promises, as well as the current challenges of these developments, are highlighted under dedicated sections. Specifically, we cover the impact of (i) computer vision and machine learning; (ii) advanced processing and computing; (iii) knowledge-based AI; (iv) explainable AI and causal inference; (v) physics-aware models; (vi) user-centric approaches; and (vii) the much-needed discussion of ethical and societal issues related to the massive use of ML technologies in EO. Figure 1 summarises the content and organisation of this perspective paper.



**Figure 1.** Conceptual overview of this perspective paper: (a) different levels of algorithms emerge from the areas of machine learning (ML) and interact with computer vision (CV), computer science, and statistics to learn patterns and associations from observational data. The models must integrate domain knowledge and biogeophysical constraints to advance in the modeling and understanding the Earth’s processes. The ultimate goal is to provide intelligent, interoperable, actionable, trustworthy, robust systems whose decisions should be accountable. (b) The field of AI– and specifically, the area of ML within it – interacts (and is embedded into) several systems to realise such ambitious goals, from high-performance computing platforms in digital twins to Earth system model simulations and products, a wide range of Earth observation data, the characterisation and quantification of uncertainty, and the (active) role of the users. (c) The processing chain in AI goes from modeling (e.g. classification, detection, parameter retrieval) with ML, CV and high-performance computing techniques that answer ‘what questions’, to the more ambitious goals of explainable AI, causal relations and ontologies that answer ‘what if’ questions, and finally to communicate decisions, which involves ethical issues, trust and interaction with the user.

# Modelling — Machine learning, computer vision and processing

## 1 Machine learning and computer vision for Earth observation

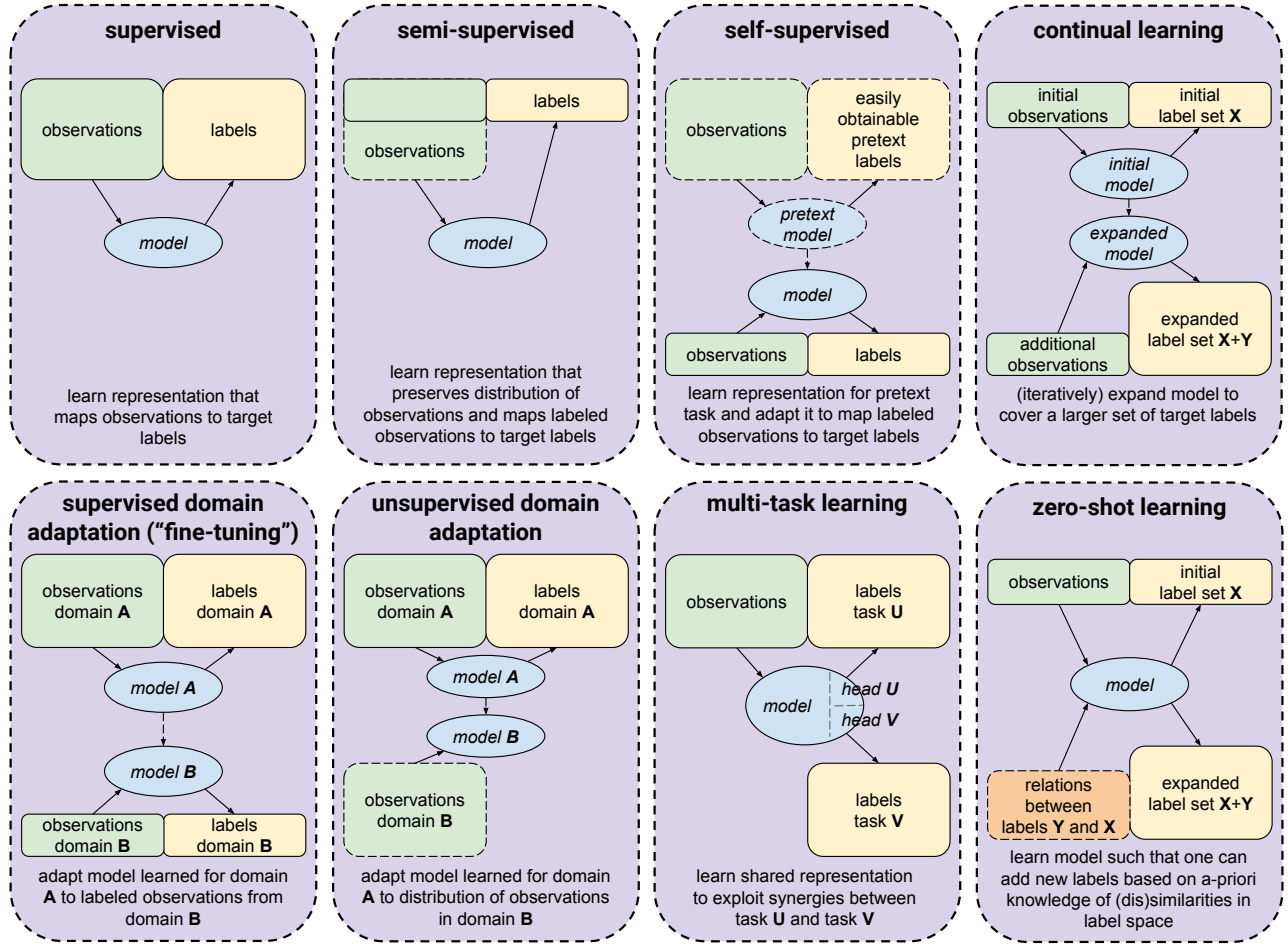
A significant portion of EO data comes in the form of images, so computer vision (CV) is a natural tool for extracting information from the data. Many techniques utilized in CV are rooted in the field of machine learning, such as neural networks. The EO community has benefited from ML-based techniques in computer vision while providing fertile ground for the development of new methods. For instance, neural network architectures developed specifically for image analysis are now commonly used in remote sensing, and the quest for human-like machine perception has inspired new research directions in EO, like monocular 3D perception or geospatial visual question answering. Conversely, some EO-specific challenges do not typically arise in CV studies and call for targeted research efforts.

### 1.1 Specific challenges of Earth observation

In remote sensing, data fusion [9] is routinely performed to jointly exploit multiple data sources (pan-sharpening, optical-SAR fusion, etc.). From the AI perspective, this is a form of representation learning where a joint latent representation shall be found for multiple data sources. It remains an open question how to optimally combine measurements from various sensors into a generic representation [10], while at the same time, we are confronted with an increasingly diverse range of data, including imaging sources, but also ground-based images [11] or social media posts [12]. Arguably, an ideal EO system promptly delivers downstream users the desired geospatial information using any suitable sensor data. This view suggests that EO services could benefit from user-centric abstractions (see Sec. 6) that are tailored to specific information extraction tasks but flexibly draw from a pool of different inputs.

Another challenge is the synthesis between the AI data-driven paradigm and the rich physical knowledge about the Earth system [7]. Physical models excel at describing well-understood relations that can be condensed into explicit, low-dimensional equations. On the other hand, the strength of statistical AI is to construct predictive models for phenomena where the mechanistic relations are poorly understood or intractable, by fitting generic, high-dimensional functions to large data samples. For many quantities of interest in EO (hydrology, vegetation dynamics, etc.) physical models exist. However extracting them from data is a perception task, so one can expect a hybrid approach to be beneficial (see Sec. 5). Not only would the results be consistent and explainable in terms of the underlying processes, but one would also need substantially less training data, yet still be more immune against overfitting.

A specific challenge in EO is the available reference data. Modern machine learning relies on large



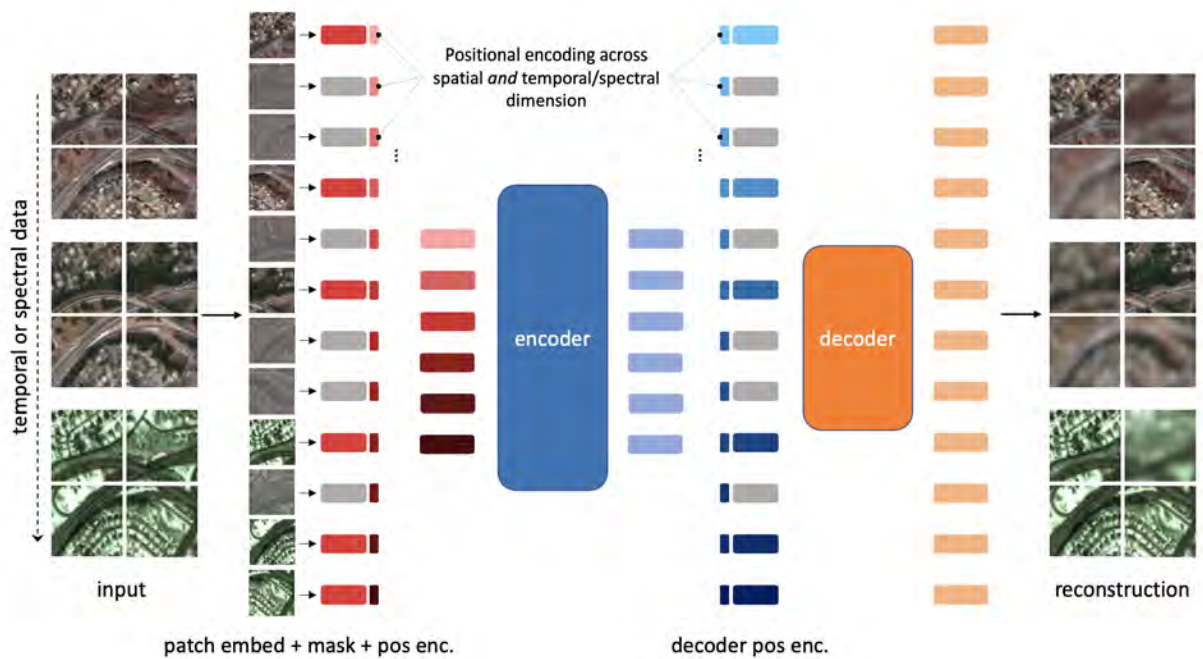
**Figure 2.** Different learning paradigms are relevant for AI applications in Earth observation, each with their own requirements and challenges. See text for details.

volumes of training data with labels (“ground truth” prediction targets) that one might not have for EO, for instance, if collecting them involves fieldwork. For spatially explicit computer vision tasks it is often also assumed that the reference comes in the form of dense 2D label maps and that the labels show a balanced, representative sample. EO reference data is often not as “analysis-ready”. First, the reference may be imperfectly aligned to the images in terms of both location and scale, and it may refer to semantics not visible in the data, like land parcels. Second, ground truth may come as sparse, point-wise observations without surrounding context, making it harder to learn spatial patterns [13]. Third, the data distributions are often extremely unbalanced with long tails [14], where rare, under-represented cases may be important (e.g. very high vegetation, uncommon crop diseases, weather extremes, etc.).

## 1.2 General problem settings relevant for Earth observation

Several problem settings from the general machine learning domain also emerge in the context of CV and EO. Figure 2 shows several of these and how they relate to the more standard supervised and





**Figure 3.** Pre-trained transformer models for Earth observation are typically trained via the masked auto-encoder (MAE) principle. Image from [29] (©2022, BY-SA).

semi-supervised learning setting. Arguably a main frontier of present-day AI is the step from supervised to unsupervised (learning with no labels) and *self-supervised learning* (SSL, learning in a supervised way with no labels by crafting a ‘pretext’ learning task from the data itself – e.g. mask part of the image and predict the masked part, predict the season the image was acquired in, etc.). The current generation of computer vision systems relies primarily on supervised learning from large volumes of *labelled* training data. Biological vision [15], as well as information-theoretic considerations [16], suggest that visual perception can be learned with fewer annotations by abstracting inherent structure in the (unlabeled) data into a representation. While progress has been made, especially for language models [17], unlabeled data are still not utilized to their full potential. There is a clear opportunity for future research, as in EO applications there are typically virtually unlimited quantities of unlabeled data available. When combined with advances in self-supervised learning (e.g., learning via data augmentations [18, 19], or by reconstructing masked image portions [20]) and large-scale pre-training, the abundance of unlabeled data is paving the way for foundational EO models that can be used across a wide range of retrieval tasks (Figure 3); potentially also including language or other information sources beyond EO sensor data [21–23]. Both the access to large quantities of diverse, but unlabeled data and the synergies offered by multi-modality are in line with recent trends towards generative visual AI [24, 25] and foundational vision models [26–28].

Due to the high amount of data that is being gathered (e.g., by satellites), *continuously learning* new

---

information (e.g. new LULC classes, sudden/abrupt changes) from a stream of recently acquired EO data is essential. This requires updating the model parameters by considering the recently acquired training images. On the one hand, re-training the entire ML model is impractical for large-scale EO applications due to computational and storage limits. On the other hand, updating the model parameters based on only recently available training samples may degrade the model’s ability to characterize the previously seen data, an issue known as catastrophic forgetting [30]. Few studies have addressed continual (deep) learning [31–34]– models that can make evolve their learning problem (e.g. the classes) without forgetting the one used for the original training – and meta-learning [35] – models that instead of focusing on a single problem learn how to efficiently fine tune into new problems with minimal supervision – for RS. Their primary goal is to overcome the forgetting of already learned LULC classes and leverage knowledge learned during earlier training rounds to obtain better performance or faster training for new LULC classes, which are considered new concepts. However, these methods require an accurate set of labelled data for the new classes, which may not be available in large-scale EO applications. Additionally, current methods assume that the concept boundaries are known. Such a setup is not well-suited for operational EO applications, where the concept boundaries are unknown a priori and need to be discovered. Further advances in continual learning are, therefore, likely to benefit the EO community.

However, many other learning paradigms can also be explored. *Semi-supervised domain adaptation* (or simply transfer learning) is a setting in which a neural network is trained on a large dataset and later adapted to a smaller dataset [36]. This method turns out to be effective for cases where only limited labelled data of the concept is available, but many labelled data from a related concept are available. *Unsupervised domain adaptation* [37] is a more complex problem setting in which there are no labels available for the target domain. Xu et al. [38] survey several remote sensing applications in this problem setting. There are several other prominent problem settings for which general ML methods have been developed that can be applied to EO applications, including but not limited to: few-shot learning [39] to expand the output range to previously unseen cases with only a few examples; zero-shot learning [40], where the representation should have enough metric structure to handle output values outside the training range; and meta-learning to obtain flexible representations that can be transferred to new tasks with few data [41, 42].

### 1.3 AutoML for Earth Observation

Many machine learning models used in EO are based on neural networks. To make them perform well, one must tune many architectural parameters, on top of the training hyperparameters [43]. The goal of automated machine learning (AutoML) is to support the human in the data science loop by automating the design of a machine learning pipeline, including the choice of a suitable model, its architectural optimisation and the tuning of its hyperparameters [44]. So far AutoML has largely been driven by benchmarks derived from standard CV applications (e.g., the NAS-Bench series [45]) and has neglected EO-specific data such as BigEarthNet [46]. Common characteristics of EO tasks are large amounts

---

of observational data with relatively few high-quality labels and data distributions that substantially depend on the acquisition time and on the geographic location. To best support EO, AutoML must, therefore, go beyond the standard paradigm of selecting models and optimising their hyperparameters and should address more intricate topics like AutoML for heterogeneous spatiotemporal data. This includes, for example, investigations on which type of search spaces would lead to good results within the constraints of the application [47–49], or of the impact of acquiring more labelled data by modelling and analysing the learning curves of certain algorithms [50, 51].

Moreover, recent advances in AutoML enable the configurations of models that are interpretable (c.f. Sec. 4), robust against input perturbations [52, 53] and that have well-calibrated uncertainty scores; properties that will also benefit EO applications. To realise the potential of AutoML for EO, new EO-specific datasets are needed. While initiatives like BigEarthNet [46] can be regarded as first steps towards that goal, no EO-related benchmarks have yet entered the mainstream to a point where researchers interested primarily in AutoML would adopt them and tackle their specific challenges [54, 55].

#### 1.4 Beyond computer vision

Finally, we argue that CV should not be viewed in isolation, but as a tool embedded in the larger context of Earth system monitoring and decision support. From that perspective, an important question is how to best integrate AI systems into the EO infrastructure. The variety of sensors used in EO is high – arguably higher than in most other applications of CV. This creates a need for vision systems that can natively handle multi-sensor settings and exploit synergies between sensing modalities. To do so, it may not be ideal to add CV capabilities to an EO program ex-post, instead they should inform the system design already in the planning phase. For instance, one could plan space missions with a view towards AI methods able to combine their respective observations. We postulate that systematically planned combinations, rather than the present, opportunistic practice [13, 56], could greatly increase the impact of many satellite missions. Also for onboard AI processors deployed in space, e.g., to coordinate constellations or to preprocess data, it is obviously mandatory to include CV in the planning (see Sec. 2.4).

Once AI is included in the system design, the question arises as to why it should be limited to perception and mapping. We speculate that, in the long run, CV will only be a component of much larger AI systems that go all the way to high-level reasoning and decision support. Some AI researchers claim [57] that to that end one may have to integrate statistical AI with relational, symbolic AI (see Sec. 3).

---

### Promising directions for future research

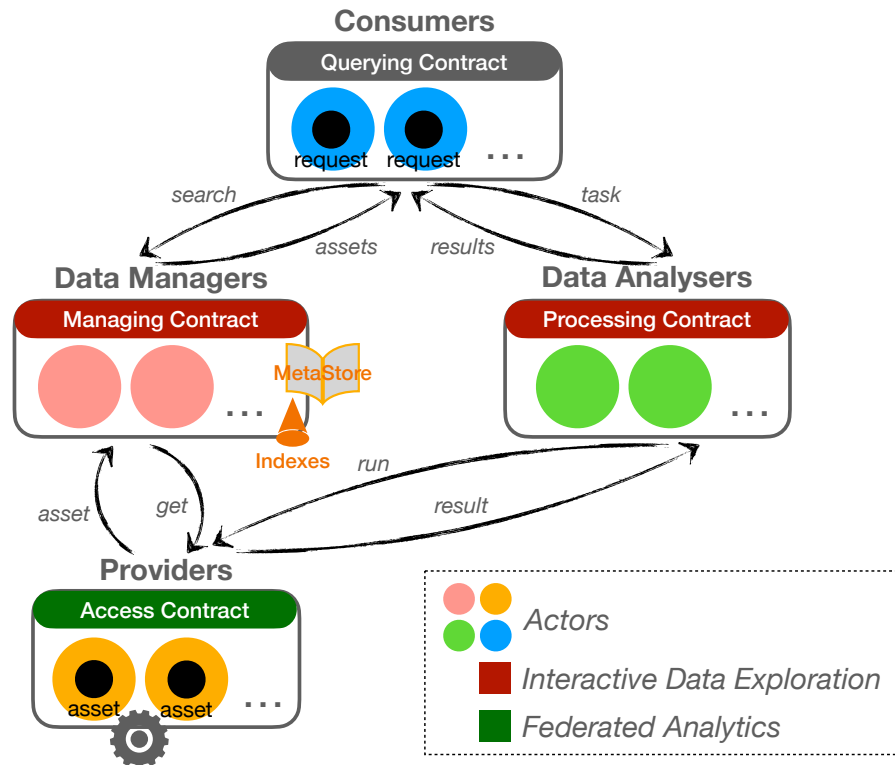
- Multi-modal EO models and cross-modal representation learning to move from sensor-specific to task-specific models
- Self-supervised learning across varying combinations of tasks and/or spectral bands to build foundation models specialized in EO tasks
- Visual generative AI and conditional generative modelling for EO image analysis to tackle tasks as downscaling, sensor-to-sensor translation, data interpolation or simulation
- AutoML solutions specifically tailored towards EO problem settings
- Domain adaptation, low-shot and continual learning to expand AI-based EO to new, data-scarce and/or dynamic applications
- AI-friendly mission planning and EO/AI co-design to better integrate AI into space-based EO

## 2 Advanced processing and computing

This section describes the state-of-the-art and the most promising research directions in the efficient processing of very large and heterogeneous EO data.

### 2.1 Earth observation ecosystems – new system architectures

Recently, different EO exploitation platforms — such as the Google Earth Engine and Sentinel Hub — have emerged to ease large-scale analyses by offering a certain level of processing and data abstraction [58]. According to our knowledge, all exploitation platforms rely on a heterogeneous set of technologies with varying sets of interfaces and data formats, making the federated use of these platforms difficult [59]. We believe that EO technology must be interoperable and accessible to everyone in a seamless manner. The AgoraEO project explores this direction toward a decentralized, open, and unified EO ecosystem, where one can share, find, compose, and efficiently execute cross-platform EO assets [59]. Figure 4 illustrates an architecture for an EO ecosystem based on the actor model, which provides the flexibility to shape an EO ecosystem to applications needs [59]. However, building EO ecosystems is still a big challenge. First, assets are highly heterogeneous, ranging from data archives, processing tools (e.g. SNAP and GDAL), and ML algorithms to high-performance computing clusters and human expertise. This high heterogeneity renders asset management (storing, querying, and composition) an open problem [60]. Second, most EO analytics are exploratory, which requires fast query results. Third, data assets are naturally scattered across multiple sites and platforms, making AI-based analytics hard. It is necessary to devise a framework to support federated analytics.



**Figure 4.** An Actor model-based EO ecosystem architecture. Image from [59] (©2021, BY-SA).

## 2.2 Federated learning

An increasing number of EO applications operate in highly distributed settings, especially with the emergence of EO ecosystems. Nowadays, it is not rare to find EO data highly scattered across multiple sites. However, standard ML algorithms require centralising the data before training, which is problematic in the EO settings due to its high volume. Additionally, centralising datasets might be impossible due to data constraints, such as privacy regulations or legal issues. Federated learning is a way to address this problem of associated governance, privacy concerns, and technical challenges [61, 62, 65]. It aims to learn a local model on each distributed site (i.e., client) and aggregate the resulting local models to a shared global one [63]. It reduces data transfer costs to a minimum, and models can generalise across EO datasets, leading to more accurate and less biased models. Applying federated learning to geo-distributed EO data is challenging because federated learning relies on the general assumption that the underlying data are independent and identically distributed (the i.i.d. assumption). This is rarely the case in operational EO data analysis scenarios. Furthermore, in RS, the data naturally tends to be remarkably heterogeneous – not only because of the variety of sensors, modalities, and characteristics of RS missions but even within the same product. Most of the existing federated learning algorithms have limited capabilities to deal with these conditions, leading to a sub-optimal performance [64]. To overcome this issue, the first multi-modal federated learning framework in RS is introduced in [65]. This framework learns from decentralized multi-modal RS archives based on three modules: 1) multi-modal

---

fusion module (which performs iterative model averaging to learn without accessing data on clients in the case that clients are associated with different data modalities); 2) feature whitening module (which aligns the representations learned among the different clients); and 3) mutual information maximization module (which maximizes the similarity of images from different modalities). Although this framework is found effective for multi-modal learning from decentralized and unshared data archives in RS, there is a need for the development of federated learning models that are associated to low local training and aggregation complexity, high learning efficiency and also low communication cost.

### 2.3 AI interoperability

The need for distributed or federated analytics and learning is increasing rapidly. More and more applications require building models in a distributed or federated fashion (i.e. over multiple data sources). Interoperability, in terms of data and model communication, as well as across different areas to achieve a common goal, e.g., building a global model, is critical in these applications. Without interoperability, achieving common goals efficiently and effectively is almost impossible. Primarily, we expect interoperability as a de facto standard in EO ecosystems, where multiple participants can join to contribute with their assets. This calls for standards representing input data (e.g. satellite imagery) and intermediate data (e.g. local models) to optimise multiple data sources and information models. Unfortunately, most data representations and systems are not interoperable at all. Examples in this direction are the Hub project by ActiveLoop, which aims at providing an AI database, i.e. a common storage representation, for input data [66], the Open Neural Network eXchange (ONNX) format, an open-source standard for machine and deep learning [67], or OpenML [68], an ecosystem for sharing ML datasets, algorithms and results. However, real interoperability is still far from reality in EO. AI interoperability must be solved if we expect EO ecosystems to succeed.

### 2.4 On-board data processing

One of the most important goals of EO is to transfer satellite data processing from Earth (ground segment) to space (space segment) through the development of so-called onboard payload data processing. Indeed, quite often, the raw data generated by modern instruments is more than what can be transmitted to the ground. This makes using various signal processing and compression techniques necessary to reduce the amount of data transmitted, especially for small satellites. However, onboard processing requires several characteristics, such as flexibility, robustness and low computation burden. Moreover, the design of optimum schemes generally involves an appropriate combination of software and hardware solutions. For this reason, AI models represent a valid technology [69]. In this case, the highest processing demand is required during the training phase (which is operated on the ground segment), while less effort is necessary to yield the onboard products with hardware prototyping that can be done in fast loops. We would like to note that for the onboard implementation, the computational complexity of deep neural networks can be reduced using various techniques, which further facilitate their application under the constraints of the space environment. As an example, a quantization technique that aims at

---

reducing, with a negligible loss of accuracy, the precision of the numerical representation of the weights (e.g. from variables coded in single-precision floating-point format to those coded in signed integer format) is found effective for this purpose [70]. ESA has already launched a first demonstration mission named Phi-sat-1, characterized by onboard AI processing for cloud detection [71]. More challenging and diverse AI-based applications are being developed for implementation on the payload of the Phi-sat-2 mission [72]. The on-board reconfiguration of the AI components, either from the data perspective (data-centric) or from the module perspective (e.g. slimmable networks [73]) represents one of the future trends of this technology and it is aligned with the more general goal of the next on-board imaging systems in space missions, which is to achieve partial or complete autonomy [74].

#### Promising directions for future research

- Development of decentralized, open and interoperable EO ecosystems
- Development of federated learning algorithms that are not only robust to training data heterogeneity across clients but also associated with low local training and aggregation complexity, high learning efficiency, and low communication cost
- Development of lightweight AI models for on-board processing of EO data



---

## Part II

# Understanding — Physics-machine learning interplay, causality and ontologies

## 3 Knowledge-based AI and Earth observation

Knowledge-based AI includes knowledge representation, reasoning (incl. constraint satisfaction) and planning (incl. scheduling). We discuss several opportunities for knowledge-based AI relevant to EO. The top row of Figure 5 illustrates the main avenues in this direction.

### 3.1 Decision modelling and decision support with Earth observation

A prototypical example of the knowledge-based approach to AI is expert systems, designed to mimic the behaviour of human experts in narrow areas of expertise [82]. They consist of a knowledge base (capturing human knowledge) and an inference engine (emulating human reasoning) and overall support the decision-making process. The spirit of expert systems has survived in its purest form at the intersection of AI and decision support and is exemplified by the DEX (Decision EXpert) approach [83]. DEX is a hierarchical multi-criteria decision modelling method, which represents decision preferences on alternatives described by qualitative variables in the form of decision rules, very similar to the IF-THEN rules of expert systems, and supports a variety of knowledge-explanation methods. DEX has been successfully and extensively used in various areas, with increasing applications concerning agriculture and its environmental impacts. One example of such an application is SoilNavigator, a field-scale decision support system for assessing and managing soil functions [84]. Due to the lack of measured data on soil functions, the SoilNavigator decision support system assesses all of them in a knowledge-based manner.

The potential of EO data in decision-making is increasingly recognised. For example, the OECD [85] has prepared a document entitled “EO for decision making”, which clearly states the importance of EO data for decision making, especially when combining it with other geospatial (e.g. ground data on air pollution), administrative and socio-economic data, in the context of assessing environmental risks and their impacts on humans and ecosystems. Combining knowledge-based approaches to decision support with EO data or downstream products derived by machine learning will unlock a whole new range of applications of AI that can help in achieving sustainable development goals [86].

### 3.2 Knowledge-based machine learning and Earth observation

Most ML approaches today are data-driven and require large quantities of data to produce accurate and reliable models, which are typically not transparent and do not provide explanations. However, ML



---

approaches that make use of formally represented domain knowledge, in addition to data, have a long tradition in the field, starting with explanation-based learning [87]. Inductive logic programming [88] uses logic programs to represent data (examples), domain knowledge (background knowledge) and learned models (hypotheses). This gives rise to a broader learning paradigm known as relational learning or relational data mining [89]. Combining logical representations with probability theory has led to the development of approaches that learn probabilistic models expressed in logic within the paradigm of statistical relational learning [90] and statistical relational AI [91].

The knowledge used in ML, in addition to data, can take different forms. Taxonomies are a form of domain knowledge that represents relations among concepts in the domain of study: in predictive modelling and classification, such taxonomies can represent hierarchical relations among the classes predicted. Typically, examples in this context have multiple labels from the taxonomy and the task at hand is known as hierarchical multi-label classification [92]. Land cover classification schemes [93] are typically hierarchical. However, most ML approaches to land cover classification ignore this and learn models for each land cover class separately. Multi-label classification of land cover [94–96] and classification with label hierarchies [97, 98] is attracting more and more attention.

Other forms of knowledge that ML approaches can use include constraints, as considered in constraint-based data mining [99]. Rules can be hard or soft and can (in predictive modelling) impose preferences on the predictive power, form (language constraints), or size of the learned models. Constraints on clustering can, for example, specify that specific data points should belong to the same cluster.

Last but not least, explanation and the use of existing domain knowledge play a prominent role in computational scientific discovery [100], which includes the topic of automated modelling of dynamical systems in science and engineering. Approaches such as process-based modelling, exemplified by the tool ProBMoT [101], integrate data-driven and knowledge-driven approaches into modelling, where ontological domain knowledge in the form of templates can be used as model parts.

### 3.3 Ontologies for open science in machine learning and Earth observation

Ontologies formally represent knowledge about a domain of interest, focusing on entities in the field and the relations between them. They mainly comprise classes (or concepts), properties (or attributes), instances (or class members) and relationships. Ontologies are crucial in providing controlled vocabularies for semantic annotation of datasets, which allows datasets to be FAIR, i.e. findable, accessible, interoperable and reusable [102]. With the rapidly increasing amounts of EO data, datasets must be FAIR according to the aforementioned criteria to maximise their added value.

Currently, EO data are primarily managed through spatial data infrastructures, which facilitate the exchange, sharing, use and integration of geospatial data based on interoperability standards. Spatial data infrastructures follow the FAIR principles in some aspects (e.g. have registered or indexed meta-data in a searchable resource and are also accessible, even when the data are no longer available). However, they are not yet entirely FAIR-compliant [103], mainly due to the lack of ontologies for EO

---

data that would provide vocabularies for meta-data that follow FAIR principles. We expect significant further development of ontologies for EO, ML, and other areas of AI. These should be connected to commonly accepted upper-level ontologies, facilitating their interoperability. Linking the ontologies for EO and AI will significantly facilitate open science at the intersection of EO and AI. As a recent example, [104] proposed a neuro-symbolic paradigm for a self-supervised ontology learning for natural hazard identification and monitoring.

#### Promising directions for future research

- Utilize Earth observation data to derive better decision support systems
- Utilize domain knowledge to be included as input in machine learning systems
- Development of extensive ontologies for Earth observation

## 4 Explainable AI and causal inference

Experimentation and observational data analysis form the foundation for accurately modelling and understanding physical phenomena [81]. However, model fitting alone is often inadequate, and there is a desire to comprehend and characterize the system’s behaviour. The scientific consistency, reliability, and explainability of results generated with AI models are crucial when dealing with complex systems such as the Earth or its climate systems [8, 80, 105]. Developing robust models that can be visualized, queried, and interpreted is essential. Achieving transparency, interpretability, and explainability in ML models for geosciences and RS is key for several reasons: 1) to ensure the model’s performance consistency, generalizability, and robustness, 2) to enhance understanding of the Earth’s system and 3) to foster broader adoption and confidence among domain scientists. However, model interpretability is challenging because the model assumes a specific causal relation (inputs cause outputs) that may not be correct or complete. Retrieving causal relations from observations and quantifying the effects of certain interventions has been a longstanding human goal and remains a key objective when employing AI approaches, particularly ML techniques, for scientific inquiries [106]. This section will explore recent advances in explainable AI and causal inference for EO. The second row of Figure 5 summarizes the main concepts presented.

### 4.1 Explainable AI

Due to their size and complexity, many successful neural network models pose a challenge in comprehending or studying their inner workings at a level of abstraction that provides insight into their predictions. Various explainable AI (XAI) methods exist for different learning paradigms [107, 108]. These methods encompass *post-hoc interpretability*, employing advanced feature ranking techniques [109, 110] or analyzing activation maps within the learned network (e.g., GradCam or layer-wise relevance propagation), as well as *interpretability by design*, involving networks explicitly designed to be explainable [111]. These techniques can generate spatially explicit and temporally resolved explanations of what AI models have

---

learned, emphasizing specific spatial and temporal scales [110].

“Post-hoc” methods, including backpropagation-based or feature attribution techniques like layer-wise relevance propagation, DeepLift, and integrated gradients [107, 108], also encompass input feature ablation, occlusion, and perturbation-based approaches such as model reliance and Shapley value sampling. These methods address the challenge of meaningfully relating a complex neural architecture’s output or intermediate activations to its input features. Attention mechanisms and transformers, which have shown promise in Earth sciences [112], provide ante-hoc explanations by assigning coefficients to different spatio-temporal locations, indicating their importance when the model generates an output. These coefficients form saliency maps, facilitating interpretation. In contrast to traditional recurrent neural networks, attention mechanisms allow insight into long-range interactions in space-time series.

“Explainable by design” methods embed interpretability into the model. Examples include models based on semantic bottlenecks, where intermediate processing pipelines (e.g., a classifier of concepts relevant to the downstream task) are learned concurrently with the final task. Recent literature in EO proposes such models to understand subjective appreciation of landscapes [113], align forest type classifiers with foresters’ decision-making processes [114], or unsupervised exploration of what defines a wild landscape [115].

## 4.2 Causal inference

Despite the impressive predictive capabilities of current machine and deep learning methods, there is limited actual learning, as these algorithms excel in fitting complex functions without providing a clear understanding of underlying causal relations [78, 81, 105, 106, 116]. Analyzing Earth’s complex systems with multivariate and non-gridded datasets faces missing data, non-linearities, and non-stationarities. Variables and processes are often coupled in space and time, with (tele)connections being long-range, discontinuous, and varying in strength. Addressing this issue is crucial for identifying predictors, developing robust models, and ensuring correct answers for the right reasons.

Causal inference aims to discover and explain a system’s causal structure based on data, models, and assumptions [78, 106, 117]. It addresses fundamental Earth science problems: 1) hypothesis testing, 2) discovering latent factors beyond spurious correlations, and 3) understanding systems from observational data and assumptions. True understanding involves knowing the entire causal chain, enabling predictions of consequences and analysis of interventions and counterfactuals. Three main fields emerge as relevant for modelling and understanding processes in the Earth system.

**Interventional Analysis** Performing interventions on simulations, though challenging, enhances the efficiency and robustness of causal discovery. The choice of interventions requires careful consideration, respecting the spatio-temporal trajectory of the system. Physical and causal sense is crucial, ensuring a genuine distributional shift in the system’s dynamics [106].

---

**Counterfactuals Analysis** Counterfactual analysis imagines alternative outcomes based on different actions. Attributing causes of extreme events to climate change involves counterfactual thought experiments. Using algorithms like MACE, DACE, and NSGA-II, counterfactual modules facilitate model interpretation in human-like language, fostering explicit user-machine interaction [106].

**Causal Discovery** Observational causal discovery extracts cause-effect relationships from multivariate datasets, moving beyond correlation. Various methods allow learning causal relations even without explicit time involvement, including probabilistic graphical models, structure learning, dynamical systems, Granger-causality, instrumental variable paradigms, and additive noise models. Understanding complex phenomena related to climate change requires discovering teleconnection patterns, where recent deep learning approaches show promise [118].

#### Promising directions for future research

- Causal machine learning promises to integrate ML models with causal inference to better understand and predict complex earth systems dynamics observed through remote sensing
- Exploring spatiotemporal causal models can lead to significant breakthroughs in discovering the impacts (effects) of environmental factors on global change
- Investigating counterfactual scenarios through remote sensing and climate data can provide insights into the potential impacts of different interventions e.g., on socio-economic factors, aiding policymakers in making informed environmental decisions

## 5 Physics-aware machine learning

Machine learning has been widely and openly criticised because of the difficulty of interpreting models and the lack of physical consistency. It may happen that models are so big and overparameterised that accuracy compromises fulfilling physical laws, such as mass or energy conservation. Losing contact with physics should always be avoided if we want to design robust and reliable (ML) models. This disconnection has implications for the consistency of products and parameter estimation in data assimilation schemes and forecasting tools and in understanding the underlying processes from Earth's observational data.

Historically, physical modelling and ML have been treated as two fields under very different scientific paradigms (theory-driven versus data-driven). However, integrating domain knowledge and physical consistency has been suggested as a principled way to provide solid theoretical constraints on top of the observational ones [80]. The synergy between the two approaches has gained attention by either redesigning the model's architecture, augmenting the training dataset with simulations, or including physical constraints in the cost function to be minimised [80, 119]. Integrating physics in ML models may improve performance and generalisation and, more importantly, consistency and credibility. The hybridisation has an interesting regularisation effect, given that the physics limits the parameter

---

space to search and thus discards implausible models. Therefore, physics-aware ML models better combat overfitting, typically become simpler (sparser), and require less training data to achieve similar performance. Physics-aware ML thus leads to enhanced computational efficiency and constitutes a stepping stone towards achieving more interpretable and robust ML models [120, 121]. This section presents advances in physics-aware machine learning and its applications to the specificities of EO. The third row of Figure 5 summarizes the main concepts covered in the section.

### 5.1 Inverting physical models with machine learning

Biogeophysical parameter estimation and retrieval with ML aim at learning a mapping from an observed satellite spectrum to an underlying biophysical parameter. Unfortunately, relying only on observational data (both spectra and in-situ measurement) to perform regression is expensive, time-consuming, and very often infeasible [122, 123]. This problem can be better posed by exploiting the wealth of simulated data available through physical models such as radiative transfer models (RTMs) and electromagnetic models (EMs). They solve what is called the *forward problem*, that is, they generate physically meaningful observational data (e.g. spectra) from state parameters (e.g. leaf-level or canopy parameters). Once the simulations of a particular observed bio-geophysical scenario (e.g. a forest) are given, they provide the electromagnetic quantity that would be measured by the EO sensor (e.g. the back-scattering coefficients or reflectances). Using these models, we can generate as many input-output pairs as we need for training and testing our ML scheme to do the inversion, i.e. to estimate parameters from RS data [124, 125]. Another example is the prediction of quad-pol signatures for single-polarized SAR data [126]. The use of an RTM/EM model can be beneficial from different points of view: (1) they can help with data generation balancing by compensating for the lack of ground measurements; (2) they can help in understanding the bounds within which the desired parameters can be retrieved less ambiguously; and (3) they can help in the selection of the optimal inputs of the ML model.

### 5.2 Hybrid machine learning: interactions of transfer models and ML

ML algorithms and physics can be fully blended in several ways extending traditional approaches based on data assimilation. Still, simple constrained optimisation methods often work well in practice [127, 128]. Indeed, by inferring the constraints to the input/output pairs based on physical models, ML models can obtain improved performance. Another option is to learn ML emulators combined with purely data-driven algorithms for model inversion [79]. Finally, one can also consider a fully coupled neural network where layers that describe complicated and uncertain processes feed physics layers that encode known relations of intermediate feature representation with the target variables. The promise of such a hybrid modelling approach [129] is to combine the excellent predictive abilities of modern ML, fuelled by the “unreasonable effectiveness of data” [130], with the causal consistency, explainability, and data efficiency of classical physical modelling.

Physical models require setting parameters that are rarely derived from first-domain principles. ML

---

models can learn such parameterisations. For example, instead of assigning vegetation parameters empirically to plant functional types in Earth system models, one can learn these parameterisations from proxy covariates using ML, and thus achieve flexibility, adaptability, dynamics, and context-dependent properties [131].

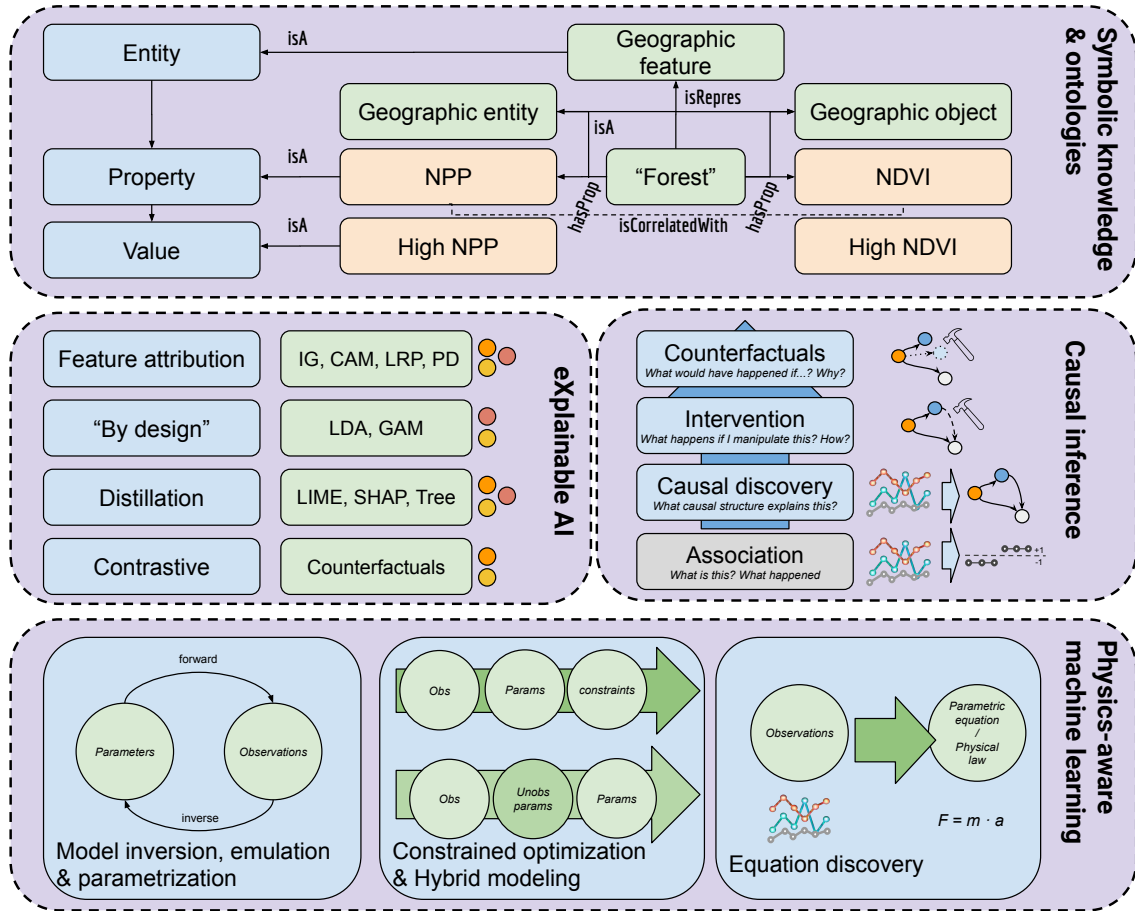
Emulating models in geosciences, climate sciences, and remote sensing is gaining popularity [80, 122, 123]. Emulators are ML models that mimic the forward physical models using a small yet representative data set of simulations. Once trained, emulators can provide fast-forward simulations, allowing improved model inversion and parameterisations. However, replacing physical models with ML models requires running expensive offline evaluations first, and alternatives exist that construct the model and choose the proper simulations iteratively [79, 132].

### 5.3 Discovering equations from data

As in many fields of science and engineering, Earth system models describe processes with a set of differential equations encoding our prior belief about the dynamics and variable interactions. Learning ordinary and partial differential equations (ODE/PDEs) from stochastic observations is perhaps one of the most challenging problems nowadays in statistical learning. Several approaches exist for solving this problem, ranging from equation-free modelling [133] and empirical dynamic modelling [116] to automated inference of dynamics [134]. Imposing sparsity on equation discovery is a sensible assumption that allows us to identify the governing equations of the biosphere more sharply [135]. Other criteria, such as compositionality, unification of theories, and interestingness, can help in learning laws and equations from data while trying to alleviate the *equivifinality* (non-identifiability) issue, that is, finding many equally accurate models for the wrong reasons. For a comprehensive overview of equation discovery techniques with case studies in Earth and climate sciences, we refer the interested reader to [81].

#### Promising directions for future research

- Using hybrid machine learning to emulate costly radiative transfer models (RTMs) to improve the speed and accuracy of EO data analysis, allowing more trustworthy and interpretable simulations
- Hybrid approaches allow for more precise parameter estimation in Earth science models that enable more accurate predictions, calibrated uncertainties, and a better understanding of environmental processes
- Hybrid models yield sparser, more interpretable ML models that enhance transparency and usability in EO, making complex data more accessible and actionable



**Figure 5.** Methodologies for understanding complex systems such as the Earth. **Top:** Ontologies can represent both symbolic and numeric knowledge, reason only based on cognitive semantics and share knowledge on the interpretation of data, e.g. remote sensing images. Following [75], we can describe complex concepts such as “forest” as a particular ‘Entity’ with some attributes or properties (e.g. NPP and NDVI levels) which can have associated values (e.g. high or low). This object can also be defined according to many intertwined properties and value levels. **Middle:** Many XAI methods are available to obtain explanations from ML models, typically categorised in feature attribution methods, explainable by design, distillation-based and contrastive [76], which fall under post-hoc methods (indicated by the orange dot), local (indicated by the yellow dot) or model-agnostic (indicated by the red dot). However, XAI only explains what the ML model learned, which can be correct for wrong reasons. Causality advances in the interpretability [77, 78], and proposes steps beyond the mere association rules developed in conventional ML -that can only answer ‘what’ questions- by advancing in discovering causal structures in the data, answering ‘how’ (intervention analysis) and ‘why’ (counterfactuals) queries. **Bottom:** ML and domain knowledge can interact in many ways [79]: from inverting or emulating complex codes (e.g. RTM or climate models) to the estimation of parameters constrained by physics-based losses or ML-model coupling in hybrid modelling that allows estimating unobserved parameters [80]; and to the direct discovery of fully interpretable equations and laws from data [81].



# Communicating — Machine-user interaction, trustworthiness & ethics

## 6 User-centric Earth observation

Making EO data accessible in catalogues is not enough. Improving users' experience must be high on the agenda too. Such experience is currently still not optimal, since EO data analysis and interpretation is still performed in a laborious way, by repeated cycles of trial and error, without reaching the desired degree of flexibility and robustness. Furthermore, the overwhelming amount of data puts strong limitations on the extent (geographic coverage within acceptable costs), depth (causal analysis, evaluation of the implications) and response time of human interpretation. As a consequence, a significant amount of EO data remains unused [136]. User-centric EO aims at enabling the interpretation of EO data content, associating it with other sources of information, understanding user inquiries in an interactive dialogue (thus refining the expression of needs), distil data content, and suggesting the most appropriate information items or the interpretation alternatives.

### 6.1 Decreasing the user load with active learning

The first form of interaction is to be found in the support of human operators creating the labels. Extracting labels directly from images without guidance is time-consuming and can be difficult to crowdsource. EO crowdsourcing may not be impossible, but is hampered by factors such as the lack of agreement about class definitions among communities, domains and even users. The label collection process can be optimized via human-machine interaction with active learning [137]. The active learning process is iterative: at each iteration, the learning algorithm automatically chooses the most informative unlabeled samples for manual labelling from a human expert (supervisor) and the algorithm is retrained with the additional labelled samples. This way, the unnecessary and redundant labelling of samples that are not informative for the classifier is avoided, significantly reducing the labelling cost and time. Questions of sample informativeness (uncertainty and diversity of samples) have been at the core of the debate. At the same time, equally important aspects such as accessibility costs [138] and user's skills in labelling [139] have been relatively neglected so far. Also, the question of the applicability of active learning methods to DL models is still an open debate, since single samples' usefulness is diluted by the batch-learning nature of deep neural networks [140] or focusing on sample diversity and misclassifications of the current model [?] are also being actively explored. A comprehensive and systematic overview of deep active learning methods is provided in [141], showing promising avenues primarily on improving sample selection strategies, optimising training methods and improving task-independent models.



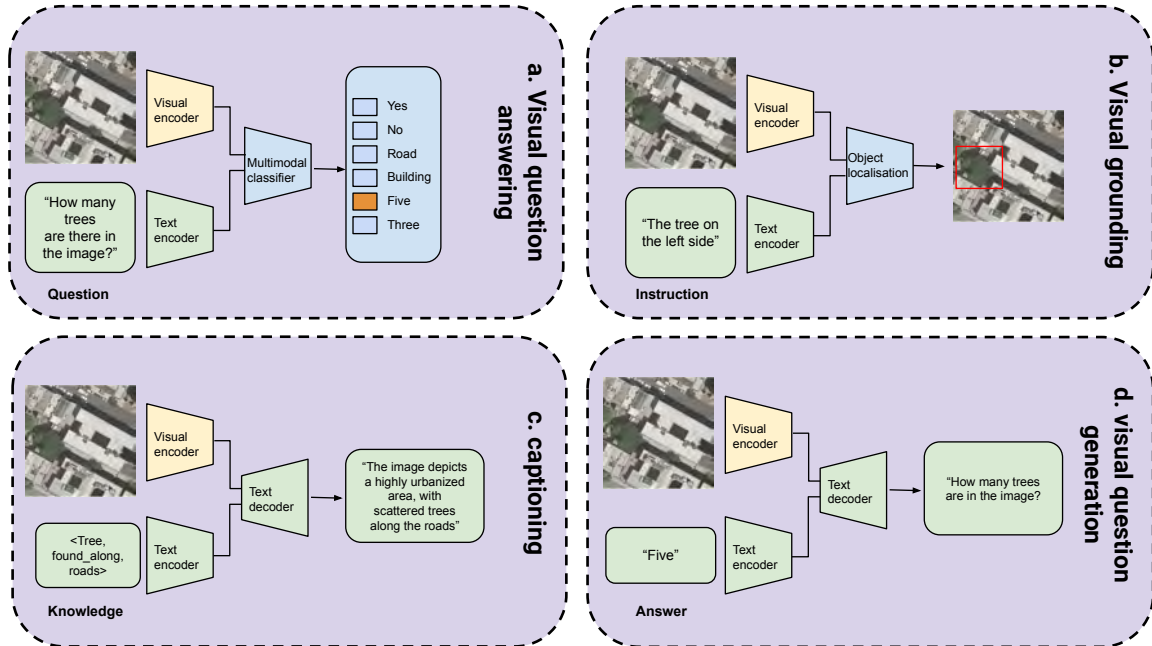
---

## 6.2 Language-based interaction

Besides label collection, interactivity can play a major role in opening access to EO information to virtually anyone. To reach such a goal, one must work on the access interface to the information itself. The engineering and modelling phase is a barrier insurmountable for many users, who cannot exploit the potential of EO because of technical impediments. Creating access in a more inclusive way, e.g. through natural language, would open EO to a much broader set of users, including decision-makers, the press and virtually all areas of society. Language opens EO access to society: methodologies ensuring this access via multimodal processing connecting natural language processing (NLP) models (understanding users' questions) and visual models (from computer vision, reviewed in the other sections) is critical. The interaction between language and remote sensing has become a very active area of research, whose main components are reviewed in [142] and are sketched in Fig. 6. In this section, we focus on two of them: visual question answering and image captioning / storytelling.

- Visual question answering [143] is a set of methodologies aiming at creating such links. The first approaches are showing proofs of concept for Sentinel resolution and VHR resolution imaging [144] or emphasising changes in EO imagery [145]. These first methods still have a large margin of progression, especially regarding the fusion strategies to mix language, the way to generate, learn and exploit from more realistic language, the usage of large language models (LLM) from NLP (e.g. BERT, GPT) [146] and the injection of EO context (or *prompting*) in these models [147]. These research lines will permit the development of digital assistants able to understand real questions from users and retrieve action-ready information from EO. The emergence of EO chatbots, scanning EO data and retrieving information on demand or dialoguing with users is an exciting avenue for both research and the industry.
- Besides answering user-specific questions, one could be interested in gathering general information about a scene for writing a journalistic story or simply having a summary of the image content for storytelling. Such image transcription in language form is often referred to as image summarisation and captioning. In this context, recent research has created large datasets, developed practical algorithms, and extended potential applications for RS image captioning tasks [148, 149]. Most existing remote sensing image captioning methods are based on encoder-decoder frameworks, which learn features from the input image, and then convert the encoded features into natural language descriptions. Attention mechanisms and transformers are extensively explored to this end.

Beyond tasks where generating answers or a description of an image is the final goal, image/language alignment is also an interesting usage of language models, where language-based representations (learned with vision language models such as CLIP [150]) are used to build a bridge between semantic concepts across modalities: in [151], representations from language models are aligned to those from remote sensing images by first aligning them to ground images co-located with the remote sensing ones. In [152], authors learn semantic segmentation models, where intermediate features are aligned to text embeddings

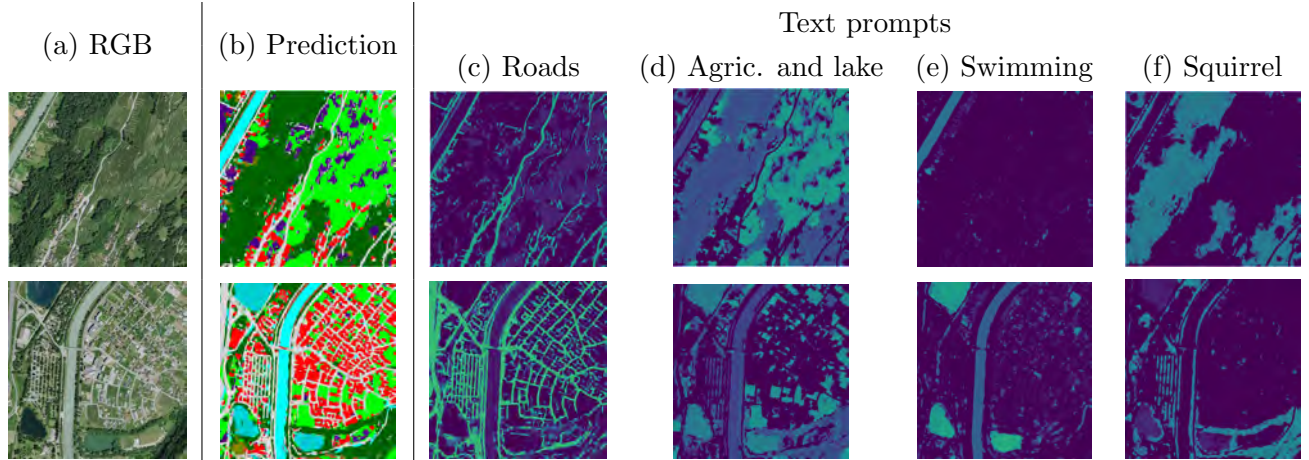


**Figure 6.** Main current advances at the interface of vision and language for remote sensing: a) visual question answering: answering a text question about an image, b) visual grounding: identifying objects in the image with text; c) captioning: describe the content of the image, possibly using some external knowledge; d) visual question generation: generate a question about the image that has the specified answer (could also use external knowledge). In the plots, yellow parts are visual encoders, green are text-based encoders/decoders, and blue are ML models for classification / object detection. Inspired by [142].

from LLMs, allowing to activate image regions with respect to unseen concepts that are semantically close, but unseen during training, in a zero-shot manner (Figure 7 illustrates the out of vocabulary search abilities of a model initially meant for land cover mapping after its feature representations are aligned to those of a language model; since the features of water pixels now have representations close to those of the word ‘water’, one can now search for new concepts by simply typing the corresponding words as text prompts). In [153], authors introduce the task of referring remote sensing image segmentation (RRSIS), aiming at segmenting out the objects in a remote sensing image to which a given natural language expression refers.

### 6.3 Enhanced visualisation

The exploration of big multimodal data, including EO data multi-mission images, meta-data, or text calls for new visualisation techniques able to represent data in a way that eases interpretation and explanations, but also the discovery of unexpected information, anomalies and causal relations (see Section 4). Succeeding in such exploration implies designing visualisation techniques, visual analytics,



**Figure 7.** Image-text concepts alignment from [152]: given a model processing RGB images (a) and classifying landcover classes (b), the per pixel image features are aligned with a contrastive learning objective to the text concept vectors of GLoVe [154]. Once aligned the image can be queried with text prompts for activation of concepts unseen during training, from semantically close to the landcover classes (c-d) to more far away, yet semantically related ones, such as ‘swimming’ (e) or ‘squirrel’. Activations in (c-f) are normalized on a scale from 0 to 1, with dark blue colour for low values and light green for higher values. (©, 2023 IEEE).

human-computer interaction, and augmented or virtual reality [155].

Early geographic EO visualisation attempts were ‘virtual flights’ on 3D terrains, rendered using 3D models from stereo or InSAR observations. Artefacts and noise removal, as well as multi-modal and multi-resolution representation in the RGB space, were the core issues. More recent visual data mining employs immersive visualisation [156] of the EO information content with high user interaction. Immersive video environments like CAVE (a cube-sized room, where data are projected on 3 to 6 inner walls and the user is standing in its middle) or VR sets enriched with AI methods were used for multi-dimensional space adaptive projections. These visualisation methods and the analytics used interfaces must be predictive and adaptive, learning and anticipating the user behaviour and collaborating with them. This means understanding the user’s intentions and context and establishing a dialogue to transform non-visual sensor data and information into understandable representations, potentially exploiting the approaches described in Section 6.2.

---

#### Promising directions for future research

- Multimodal vision language models used to interactively query beyond concepts seen during training and adapt models across nomenclatures
- Involving information beyond pixels in the reasoning process, for example involving knowledge graphs mined with NLP tools [157]
- Interactive (chat)bots able to retrieve information on demand, following reasoning logic and feedback from users

## 7 Earth observation and society: the growing relevance of ethics

In this final section, we consider ethical issues linked to Artificial Intelligence, Earth Observation (AI4EO) and society under the following heads: (i) AI and Data Ethics (ii) Ethical Opportunities for Sustainability: Environment and Health

### 7.1 AI and Data Ethics

Ethical issues related to combined uses of AI and EO can emerge either in the context of AI ethics or Data Ethics. The principles linked to each of these overlap. However, it is useful to understand them separately.

#### 7.1.1 AI Ethics

Jobin et al. [158] recently categorized AI ethics guidelines that most commonly appear in various national documents under 11 heads. Combining insights from Jobin et al. and four other relevant sources, Kochupillai et al. [159] created 6 major categories of ethical duties (see Table 1) and corresponding ethical consequences that are most relevant for the AI4EO community.

Several of these categories overlap with the principles of trustworthy AI compiled in the EU Ethics Guidelines for Trustworthy AI [162]. However, ethics guidelines, even when neatly categorized, may often appear vague or too high-level to EO scientists. [163] To gain a better and more practical understanding of ethics, it is useful to know that, at the simplest level, ethics is about minimising harm and maximising benefits using appropriate (fair, honest) means. This may not, however, be an easy or straightforward exercise in the context of everyday AI4EO research. In all instances involving ethical issues, several approaches can be utilized to reach an ethically acceptable decision or outcome. For example, dominant Western approaches to ethics consider the duties of the actor and consequences of the action/inaction (so-called deontological and consequentialist approaches to ethics) to determine their acceptability. [164]

Yet, these approaches may be impractical in emerging technologies such as AI4EO; scientists may not be able to clearly predict the future consequences of their research in the present and may not be fully aware of their duties. Recent work recommends combining approaches emerging from Eastern philosophy,

**Table 1.** Six categories of ethical duties (from [159]).

1	Privacy	In the context of EO, this includes (i) Non-stigmatisation; (ii) respecting or protecting autonomy or freedom of persons and communities involved; (iii) Self-determination, including national sovereignty, and (iv) ensuring proper data governance, including clear rules regarding data ownership and licensing.
2	Honesty	Includes (i) transparency (including how representative the data is), (ii) explainability, and (iii) data veracity. It also requires good data segregation practices, e.g. preventing the overlap of test and training data.
3	Integrity	Includes (i) ensuring technical robustness; (ii) proper uncertainty quantification and reporting of the limitations of the model; and (iii) safety/security, including ensuring national security (e.g. by selective disclosure of sensitive data) as well as safety vulnerability of DL models to adversarial examples [160] or patch attacks [161].
4	Fairness	Includes. (i) Non-bias (including non-bias in training data); (ii) Non-discrimination and diversity (including adequately diverse and representative datasets) (iii) Socio-cultural sensitivities, e.g. while labelling data; (iv) Democratic creation of standards, e.g. when applying labels or using the data for policy decisions.
5	Responsibility	Includes (i) Human agency and oversight, or human-in-the-loop; (ii) duty of care, especially to understand corresponding ground realities; (iii) accountability; and (iv) ensuring social security and cohesion.
6	Sustainability	Includes (i) Scientific, social, and cultural sustainability (e.g. preventing “colonialist science.”); and (ii) Environmental sustainability.

particularly the intentions/desires of the individual researcher and/or the research project/program as a whole. Such considerations would include the evaluation of whether each action or research step corresponds with ethically well thought-out intentions or desires [165].

Consider the ethical principle of justice and fairness, which requires, for example, that appropriate measures be taken to ensure consistency and accuracy when applying sensitive data labels, such as slums. Even when accurate, such labels can lead to stigmatisation or even to uprooting people from their living spaces. Accordingly, taking into account the intention of the research or researcher, presumably, to improve living conditions in congested urban areas, ethics calls for more constructive labelling. This may also entail re-thinking research designs that can go beyond the labelling of slums and help uplift or improve the socio-economic conditions of marginalized segments. The principle of fairness, therefore, can be applied not only at the time of labelling data but even earlier in the research life-cycle, i.e. at the time of selecting research focus areas and research questions. Ethics is equally important at other stages of research - from allocating research funding to converting research findings to practical applications.

Similarly, the issue of responsibility, which includes human agency or human-in-the-loop features in AI designs, needs careful consideration. While active involvement of human actors (diverse stakeholders) in every aspect of AI4EO research may not be possible, human agency and oversight must not be limited within the confines of the system and its narrow research question. Such oversight is also needed to understand the context in which the designed system is going to be implemented or introduced and

---

the impact that it may have on real people, their lives, livelihoods, cultural sensitivities, and their environment. These insights are crucial to ensuring that the intended purpose of the research and development activity is actually going to be accomplished on the ground.

If the human agency and oversight requirement does not evolve considerably in the AI4EO research context, intricate human realities, relationships, and expediencies can be overlooked or ignored, leading to significant ethically undesirable consequences. [159] Ethics, therefore, calls for the humanisation of projects and supplementing any simulations and recommendations with multi-disciplinary social science and ground realities research. Such realities are not limited to ground truth or in situ data as understood in conventional EO language, but also include the socio-cultural, economic, and political context of the region that can make the consequences of implementing recommendations emerging from AI ethical or unethical.

In this context, the rapid emergence of Large Language Models (LLMs) can increase the ethical risks related to inaccurate or outdated data, non-representative or biased datasets, errors, and misinformation, [166], which can be further hidden by the size and scope of recent LLMs trained on large quantities of Web data. Moreover, AI4EO has a risk of leading to a kind of colonialist science, as Earth observation happens remotely. AI4EO approaches may give scientists who study places without ever going there a false impression that they know the place and its circumstances. This can lead to misunderstanding or a complete lack of understanding of local contexts and may even lead to disempowerment or discrimination against local communities, their rights over ancestral lands or occupations, their cultural values, and even their means of livelihood. Policy decisions that are made by relying on such science can also lead to incredible socio-economic inequity and violation of human rights.

### 7.1.2 Data Ethics

Recently, the Data Ethics Commission set up by Germany's Federal Government on 18 July 2018 [167] highlighted 6 general ethical and legal principles from 8 guiding motifs. These general ethical principles are human dignity, self-determination, security, democracy, justice and solidarity, and sustainability. These principles signal the central relevance of the human-centred and value-oriented design of technology. However, finding a balance between potentially divergent principles is a complex matter requiring case-by-case attention.

Consider the example of agriculture, where the EU Code of Conduct for agricultural data sharing emphasizes that “The farmer remains at the heart of the collection, processing, and management of agricultural data” [168]. The Code recognizes that farmers who share data have the right to know and participate in the uses to which their data is put. The challenge in EO and AI4EO research is that farmers' data is remotely collected, often without the knowledge of farmers. If the data collected is used merely to provide information to farmers, there is no ethical concern. However, when the data starts being used for the formulation of policies that (adversely) impact individual farmers' rights, the principles of the EU Code of Conduct will become very relevant. In fact, the use of AI in combination



---

with EO in such instances will also trigger the application of the proposed EU AI Act, which may label such AI use cases as “high risk” and, therefore, subject to extensive regulation.

## 7.2 Ethical Opportunities for Sustainability: Environment and Health

Beyond ethical issues, however, AI4EO research can also contribute to accomplishing various UN Sustainable Development Goals. This can range from assisting in the monitoring the compliance with or deviation from the SDGs or can actively support their accomplishment, e.g. to help expand or protect biodiversity. Recent research suggests that EO and Blockchain technology can be combined to support conservation efforts in Namibia [2]. EO can also help save lives during natural disasters [1], help monitor climate change [3] and sea levels [4].

Similarly, AI4EO applications also offer valuable benefits to public health. Epidemiological studies on the health effects of environmental hazards (such as air pollution and heatwaves) rely on observations from ground-based stations. The spatial distribution of the monitoring network is very heterogeneous and densely located in major cities, making it difficult to accurately represent the human exposure situation across the national territory. Recent AI4EO developments have focused on reconstructing historical air pollution levels at high spatial-temporal resolution using an extensive list of satellite-based data sources. Several single-learner [169–171] and ensemble [172] ML models have also been explored to fill the gaps of air pollution concentrations, thus maximising the human exposure evidence. These spatiotemporal gap-filled environmental datasets are then linked to multiple health outcomes to estimate the associated health-risk exposure to environmental hazards and to quantify the total health burden (e.g., deaths) [173–175].

### Promising directions for future research

- Research on means to conducting more inclusive and multi-disciplinary research in AI4EO, facilitating greater input of ground realities, including socio-cultural, economic and political realities
- Research on means by which multi-cultural approaches can be used to design novel frameworks that facilitate interaction between EO scientists and ethics experts
- Multi-disciplinary research for the creation of novel educational curricula that help EO scientists understand how to apply ethics in their day-to-day research

## Conclusions

The field of AI for EO is rapidly advancing due to the increasing availability of large amounts of sensor data at various resolutions. This perspective paper has provided an overview of the essential scientific tools and approaches supporting the transition from raw EO data to usable information. Since its inception, computer science and signal processing have been necessary in this field. Integrating machine

---

learning, deep learning and computer vision techniques has recently opened up numerous avenues for further, substantial advances in EO, with potential major impacts in downstream applications such as agriculture, food security, or environmental health.

Here, we have discussed current challenges that still prevent the field from reaching full maturity and highlighted promising directions for future research leveraging approaches from computer vision, explainable AI, physics-aware models, user-centric approaches, and advanced processing and computing. Additionally, we have emphasised the importance of carefully considering ethical issues related to the massive use of ML technologies in EO, which is further heightened in the context of large generative ML models for EO. This overview aims to inspire the EO community to achieve the necessary transformative advances in the field and unlock future technical achievements. These are necessary to fulfil the monitor capabilities of using Earth observation to take the pulse of Earth.

## References

1. P. C. Oddo and J. D. Bolton, “The value of near real-time earth observations for improved flood disaster response,” *Frontiers in Environmental Science*, vol. 7, pp. 2–3, 2019.
2. D. Oberhauser, “Blockchain for environmental governance: Can smart contracts reinforce payments for ecosystem services in namibia,” *Frontiers in Blockchain*, vol. 2, 2019.
3. J. Yang, P. Gong, R. Fu, M. Zhang, J. Chen, S. Liang, B. Xu, J. Shi, and R. Dickinson, “The role of satellite remote sensing in climate change studies,” *Nature Climate Change*, vol. 3, no. 10, p. 875–883, 2013.
4. R. Gens, “Remote sensing of coastlines: Detection, extraction and monitoring,” *International Journal of Remote Sensing*, vol. 31, no. 7, p. 1819–1836, 2010.
5. X. X. Zhu, D. Tuia, L. Mou, G.-S. Xia, L. Zhang, F. Xu, and F. Fraundorfer, “Deep learning in remote sensing: A comprehensive review and list of resources,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 5, no. 4, pp. 8–36, 2017.
6. C. Persello, J. D. Wegner, R. Haensch, D. Tuia, P. Ghamisi, M. Koeva, and G. Camps-Valls, “Deep learning and earth observation to support the sustainable development goals: Current approaches, open challenges, and future opportunities,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 10, no. 2, pp. 172–200, 2022.
7. G. Camps-Valls, D. Tuia, X. X. Zhu, and M. Reichstein, *Deep learning for the Earth sciences: A comprehensive approach to remote sensing, climate science and geosciences*. Wiley and Sons, 2021.



- 
8. D. Tuia, R. Roscher, J. D. Wegner, N. Jacobs, X. X. Zhu, and G. Camps-Valls, "Towards a collective agenda on ai for earth science data analysis," *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 2, pp. 88–104, 2021.
  9. M. Schmitt and X. X. Zhu, "Data fusion and remote sensing: An ever-growing relationship," *IEEE Geoscience and Remote Sensing Magazine*, vol. 4, no. 4, pp. 6–23, 2016.
  10. L. Scheibenreif, M. Mommert, and D. Borth, "Contrastive self-supervised data fusion for satellite imagery," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 3, pp. 705–711, 2022.
  11. S. Lefèvre, D. Tuia, J. D. Wegner, T. Produit, and A. S. Nassar, "Toward seamless multiview scene analysis from satellite to street level," *Proceedings of the IEEE*, vol. 105, no. 10, pp. 1884–1899, 2017.
  12. B. Jongman, J. Wagemaker, B. R. Romero, and E. C. De Perez, "Early flood detection for rapid humanitarian response: harnessing near real-time satellite and twitter signals," *ISPRS International Journal of Geo-Information*, vol. 4, no. 4, pp. 2246–2266, 2015.
  13. N. Lang, W. Jetz, K. Schindler, and J. D. Wegner, "A high-resolution canopy height model of the earth," *Nature Ecology & Evolution*, 2023.
  14. F. Waldner, Y. Chen, R. Lawes, and Z. Hochman, "Needle in a haystack: Mapping rare and infrequent crops using satellite imagery and data balancing methods," *Remote Sensing of Environment*, vol. 233, p. 111375, 2019.
  15. L. Fei-Fei, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 4, pp. 594–611, 2006.
  16. G. E. Hinton, T. J. Sejnowski *et al.*, *Unsupervised learning: foundations of neural computation*. MIT press, 1999.
  17. T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in Neural Information Processing Systems 33*, 2020.
  18. T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proceedings of the 37th International conference on machine learning, ICML 2020*. PMLR, 2020, pp. 1597–1607.
  19. X. Chen and K. He, "Exploring simple siamese representation learning," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021*, 2021, pp. 15 750–15 758.
  20. K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *2022 IEEE/CVF Conference on Computer Vision and Pattern recognition*, 2022, pp. 16 000–16 009.

- 
21. G. Tseng, R. Cartuyvels, I. Zvonkov, M. Purohit, D. Rolnick, and H. Kerner, “Lightweight, pre-trained transformers for remote sensing timeseries,” *arXiv preprint arXiv:2304.14065*, 2023.
  22. J. Irvin, L. Tao, J. Zhou, Y. Ma, L. Nashold, B. Liu, and A. Y. Ng, “Usat: A unified self-supervised encoder for multi-sensor satellite imagery,” *arXiv preprint arXiv:2312.02199*, 2023.
  23. X. Sun, P. Wang, W. Lu, Z. Zhu, X. Lu, Q. He, J. Li, X. Rong, Z. Yang, H. Chang *et al.*, “Ringmo: A remote sensing foundation model with masked image modeling,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, 2022.
  24. J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 6840–6851.
  25. R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022*, pp. 10 684–10 695.
  26. R. Bachmann, D. Mizrahi, A. Atanov, and A. Zamir, “MultiMAE: Multi-modal multi-task masked autoencoders,” in *European Conference on Computer Vision, 2022*, pp. 348–367.
  27. A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, “Segment anything,” in *IEEE/CVF International Conference on Computer Vision, 2023*, pp. 4015–4026.
  28. J. Jakubik, S. Roy, C. E. Phillips, P. Fraccaro, D. Godwin, B. Zadrozny, D. Szwarcman, C. Gomes, G. Nyirjesy, B. Edwards, D. Kimura, N. Simumba, L. Chu, S. K. Mukkavilli, D. Lambhate, K. Das, R. Bangalore, D. Oliveira, M. Muszynski, K. Ankur, M. Ramasubramanian, I. Gurung, S. Khallaghi, H. S. Li, M. Cecil, M. Ahmadi, F. Kordi, H. Alemohammad, M. Maskey, R. Ganti, K. Weldemariam, and R. Ramachandran, “Foundation Models for Generalist Geospatial Artificial Intelligence,” *arXiv preprint 2310.18660*, 2023.
  29. Y. Cong, S. Khanna, C. Meng, P. Liu, E. Rozi, Y. He, M. Burke, D. B. Lobell, and S. Ermon, “Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery,” in *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, 2022*.
  30. J. Xu and Z. Zhu, “Reinforced continual learning,” in *Advances in Neural Information Processing Systems, NeurIPS, 2018*.
  31. O. Tasar, Y. Tarabalka, and P. Alliez, “Continual learning for dense labeling of satellite images,” in *International Geoscience and Remote Sensing Symposium, IGARSS, 2019*.
  32. N. Ammour, Y. Bazi, H. Alhichri, and N. Alajlan, “Continual learning approach for remote sensing scene classification,” *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022.

- 
33. G. Lenczner, A. Chan-Hon-Tong, N. Luminari, and B. Le Saux, “Weakly-supervised continual learning for class-incremental segmentation,” *CoRR*, vol. abs/2201.01029, 2022.
  34. S. Srivastava, M. Berman, M. Blaschko, and D. Tuia, “Adaptive compression-based lifelong learning,” in *British Machine Vision Conference, BMVC*, 2019.
  35. M. Russwurm, S. Wang, B. Kellenberger, R. Roscher, and D. Tuia., “Meta-learning to address diverse earth observation problems across resolutions,” *Nature Communications Earth & Environment*, vol. 5, no. 37, 2024.
  36. W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang, “A closer look at few-shot classification,” in *International Conference on Learning Representations, ICLR*, 2019.
  37. B. Sun and K. Saenko, “Deep coral: Correlation alignment for deep domain adaptation,” in *European Conference on Computer Vision, ECCV Workshops*, 2016.
  38. M. Xu, M. Wu, K. Chen, C. Zhang, and J. Guo, “The eyes of the gods: A survey of unsupervised domain adaptation methods based on remote sensing data,” *Remote Sensing*, vol. 14, no. 17, p. 4380, 2022.
  39. F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, “Learning to compare: Relation network for few-shot learning,” in *Computer Vision and Pattern Recognition, CVPR*, 2018.
  40. Y. Xian, B. Schiele, and Z. Akata, “Zero-shot learning – the good, the bad and the ugly,” in *Computer Vision and Pattern Recognition, CVPR*, 2017.
  41. P. Brazdil, J. N. van Rijn, C. Soares, and J. Vanschoren, *Metalearning: Applications to Automated Machine Learning and Data Mining*, 2nd ed. Springer, 2022.
  42. M. Huisman, J. N. van Rijn, and A. Plaat, “A survey of deep meta-learning,” *Artificial Intelligence Review*, vol. 54, no. 6, pp. 4483–4541, 2021.
  43. C. White, M. Safari, R. Sukthanker, B. Ru, T. Elsken, A. Zela, D. Dey, and F. Hutter, “Neural architecture search: Insights from 1000 papers,” *arXiv preprint [arXiv:2301.08727](https://arxiv.org/abs/2301.08727)*, 2023.
  44. F. Hutter, L. Kotthoff, and J. Vanschoren, *Automated machine learning: methods, systems, challenges*. Springer, 2019.
  45. C. Ying, A. Klein, E. Christiansen, E. Real, K. Murphy, and F. Hutter, “Nas-bench-101: Towards reproducible neural architecture search,” in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, ser. Proceedings of Machine Learning Research, vol. 97, 2019, pp. 7105–7114.

- 
46. G. Sumbul, M. Charfuelan, B. Demir, and V. Markl, “Bigearthnet: A large-scale benchmark archive for remote sensing image understanding,” in *International Geoscience and Remote Sensing Symposium, IGARSS*, 2019.
  47. F. Hutter, H. H. Hoos, and K. Leyton-Brown, “An efficient approach for assessing hyperparameter importance,” in *Proceedings of the 31th International Conference on Machine Learning, ICML 2014*, 2014, pp. 754–762.
  48. V. Perrone, H. Shen, M. W. Seeger, C. Archambeau, and R. Jenatton, “Learning search spaces for bayesian optimization: Another view of hyperparameter transfer learning,” in *Advances in Neural Information Processing Systems 32*, 2019, pp. 12 751–12 761.
  49. J. N. van Rijn and F. Hutter, “Hyperparameter importance across datasets,” in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018*. ACM, 2018, pp. 2367–2376.
  50. F. Mohr and J. N. van Rijn, “Learning curves for decision making in supervised machine learning – a survey,” *CoRR*, vol. abs/2201.12150, 2022.
  51. —, “Fast and informative model selection using learning curve cross-validation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 8, pp. 9669–9680, 2023.
  52. M. König, H. H. Hoos, and J. N. van Rijn, “Accelerating adversarially robust model selection for deep neural networks via racing,” in *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-24)*. AAAI Press, 2024, pp. 21 267–21 275.
  53. S. H. Silva and P. Najafirad, “Opportunities and challenges in deep learning adversarial robustness: A survey,” *arXiv preprint [arXiv:2007.00753](https://arxiv.org/abs/2007.00753)*, 2020.
  54. N. R. Palacios Salinas, M. Baratchi, J. N. van Rijn, and A. Vollrath, “Automated machine learning for satellite data: integrating remote sensing pre-trained models into automl systems,” in *Machine Learning and Knowledge Discovery in Databases*, ser. Lecture Notes in Computer Science, vol. 12979. Springer, 2021, pp. 447–462.
  55. J. Wasala, S. Marselis, L. Arp, H. Hoos, N. Longép e, and M. Baratchi, “Autosr4eo: An automl approach to super-resolution for earth observation images,” *Remote Sensing*, vol. 16, no. 3, 2024.
  56. M. Gašparović, D. Medak, I. Pilaš, L. Jurjević, and I. Balenović, “Fusion of Sentinel-2 and Planetscope imagery for vegetation detection and monitoring,” *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences*, vol. 42, no. 1, pp. 155–160, 2018.
  57. F. Chollet, “On the measure of intelligence,” *arXiv preprint [arXiv:1911.01547](https://arxiv.org/abs/1911.01547)*, 2019.

- 
58. M. Sudmanns, D. Tiede, S. Lang, H. Bergstedt, G. Trost, H. Augustin, A. Baraldi, and T. Blaschke, “Big earth data: disruptive changes in earth observation data management and analysis?” *International Journal of Digital Earth*, vol. 13, no. 7, pp. 832–850, 2020.
  59. A. Wall, B. Deiseroth, E. Tzirita Zacharitou, J.-A. Quiané-Ruiz, B. Demir, and V. Markl, “Agora-eo: A unified ecosystem for earth observation - a vision for boosting eo data literacy,” *Big Data from Space Conference*, 2021.
  60. J. Traub, Z. Kaoudi, J.-A. Quiané-Ruiz, and V. Markl, “Agora: Bringing together datasets, algorithms, models and more in a unified ecosystem [vision],” *SIGMOD Record*, vol. 49, no. 4, p. 6–11, 2021.
  61. J. Konečný, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon, “Federated learning: strategies for improving communication efficiency,” in *Neural Information Processing Systems, NeurIPS workshops*, 2016. [Online]. Available: <https://arxiv.org/abs/1610.05492>
  62. “Federated learning: Collaborative machine learning without centralized training data,” Google AI Blog: <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>, accessed: 2022-02-28.
  63. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep nnetworks from decentralized data,” in *International Conference on Artificial Intelligence and Statistics, AISTATS*, vol. 54, 2017, pp. 1273–1282.
  64. Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, “Federated learning with non-iid data,” *arXiv preprint arXiv:1806.00582*, 2018.
  65. B. Büyüktaş, G. Sumbul, and B. Demir, “Learning across decentralized multi-modal remote sensing archives with federated learning,” in *IEEE International Geoscience and Remote Sensing Symposium, IGARSS 2023*. IEEE, 2023, pp. 4966–4969.
  66. “Hub: A common database for ai,” <https://www.activeloop.ai/>, accessed: 2022-02-28.
  67. “Open neural network exchange,” <https://onnx.ai/>, accessed: 2022-02-28.
  68. J. Vanschoren, J. N. van Rijn, B. Bischl, and L. Torgo, “Openml: networked science in machine learning,” *SIGKDD Explorations*, vol. 15, pp. 49–60, 2014.
  69. M. Ziaja, P. Bosowski, M. Myller, G. Gajoch, M. Gumiela, J. Protich, K. Borda, D. Jayaraman, R. Dividino, and J. Nalepa, “Benchmarking deep learning for on-board space applications,” *Remote Sensing*, vol. 13, no. 19, 2021.
  70. C. Gernigon, S.-I. Filip, O. Sentieys, C. Coggiola, and M. Bruno, “Low-precision floating-point for efficient on-board deep neural network processing,” in *2023 European Data Handling & Data Processing Conference (EDHPC)*. IEEE, 2023, pp. 1–8.

- 
71. J. Murphy, J. E. Ward, and B. M. Namee, “Machine learning in space: A review of machine learning algorithms and hardware for space applications,” in *The 29th Irish Conference on Artificial Intelligence and Cognitive Science 2021, Dublin, Republic of Ireland, December 9-10, 2021*, ser. CEUR Workshop Proceedings, vol. 3105. CEUR-WS.org, 2021, pp. 72–83.
  72. G. Guerrisi, F. D. Frate, and G. Schiavon, “Artificial intelligence based on-board image compression for the  $\Phi$ -sat-2 mission,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 16, pp. 8063–8075, 2023.
  73. C. Li, G. Wang, B. Wang, X. Liang, Z. Li, and X. Chang, “Dynamic slimmable network,” in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 8607–8617.
  74. C. Page, K. Cahoy, and E. Gizzi, “Developing intelligent space systems: A survey & rubric for future missions,” in *37th Annual Small Satellite Conference*, 2023.
  75. D. Arvor, M. Belgiu, Z. Falomir, I. Mougenot, and L. Durieux, “Ontologies to interpret remote sensing images: why do we need them?” *GIScience & remote sensing*, vol. 56, no. 6, pp. 911–939, 2019.
  76. G. Ras, N. Xie, M. Van Gerven, and D. Doran, “Explainable deep learning: A field guide for the uninitiated,” *Journal of Artificial Intelligence Research*, vol. 73, pp. 329–396, 2022.
  77. C. Rudin, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nature machine intelligence*, vol. 1, no. 5, pp. 206–215, 2019.
  78. J. Pearl, *Causality: Models, Reasoning and Inference*, 2nd ed. Cambridge University Press, 2009.
  79. G. Camps-Valls, L. Martino, D. H. Svendsen, M. Campos-Taberner, J. Muñoz-Marí, V. Laparra, D. Luengo, and F. J. García-Haro, “Physics-aware gaussian processes in remote sensing,” *Applied Soft Computing*, vol. 68, pp. 69–82, 2018.
  80. M. Reichstein, G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and Prabhat, “Deep learning and process understanding for data-driven earth system science,” *Nature*, vol. 566, pp. 195–204, 2019.
  81. G. Camps-Valls, A. Gerhardus, U. Ninad, G. Varando, G. Martius, E. Balaguer-Ballester, R. Vinuesa, E. Diaz, L. Zanna, and J. Runge, “Discovering causal relations and equations from data,” *Physics Reports*, vol. 1044, pp. 1–68, 2023.
  82. P. Jackson, *Introduction to Expert Systems*, 3rd ed. Addison-Wesley, 1999.
  83. M. Bohanec, “Dex (decision expert): A qualitative hierarchical multi-criteria method,” in *Multiple Criteria Decision Making*. Springer Singapore, 2022, pp. 39–78.

- 
84. M. Debeljak, A. Trajanov, V. Kuzmanovski, J. Schröder, T. Sandén, H. Spiegel, D. P. Wall, M. V. de Broek, M. Rutgers, F. Bampa, R. E. Creamer, and C. B. Henriksen, “A field-scale decision support system for assessment and management of soil functions,” *Frontiers in Environmental Science*, vol. 7, p. 115, 2019.
  85. “Earth observation for decision making,” <http://www.oecd.org/environment/indicators-modelling-outlooks/earth-observation-for-decision-making.htm>, accessed: 2022-03-22.
  86. R. Vinuesa, H. Azizpour, I. Leite, M. Balaam, V. Dignum, S. Domisch, A. Felländer, S. D. Langhans, M. Tegmark, and F. F. Nerini, “The role of artificial intelligence in achieving the sustainable development goals,” *Nature Communications*, vol. 11, no. 1, pp. 1–10, 2020.
  87. T. M. Mitchell, R. M. Keller, and S. T. Kedar-Cabelli, “Explanation-based generalization: A unifying view,” *Machine Learning*, vol. 1, no. 1, pp. 47–80, 1986.
  88. N. Lavrač and S. Džeroski, *Inductive logic programming. Techniques and applications*, ser. Ellis Horwood series in artificial intelligence. Ellis Horwood, 1994.
  89. S. Džeroski and N. Lavrač, *Relational data mining*. Springer, 2001.
  90. L. Getoor and B. Taskar, *Introduction to Statistical Relational Learning*. The MIT Press, 2007.
  91. L. D. Raedt, K. Kersting, S. Natarajan, and D. Poole, “Statistical relational artificial intelligence: Logic, probability, and computation,” *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 10, no. 2, pp. 1–189, 2016.
  92. C. N. Silla and A. A. Freitas, “A survey of hierarchical classification across different application domains,” *Data Mining and Knowledge Discovery*, vol. 22, no. 1-2, pp. 31–72, 2010.
  93. H. Yang, S. Li, J. Chen, X. Zhang, and S. Xu, “The standardization and harmonization of land cover classification systems towards harmonized datasets: A review,” *ISPRS International Journal of Geo-Information*, vol. 6, no. 5, p. 154, 2017.
  94. R. Stivaktakis, G. Tsagkatakis, and P. Tsakalides, “Deep learning for multilabel land cover scene categorization using data augmentation,” *IEEE Geoscience and Remote Sensing Letters*, vol. 16, no. 7, pp. 1031–1035, 2019.
  95. G. Sumbul and B. Demir, “A deep multi-attention driven approach for multi-label remote sensing image classification,” *IEEE Access*, vol. 8, pp. 95 934–95 946, 2020.
  96. L. Möllenbrok, G. Sumbul, and B. Demir, “Deep active learning for multi-label classification of remote sensing images,” *IEEE Geoscience and Remote Sensing Letters*, 2023.



- 
97. D. Tuia, J. Muñoz-Marí, M. Kanevski, and G. Camps-Valls, “Structured output SVM for remote sensing image classification,” *Journal of Signal Processing Systems*, vol. 65, no. 3, pp. 457–468, 2011.
  98. M. O. Turkoglu, S. D’Aronco, G. Perich, F. Liebisch, C. Streit, K. Schindler, and J. D. Wegner, “Crop mapping from image time series: deep learning with multi-scale label hierarchies,” *Remote Sensing of Environment*, vol. 264, p. 112603, 2021.
  99. S. Džeroski, B. Goethals, and P. Panov, Eds., *Inductive Databases and Constraint-Based Data Mining*. Springer New York, 2010.
  100. S. Džeroski and L. Todorovski, Eds., *Computational Discovery of Scientific Knowledge*. Springer Berlin Heidelberg, 2007.
  101. N. Simidjievski, L. Todorovski, J. Kocijan, and S. Džeroski, “Equation discovery for nonlinear system identification,” *IEEE Access*, vol. 8, pp. 29 930–29 943, 2020.
  102. M. D. Wilkinson, “The fair guiding principles for scientific data management and stewardship,” *Scientific Data*, vol. 3, p. 160018, 2016.
  103. G. Giuliani, H. Cazeaux, P.-Y. Burgi, C. Poussin, J.-P. Richard, and B. Chatenoux, “Swissenvéo: A fair national environmental data repository for earth observation open science,” *Data Science Journal*, vol. 20, 2021.
  104. I. Grujdin and M. Datcu, “Self-learning ontology for natural hazards,” in *IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023, pp. 2227–2230.
  105. J. Runge, S. Bathiany, E. Bollt, G. Camps-Valls, D. Coumou, E. Deyle, C. Glymour, M. Kretschmer, M. Mahecha, J. Muñoz-Marí, E. van Nes, J. Peters, R. Quax, M. Reichstein, M. Scheffer, B. Schölkopf, P. Spirtes, G. Sugihara, J. Sun, K. Zhang, and J. Zscheischler, “Inferring causation from time series with perspectives in earth system sciences,” *Nature Communications*, vol. 10, no. 2553, 2019.
  106. J. Runge, A. Gerhardus, G. Varando, V. Eyring, and G. Camps-Valls, “Causal inference for time series,” *Nature Reviews Earth & Environment*, vol. 10, p. 2553, 2023.
  107. W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Müller, “Toward interpretable machine learning: Transparent deep neural networks and beyond,” *arXiv preprint [arXiv:2003.07631](https://arxiv.org/abs/2003.07631)*, 2020.
  108. R. Roscher, B. Bohn, M. F. Duarte, and J. Garcke, “Explainable machine learning for scientific insights and discoveries,” *IEEE Access*, vol. 8, pp. 42 200–42 216, 2020.
  109. A. Wolanin, G. Mateo-Garcia, G. Camps-Valls, L. Gomez-Chova, M. Meroni, G. Duveiller, Y. Liangzhi, and L. Guanter, “Estimating and understanding crop yields with explainable deep



- 
- learning in the indian wheat belt,” *Environmental Research Letters*, vol. 15, no. 2, pp. 1–12, 2020.
110. J. E. Johnson, V. Laparra, A. Pérez-Suay, M. Mahecha, and G. Camps-Valls, “Kernel methods and their derivatives: Concept and perspectives for the earth system sciences,” *PLOS One*, vol. 15, no. 10, p. p.e0235885, 2020.
111. D. Marcos, S. Lobry, and D. Tuia, “Semantically interpretable activation maps: What-where-how explanations within cnns,” in *International Conference on Computer Vision Workshops, ICCV Workshops*, 2019, pp. 4207–4215.
112. M. Rußwurm and M. Körner, “Self-attention for raw optical satellite time series classification,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 169, pp. 421–435, 2020.
113. A. Levering, D. Marcos, and D. Tuia, “On the relation between landscape beauty and land cover: A case study in the uk at sentinel-2 resolution with interpretable ai,” *ISPRS journal of Photogrammetry and Remote Sensing*, vol. 177, pp. 194–203, 2021.
114. T.-A. Nguyen, B. Kellenberger, and D. Tuia, “Mapping forest in the swiss alps treeline ecotone with explainable deep learning,” *Remote Sensing of Environment*, vol. 281, p. 113217, 2022.
115. T. Stomberg, I. Weber, M. Schmitt, and R. Roscher, “Jungle-net: Using explainable machine learning to gain new insights into the appearance of wilderness in satellite imagery,” *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 3, pp. 317–324, 2021.
116. E. Diaz, J. E. Adsuara, A. Moreno-Martinez, M. Piles, and G. Camps-Valls, “Inferring causal relations from observational long-term carbon and water fluxes records,” *Scientific Reports*, vol. 12, p. 1610, 2022.
117. J. Peters, D. Janzing, and B. Schölkopf, *Elements of Causal Inference - Foundations and Learning Algorithms*, ser. Adaptive Computation and Machine Learning Series. MIT Press, 2017.
118. G. Varando, M.-A. Fernández-Torres, and G. Camps-Valls, “Learning granger causal feature representations,” in *International Conference on Machine Learning, ICML Workshops*, 2021.
119. D. H. Svendsen, L. Martino, M. Campos-Taberner, F. J. García-Haro, and G. Camps-Valls, “Joint gaussian processes for biophysical parameter retrieval,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 3, pp. 1718–1727, 2017.
120. W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K.-R. Müller, *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, ser. Lecture Notes in Computer Science. Springer, 2019, vol. 11700.

- 
121. L. Von Rueden, S. Mayer, J. Garcke, C. Bauckhage, and J. Schuecker, “Informed machine learning—towards a taxonomy of explicit integration of knowledge into machine learning,” *Learning*, vol. 18, pp. 19–20, 2019.
  122. G. Camps-Valls, J. Verrelst, J. Muñoz-Marí, V. Laparra, F. Mateo-Jiménez, and J. Gomez-Dans, “A survey on gaussian processes for earth-observation data analysis: A comprehensive investigation,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 4, no. 2, pp. 58–78, 2016.
  123. G. Camps-Valls, D. Sejdinovic, J. Runge, and M. Reichstein, “A perspective on gaussian processes for earth observation,” *National Science Review*, vol. 6, no. 4, pp. 616–618, 2019.
  124. F. Del Frate and D. Solimini, “On neural network algorithms for retrieving forest biomass from SAR data,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 1, pp. 24–34, 2004.
  125. P. Sellitto, F. D. Frate, D. Solimini, and S. Casadio, “Tropospheric ozone column retrieval from esa-envisat sciamachy nadir uv/vis radiance measurements by means of a neural network algorithm,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50, pp. 998–1011, 2012.
  126. M. Datcu, Z. Huang, A. Anghel, J. Zhao, and R. Cacoveanu, “Explainable, physics-aware, trustworthy artificial intelligence: A paradigm shift for synthetic aperture radar,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 11, no. 1, pp. 8–25, 2023.
  127. K. Kashinath, A. Albert, R. Wang, M. Mustafa, and R. Yu, “Physics-informed spatio-temporal deep learning models,” *Bulletin of the American Physical Society*, vol. 64, 2019.
  128. J. Wu, K. Kashinath, A. Albert, M. Prabhat, and H. Xiao, “Physics-informed generative learning to emulate unresolved physics in climate models,” in *American Geoscience Union Fall Meeting*, 2018.
  129. G. Camps-Valls, M. Reichstein, X. Zhu, and D. Tuia, “Advancing Deep Learning for Earth Sciences: From Hybrid Modeling to Interpretability,” in *International Geoscience and Remote Sensing Symposium*, 2020.
  130. A. Halevy, P. Norvig, and F. Pereira, “The unreasonable effectiveness of data,” *IEEE Intelligent Systems*, vol. 24, no. 2, pp. 8–12, 2009.
  131. Á. Moreno-Martínez, G. Camps-Valls, J. Kattge, N. Robinson, M. Reichstein, P. van Bodegom, K. Kramer, J. H. C. Cornelissen, P. Reich, M. Bahn *et al.*, “A methodology to derive global maps of leaf traits using remote sensing and climate data,” *Remote Sensing of Environment*, vol. 218, pp. 69–88, 2018.
  132. D. H. Svendsen, L. Martino, and G. Camps-Valls, “Active emulation of computer codes with gaussian processes – application to remote sensing,” *Pattern Recognition*, vol. 100, no. 107103, pp. 1–12, 2020.

- 
133. H. Ye, R. J. Beamish, S. M. Glaser, S. C. Grant, C.-h. Hsieh, L. J. Richards, J. T. Schnute, and G. Sugihara, "Equation-free mechanistic ecosystem forecasting using empirical dynamic modeling," *Proceedings of the National Academy of Sciences*, vol. 112, no. 13, pp. E1569–e1576, 2015.
  134. B. C. Daniels and I. Nemenman, "Automated adaptive inference of phenomenological dynamical models," *Nature Communications*, vol. 6, p. 8133, 2015.
  135. S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proceedings of the National Academy of Sciences*, vol. 113, no. 15, pp. 3932–3937, 2016.
  136. L. Filchev, L. Pashova, V. Kolev, and S. Frye, "Challenges and solutions for utilizing earth observations in the "big data" era," in *BigSkyEarth Conference*, 2018.
  137. D. Tuia, F. Ratle, F. Pacifici, M. Kanevski, and W. J. Emery, "Active learning methods for remote sensing image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47(7), pp. 2218–2232, 2009.
  138. B. Demir, L. Minello, and L. Bruzzone, "Definition of effective training sets for supervised classification of remote sensing images by a novel cost-sensitive active learning method," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 2, pp. 1272–1284, 2013.
  139. D. Tuia and J. Muñoz-Marí, "Learning user's confidence for active learning," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 2, pp. 872–880, 2013.
  140. A. C. Rodríguez, S. D'Aronco, K. Schindler, and J. D. Wegner, "Mapping oil palm density at country scale: An active learning approach," *Remote Sensing of Environment*, vol. 261, p. 112479, 2021.
  141. P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, B. B. Gupta, X. Chen, and X. Wang, "A survey of deep active learning," *ACM Computing Surveys*, vol. 54, no. 9, 2021.
  142. L. Bashmal, Y. Bazi, F. Melgani, M. M. Al Rahhal, and M. A. Al Zuair, "Language integration in remote sensing: Tasks, datasets, and future directions," *IEEE Geoscience and Remote Sensing Magazine*, vol. 11, no. 4, pp. 63 – 93, 2023.
  143. S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh, "VQA: visual question answering," in *International Conference on Computer Vision, ICCV*, 2015, pp. 2425–2433.
  144. S. Lobry, D. Marcos, J. Murray, and D. Tuia, "Rsvqa: visual question answering for remote sensing data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 12, pp. 8555–8566, 2020.

- 
145. Z. Yuan, L. Mou, Z. Xiong, and X. X. Zhu, “Change detection meets visual question answering,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–13, 2022.
  146. C. Wen, Y. Hu, X. Li, Z. Yuan, and X. X. Zhu, “Vision-language models in remote sensing: Current progress and future trends,” *arXiv.org preprint arXiv:2305.05726*, 2023.
  147. C. Chappuis, V. Zermatten, S. Lobry, B. Le Saux, and D. Tuia, “Prompt-rsvqa: Prompting visual context to a language model for remote sensing visual question answering,” in *Computer Vision and Pattern Recognition, CVPR Workshops*, 2022.
  148. G. Sumbul, S. Nayak, and B. Demir, “Sd-rsic: Summarization-driven deep remote sensing image captioning,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 8, pp. 6922–6934, 2020.
  149. H. Kandala, S. Saha, B. Banerjee, and X. X. Zhu, “Exploring transformer and multilabel classification for remote sensing image captioning,” *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022.
  150. A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
  151. U. Mall, C. P. Phoo, M. K. Liu, C. Vondrick, B. Hariharan, and K. Bala, “Remote sensing vision-language foundation models without annotations via ground remote alignment,” *arXiv preprint arXiv:2312.06960*, 2023.
  152. V. Zermatten, J. C. Navarro, L. Hughes, and D. Tuia, “Text as a richer source of supervision in semantic segmentation tasks,” in *International Geoscience and Remote Sensing Symposium, IGARSS*, 2023.
  153. Z. Yuan, L. Mou, Y. Hua, and X. X. Zhu, “Rrsis: Referring remote sensing image segmentation,” *arXiv.org preprint arXiv:2306.08625*, 2023.
  154. J. Pennington, R. Socher, and C. Manning, “GloVe: Global vectors for word representation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014*. ACL, 2014, pp. 1532–1543.
  155. D. Sacha, M. Sedlmair, L. Zhang, J. A. Lee, D. Weiskopf, S. North, and D. Keim, “Human-centered machine learning through interactive visualization,” *Neurocomputing*, vol. 268, pp. 164–175, 2017.
  156. M. Babaee, S. Tsoukalas, G. Rigoll, and M. Datcu, “Immersive visualization of visual data using nonnegative matrix factorization,” *Neurocomputing*, vol. 173, pp. 245–255, 2016.
  157. L. Mi, S. Li, C. Chappuis, and D. Tuia, “Knowledge-aware cross-modal text-image retrieval for remote sensing images,” in *Proceedings of the Second Workshop on Complex Data Challenges*

---

*in Earth Observation (CDCEO 2022) co-located with 31st International Joint Conference on Artificial Intelligence and the 25th European Conference on Artificial Intelligence (IJCAI-ECAI 2022)*, ser. CEUR Workshop Proceedings, vol. 3207. CEUR-WS.org, 2022, pp. 14–20.

158. A. Jobin, M. Ienca, and E. Vayena, “The global landscape of ai ethics guidelines,” *Nature Machine Intelligence*, vol. 1, p. 389–399, 2019.
159. M. Kochupillai, M. Kahl, M. Schmitt, H. Taubenböck, and X. X. Zhu, “Earth observation and artificial intelligence: Understanding emerging ethical issues and opportunities,” *IEEE Geoscience and Remote Sensing Magazine*, 2022.
160. G. Cheng, X. Sun, K. Li, L. Guo, and J. Han, “Perturbation-seeking generative adversarial networks: A defense framework for remote sensing image scene classification,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–11, 2022.
161. X. Sun, G. Cheng, L. Pei, H. Li, and J. Han, “Threatening patch attacks on object detection in optical remote sensing images,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–10, 2023.
162. E. Commission, “Ethics guidelines for trustworthy ai: High-level expert group on artificial intelligence,” 2018. [Online]. Available: <https://wayback.archive-it.org/12090/20201227221227/https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>
163. T. Hagendorff, “The ethics of ai ethics: An evaluation of guidelines,” *Minds and machines*, vol. 30, no. 1, pp. 99–120, 2020.
164. A. J. Thomas, “Deontology, consequentialism and moral realism.” *Minerva: An Internet Journal of Philosophy*, vol. 19, 2015.
165. M. Kochupillai, “Outline of a novel approach for identifying ethical issues in early stages of ai4eo research,” *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, pp. 165–1168, 2021.
166. K. Bugbee and R. Ramachandran, “The ethics of large language models: Who controls the future of open science?” <https://www.earthdata.nasa.gov/learn/blog/ethics-large-language-models>, 2023.
167. B. d. I. Daten Ethik Kommission, “Opinion of the data ethics commission - executive summary,” [https://www.bfdi.bund.de/SharedDocs/Downloads/EN/Datenschutz/Data-Ethics-Commission.Opinion.pdf?\\_\\_blob=publicationFile&v=1](https://www.bfdi.bund.de/SharedDocs/Downloads/EN/Datenschutz/Data-Ethics-Commission.Opinion.pdf?__blob=publicationFile&v=1), p. 32, 2018.
168. COPA, COGECA, CEMA, F. Europe, CEETAR, ECPA, EFFAB, FEFAC, and ESA, “Eu code of conduct on agricultural data sharing by contractual agreement,” [https://www.copa-cogeca.eu/img/user/files/EU%20CODE/EU\\_Code\\_2018\\_web\\_version.pdf](https://www.copa-cogeca.eu/img/user/files/EU%20CODE/EU_Code_2018_web_version.pdf), 2018.

- 
169. R. Schneider, A. M. Vicedo-Cabrera, F. Sera, P. Masselot, M. Stafoggia, K. de Hoogh, I. Kloog, S. Reis, M. Vieno, and A. Gasparrini, “A satellite-based spatio-temporal machine learning model to reconstruct daily pm2.5 concentrations across great britain,” *Remote Sensing*, vol. 12, no. 22, 2020.
170. M. Stafoggia, C. Johansson, P. Glantz, M. Renzi, A. Shtein, K. de Hoogh, I. Kloog, M. Davoli, P. Michelozzi, and T. Bellander, “A random forest approach to estimate daily particulate matter, nitrogen dioxide, and ozone at fine spatial resolution in sweden,” *Atmosphere*, vol. 11, no. 3, 2020.
171. J. Chen, K. de Hoogh, J. Gulliver, B. Hoffmann, O. Hertel, M. Ketzel, M. Bauwelinck, A. van Donkelaar, U. A. Hvidtfeldt, K. Katsouyanni, N. A. Janssen, R. V. Martin, E. Samoli, P. E. Schwartz, M. Stafoggia, T. Bellander, M. Strak, K. Wolf, D. Vienneau, R. Vermeulen, B. Brunekreef, and G. Hoek, “A comparison of linear regression, regularization, and machine learning algorithms to develop europe-wide spatial models of fine particles and nitrogen dioxide,” *Environment International*, vol. 130, p. 104934, 2019.
172. Q. Di, H. Amini, L. Shi, I. Kloog, R. Silvern, J. Kelly, M. B. Sabath, C. Choirat, P. Koutrakis, A. Lyapustin, Y. Wang, L. J. Mickley, and J. Schwartz, “An ensemble-based model of pm2.5 concentration across the contiguous united states with high spatiotemporal resolution,” *Environment International*, vol. 130, p. 104909, 2019.
173. R. Schneider, P. Masselot, A. M. Vicedo-Cabrera, F. Sera, M. Blangiardo, C. Forlani, J. Douros, O. Jorba, M. Adani, R. Kouznetsov, F. Couvidat, J. Arteta, B. Raux, M. Guevara, A. Colette, J. Barré, V.-H. Peuch, and A. Gasparrini, “Differential impact of government lockdown policies on reducing air pollution levels and related mortality in europe,” *Scientific Reports*, vol. 12, no. 1, p. 726, 2022.
174. A. M. Vicedo-Cabrera, N. Scovronick, F. Sera, D. Royé, R. Schneider, A. Tobias, C. Astrom, Y. Guo, Y. Honda, D. Hondula *et al.*, “The burden of heat-related mortality attributable to recent human-induced climate change,” *Nature climate change*, vol. 11, no. 6, pp. 492–500, 2021.
175. A. Urban, C. Di Napoli, H. L. Cloke, J. Kyselý, F. Pappenberger, F. Sera, R. Schneider, A. M. Vicedo-Cabrera, F. Acquaotta, M. S. Ragettli, C. I. niguez, A. Tobias, E. Indermitte, H. Orru, J. J. Jaakkola, N. R. Rytí, M. Pascal, V. Huber, A. Schneider, F. de’ Donato, P. Michelozzi, and A. Gasparrini, “Evaluation of the era5 reanalysis-based universal thermal climate index on mortality data in europe,” *Environmental Research*, vol. 198, p. 111227, 2021.

---

## Contributions statement

- DT coordinated the writing group.
- Section contributions (in **bold**, lead writer): Introduction: **DT**; Computer vision and machine learning: **KS**, **BD**, DT, JvR, HH; Advanced processing and computing: **BD**, VM, BLS, RS, FDF; Knowledge-based AI: **SD**, JvR, HH; Explainable AI and causality: **GCV**; Physics-aware ML: **FDF**, **GCV**, KS; User-centric: **DT**, XZ, MD; Ethics: **MK**.

ESA has partially funded the authors to work on this manuscript through the ESA visiting professors program.