

TerraMind: Large-Scale Generative Multimodality for Earth Observation

Johannes Jakubik^{1,*} Felix Yang^{1,2*} Benedikt Blumenstiel^{1,*} Erik Scheurer³
 Rocco Sedona³ Stefano Maurogiovanni^{3,6} Jente Bosmans⁴ Nikolaos Dionelis⁴
 Valerio Marsocci⁴ Niklas Kopp¹ Rahul Ramachandran⁵ Paolo Fraccaro^{1,†}
 Thomas Brunschwiler^{1,†} Gabriele Cavallaro^{3,6,†} Juan Bernabe-Moreno^{1,†} Nicolas Longépé^{4,†}

¹IBM Research – Europe
²ETH Zurich
³Forschungszentrum Jülich
⁴European Space Agency Φ-Lab
⁵NASA IMPACT
⁶University of Iceland
 johannes.jakubik1@ibm.com

Accepted at ICCV'25

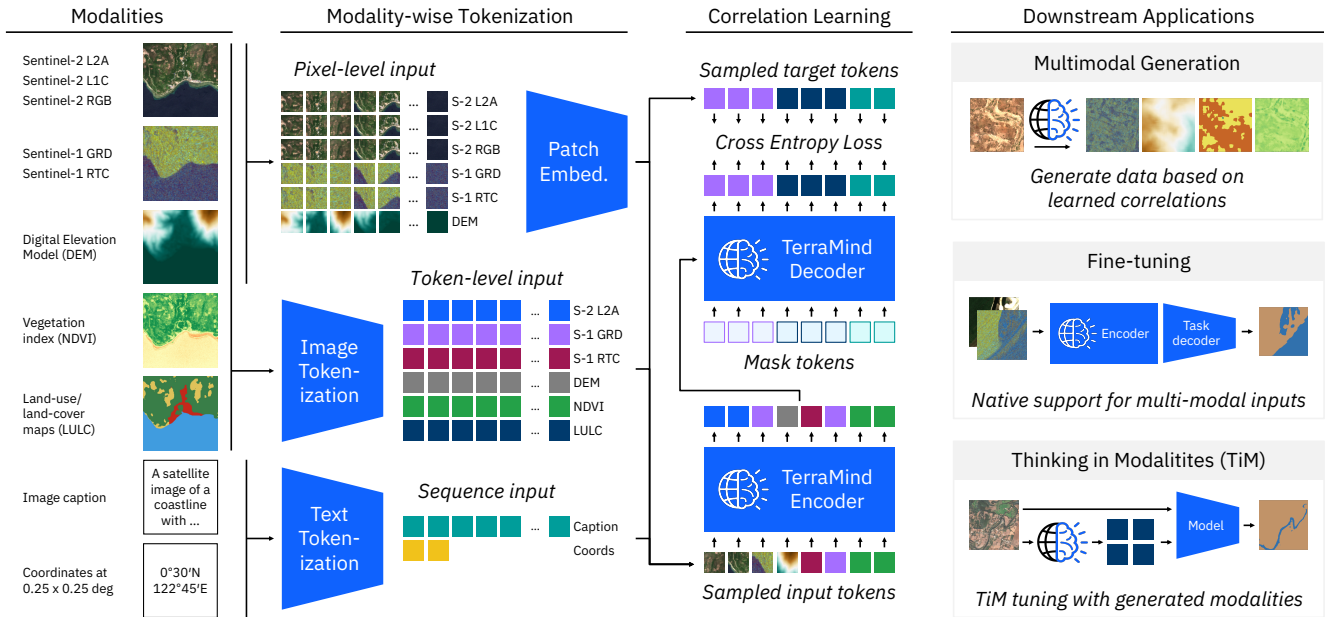


Figure 1. TerraMind represents the first any-to-any generative, and large-scale multimodal model for Earth observation pre-trained on 500 billion tokens from global geospatial data. The model digests multi-scale representations at pixel-level and token-level simultaneously. TerraMindv1 unlocks (i) generation, (ii) zero-shot and finetuning applications, and (iii) "Thinking-in-Modalities" finetuning and inference.

Abstract

We present TerraMind, the first any-to-any generative, multimodal foundation model for Earth observation (EO). Un-

like other multimodal models, TerraMind is pretrained on dual-scale representations combining both token-level and pixel-level data across modalities. On a token level, TerraMind encodes high-level contextual information to learn cross-modal relationships, while on a pixel level, TerraMind leverages fine-grained representations to capture critical

* Equal contribution
 † Equal supervision

spatial nuances. We pretrained TerraMind on nine geospatial modalities of a global, large-scale dataset. In this paper, we demonstrate that (i) TerraMind’s dual-scale early fusion approach unlocks a range of zero-shot and few-shot applications for Earth observation, (ii) TerraMind introduces “thinking in modalities” (TiM)—the capability of generating additional artificial data during finetuning and inference to improve the model output—and (iii) TerraMind achieves beyond state-of-the-art performance in community-standard benchmarks for EO like PANGAEA. The pretraining dataset, the model weights, and our code have been open-sourced under a permissive license.

1. Introduction

Earth observation (EO) increasingly benefits from multi-modality because of the important integration of complementary information from different data sources. This becomes particularly relevant as data from specific satellite missions can be unavailable for a specific time or location due to low revisiting times or weather phenomena like cloud coverage. Vice versa, for the computer vision domain, EO data is an important playground for the development of new approaches as there is significant publicly available data of very high quality and complexity. The available modalities range from sensors of different satellite missions to relevant complementary information like digital elevation.

In this work, we introduce TerraMind, the first large-scale, any-to-any generative, multimodal model for EO. With TerraMind, we introduce a dual-scale pretraining on pixel-level and token-level for generative multi-modality extending previous approaches like 4M [52] and demonstrate benefits over training primarily on tokens. TerraMind encodes high-level contextual information in tokens to enable correlation learning and scaling, while, additionally capturing important fine-grained representations using pixel-level inputs. During the pretraining, TerraMind predicts masked target tokens so that our pretraining objective boils down to a cross-modal patch classification problem that results in high-quality latent spaces. TerraMind is pretrained on a custom global-scale geospatial dataset named TerraMesh with 9 million samples that have been aligned spatiotemporally and across modalities [7]. In addition to radar and optical satellite imagery from the Copernicus Sentinel-1 (S-1) and Sentinel-2 (S-2) missions, our dataset contains task-specific modalities such as land use/land cover (LULC) and normalized difference vegetation index (NDVI) maps, metadata like digital elevation models (DEM) and geographic coordinates, and natural language in the form of captions. To the best of our knowledge, TerraMind represents the first truly generative, multimodal deep learning model for EO. Furthermore, in contrast to many current remote sensing models that utilize masked autoencoders (MAE) like [54], contrastive learning (CL),

or diffusion techniques, TerraMind uniquely demonstrates benefits of leveraging the autoregressive paradigm for Earth observation.

We provide an overview on TerraMind’s performance in the community-standard benchmark PANGAEA [49] in Figure 2 and highlight the generative capabilities of TerraMind in Figure 3. In summary, our key contributions are as follows: (i) We introduce a dual-scale approach for generative multimodal pre-training leveraging data on pixel-level and token-level, which outperforms other fusion approaches, enables generation of artificial data, and facilitates zero-shot and few-shot applications. (ii) We introduce *thinking in modalities* – similar to chain-of-thought approaches in LLMs – for multi-modal models in EO, demonstrating that infusing generated data during finetuning improves the performance. (iii) We demonstrate that TerraMind outperforms other geospatial foundation models both in unimodal and multimodal settings.

2. Related Work

Computer vision in Earth observation. Computer vision (CV) has significantly advanced Earth observation (EO) [76]. Many CV techniques, originally developed for natural image processing, have been adapted to EO [62], often with minimal modifications. A wide range of tasks benefit from these methods, including classification [16], semantic segmentation [72] (e.g., land cover mapping [20, 21]), change detection [59] (e.g., disaster response [19]), object detection [39] (e.g., vessel identification [55]), and regression (e.g., biomass estimation [53]). Deep learning architectures like CNNs [75] and Vision Transformers (ViTs) [17] have demonstrated strong performance, often surpassing tradi-




Figure 2. TerraMind outperforms other geospatial foundation models on PANGAEA benchmark [49] in finetuning. Performance is measured in mIoU and min-max scaled per dataset with max performance displayed in parentheses.

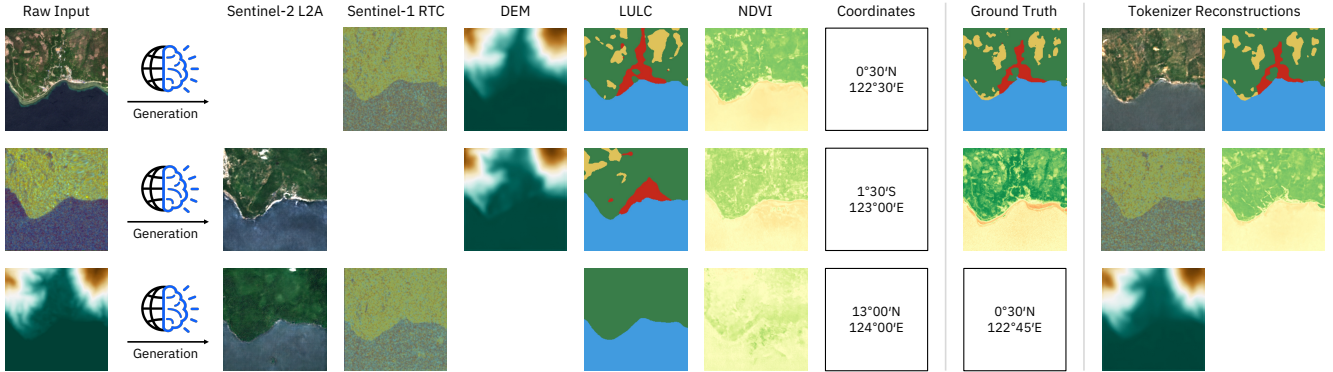


Figure 3. Chained generation example of TerraMindv1-B starting from either optical, radar, or digital elevation data. Left is input, middle is artificially generated data by TerraMind, right represents the sensor measurements and ground truths, respectively. Note that TerraMind performance is upper-bounded by the tokenizer performance, especially with respect to the diffusion decoder.

tional remote sensing (RS) methods. However, EO presents unique challenges, including diverse sensor modalities [4] and geospatial heterogeneity [46]. An emerging paradigm in EO is self-supervised learning (SSL) [64] and geospatial foundation models (GFMs) [45], which aim to leverage vast amounts of unlabeled RS data to develop general purpose task models [32]. While off-the-shelf CV models have shown promising results [36], they do not fully exploit the unique characteristics of geospatial data. Many GFMs still rely on generic CV architectures [50], which were not explicitly designed to handle the complexities of EO, such as heterogeneous sensor sources (e.g., optical, radar, DEM) [29], integrated with auxiliary data (e.g., text) [42, 47], and expert knowledge (e.g., prioritizing specific bands or indexes). In this direction, TerraMind better integrates domain-specific properties, developing a customized and expandable multimodal learning strategy.

Multimodality in CV. The current paradigm of multimodal CV is driven by the integration of diverse data streams [69], such as natural images [74], natural language text [10], temporal video data [58], and weather [70], within large foundation models [8]. Starting from the alignment of images and texts [57], these models moved beyond simple feature extraction, towards nuanced contextual understanding. The ability to combine several modalities allows for unprecedented capabilities in complex tasks [30], evidenced by the rapid advancement of multimodal Large Language Models (MLLMs) [30], that excel in tasks such as scene understanding [12], visual question answering [18], and video analysis [24]. Recent advances in architectural design [31] and large scale pre-training [11] have enabled the development of models that learn highly effective cross-modal representations [41], which can then be adapted to a wide variety of downstream tasks [66].

Multimodality in EO. Multimodality in EO originates from data fusion and is typically understood as the integra-

tion of SAR and optical data [13, 25, 28, 38] or the combination of optical data with vector data [5]. Some studies have explored alternative combinations of data. In [15], the authors introduce a contrastive framework for comparing RS images and street views. Even different optical sensors can be considered different modalities [48, 61]. Similarly, several multi-view images (i.e. multimodal) datasets [26, 44, 54] are introduced. More recent approaches combined text and images [40], both for discriminative [42] and generative [34] purposes. Lately, different GFMs are trained in a multimodal way [4, 54, 68], still focusing either on a specific set of modalities (e.g., vision [68], [3]) or tasks (e.g., generative [34]). Compared to multi-scale high-quality generation models for optical data, like MetaEarth [71], our approach allows to generate any modality from any other pretraining modality. To the best of our knowledge, no existing model has combined a wide and diverse amount of modalities both for discriminative and generative purposes, as TerraMind does. In Table 1, we provide a comparison of TerraMind with related model architectures.

3. Dataset

For the pretraining of TerraMind and its tokenizers, we create a multimodal dataset called TerraMesh, which will be open-sourced to the community. TerraMesh builds on existing datasets, which we expand by adding modalities from external data sources or by applying pseudo-labeling. We provide an overview of the aligned image modalities and a detailed dataset description in the supplementary material.

Base datasets. TerraMesh is based on SSL4EO-S12 [6, 65] and MajorTOM-Core [23], two unlabeled remote sensing datasets containing co-aligned radar and optical imagery from Sentinel-1 and Sentinel-2 satellites. SSL4EO-S12 has lower geographic coverage but is multi-seasonal. MajorTOM-Core covers most of the Earth’s land surface at a single timestamp. For MajorTOM-Core, we apply a sub-

Model	Modalities	Any-to-Any Generation	Multi-Scale Features
RemoteCLIP	optical, text	✗	✗
CROMA	optical, radar	✗	✗
AnySat	aerial, optical, radar, NAIP	✗	✗
DeCUR	optical, radar	✗	✗
DiffusionSat	optical, text	✗	✗
DOFA	optical, radar, hyperspectral, NAIP	✗	✗
MetaEarth	optical (unimodal)	✗	✓
Galileo	optical, radar, elevation, weather, location, population, ...	✗	✓
TerraMind	optical, radar, land use, elevation, vegetation index, location, text	✓	✓

Table 1. Comparison of TerraMind to other model architectures. While several model architectures have been developed to process multiple EO modalities, TerraMind represents the first truly generative model that can produce any modality from any other modality.

sampling scheme based on LULC classes and ecoregions. TerraMesh includes a total of approximately 9 million globally distributed samples from both Sentinel-1 and Sentinel-2, each measuring 264×264 pixels at 10m resolution.

Additional modalities. We obtain co-aligned yearly LULC maps by ESRI with nine land use classes. Additionally, we leverage SENSeI v2 [22] as a cloud and ice annotation model and update the ESRI LULC classes for better spatiotemporal alignment. NDVI maps are computed using the corresponding spectral bands from Sentinel-2. DEM is extracted from the Copernicus DEM 30m dataset [2], which provides global coverage of the Earth’s elevation at a 30m resolution. Captions are generated synthetically by constructing RGB images from Sentinel-2 patches using the corresponding spectral bands and processing them with LLaVA-Next [37]. A tailored prompt guides the model to describe the content of each image (see supplementary material for details). For geolocations, we round the latitude and longitude values from the center of each patch to the nearest quarter degree and save the resulting discretized coordinates as strings in a well-defined format (e.g., “lat=35.75 lon=-89.00”).

4. Methods

TerraMind pretraining is two-staged. We first pretrain unimodal tokenizer models, tokenize the modalities, and then leverage token-level and pixel-level input to pretrain the Ter-

raMind encoder-decoder architecture. In the following, we describe those individual stages.

4.1. Tokenization

We develop modality-specific tokenizers to encode each modality into a sequence of discrete tokens for pretraining and decode token sequences back to their original forms. Thus, TerraMind is in principle compatible with any modality, as long as it can be tokenized and aligned with other modalities. For reasons of space, we delegate most experiments related to the tokenizer performances to the supplementary material.

Image-like modalities. We train autoencoder-based architectures with a quantization step in the bottleneck for image-like modalities such as Sentinel-1, Sentinel-2, LULC, NDVI, and DEM. Tokenizer encoders process an input image and generate a latent representation for each 16×16 patch, which is then discretized with finite-scalar-quantization (FSQ) [51] into one of N codewords, where N represents the codebook size. All tokenizers use a vocabulary size of 16K besides the simpler LULC modality for which we use 4K. These codewords are then used by the diffusion decoder to reconstruct the original image. The benefit of leveraging diffusion decoders lies in facilitating cross-modal generation in TerraMind by transforming tokens back into images. By mapping each codeword to a unique integer in $\{0, 1, \dots, N - 1\}$, we obtain discrete tokens for each image patch. We pretrain the tokenizers in a self-supervised setting. FSQ as quantization method enhances training stability [51] compared to vector quantization [63]. FSQ improves training stability by using a fixed codebook instead of a learnable one, eliminating the need for codebook-related loss terms. Notably, FSQ is heavily influenced by ideas of neural compression [27]. For our 12 channel Sentinel2-L2A data, we achieve compression rates of over 3000x by quantizing the data. We summarize the architecture of our tokenizers in Figure 4. The main objective of the overall tokenizer is to encode image patches consistently into discrete tokens based on semantic similarity to enable cross-modal correlation learning. Thus, the loss of certain details is an expected trade-off, since the focus is on grouping similar patches rather than preserving all fine-grained features. Naturally, more accurate reconstructions facilitate the cross-modal generation, which is an important component of TerraMind – the main focus however lies on consistent cross-modal correlation learning. We provided further details on the pretraining of the tokenizers in the supplementary material and open-source code and weights for the tokenizer models.

Sequence-like modalities. We treat both captions and geolocations as text and use a single text tokenizer to process both modalities. By discretizing the geographic coordinates and representing them as strings in a structured format, we introduce special coordinate tokens into the tokenizer’s vo-

<https://planetarycomputer.microsoft.com/dataset/10-lulc-annual-v02>




Figure 4. Tokenizer for image-like modalities combining finite-scalar quantization [51] with diffusion decoding.

cabulary. This approach allows us to encode geolocations as a sequence of discrete tokens, beginning with a latitude token followed by a longitude token. For the textual data, we modify the existing WordPiece tokenizer from BERT [33] rather than fitting a new tokenizer specifically for our captions since satellite imagery and corresponding captions have far less variations than natural images.

4.2. Pre-training

Architecture. TerraMind uses a symmetric Transformer-based encoder-decoder architecture proposed in [52], which is designed to process sequences of multimodal tokens. In addition to discrete tokens, TerraMind accepts pixel-level inputs, specifically satellite imagery and digital elevation maps. For pixel-level inputs, we apply learnable patch-wise linear projections to generate patch embeddings for each 16×16 patch, similar to the approach used in ViT [17].

Dual-scale early fusion. In contrast to [52], we not only embed token-level data but additionally leverage pixel-level data across a range of input modalities to introduce a dual-scale feature representation to support the structuring of the embedding space. Both tokens and patches represent a 16×16 pixel area. Tokens represent this area via a single discrete integer value, while the image patches describe the same area with the actual floating point sensor data. During pretraining, the model, thus, not only learns a correlation *between* modalities (i.e., cross-modal learning) but also between different levels of abstraction *within* the same modality. The low-level token information enables the cross-modal correlation learning, while adding pixel level input accounts for spatial nuances. Based on dual-scale features the model further learns to better structure pixel-level data in the embedding space via the corresponding information from the discrete token. We illustrate the pretraining paradigm in Figure 5. The model is agnostic to processing tokens or patches in the input space, while the target is generally token-level data. We use six pixel-level modalities and eight token-level modalities. Optical, radar, and elevation data is present both in pixel-space and token-space.

Masking strategy. TerraMind applies a masked modeling approach in the token space following [52]. The model leverages a set of randomly selected target tokens that have




Figure 5. Illustration of the pre-training task. Given a multimodal sample, random subsets of patches, input tokens, and target tokens are sampled from a Dirichlet distribution. Based on the encoded input, the decoder predicts target tokens for the masked input.

to be reconstructed from a randomly selected set of input tokens and pixel-level data. During pre-training, we sample input and target data from a Dirichlet distribution:

$$f(\mathbf{x} | \boldsymbol{\alpha}) = \frac{1}{B(\boldsymbol{\alpha})} \prod_{i=1}^M p_i^{\alpha_i - 1} \quad (1)$$

where $\mathbf{p} = (p_1, p_2, \dots, p_M)$ is a vector of probabilities, and $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_M)$ is a vector of weighting parameters that we leverage during pretraining as a steering wheel to guide the model. $B(\boldsymbol{\alpha})$ refers to a multivariate Beta function.

We opt for masked token reconstruction to familiarize the model with the absence of entire modalities, which is crucial for a high usability of a multimodal model in Earth observation. During pre-training, the model learns an internal representation of unseen modalities which is expected to benefit a range of downstream applications. In addition, sampling input and target tokens improves the computational efficiency of the pre-training, as each token is a compressed representation of a patch with compression factors of between 250x and 3000x depending on the modality. Finally, without tokenized representations of the image-like modalities, it is challenging to learn the correlation to sequence-like modalities. The overall training objective of TerraMind boils down to a cross-modal patch-level classification problem that we represent via a cross entropy loss:

$$\mathcal{L}_{\text{CE}} = - \sum_{i=1}^N y_i \log(p_i), \quad (2)$$

where y_i is the one-hot encoded true class of token i , p_i is the predicted probability for token i , N is the total number of possible tokens. Interestingly, we can infer an upper bound loss for a random model. When a model outputs arbitrary target tokens, the cross entropy loss will collapse to the natural logarithm of the vocabulary size:

$$\mathcal{L}_{\text{CE, random}} = - \sum_{i=1}^N y_i \log\left(\frac{1}{N}\right) = \log(N), \quad (3)$$

which we validated empirically as an upper bound for the TerraMind pretraining loss. For a vocabulary size of 16K, a random model would have a validation loss of 9.68.

Scaling. We trained three versions of TerraMind scaling across model size, compute, and data. In addition, we pre-train different versions of TerraMind with respect to the number of dual-scale features. **TerraMindv1-B** is pre-trained on 500B tokens for 6 days on 32 NVIDIA A100 GPUs. The model uses dual-scale features from both token-level and pixel-level. During initial experiments, we observed significant improvements from scaling model size when switching from a tiny backbone to a small backbone to a base backbone. Therefore, we pre-trained **TerraMindv1-L** on a large backbone with 500B tokens on 32 NVIDIA A100 GPUs trained for 10 days. Finally, to better understand the effect of scaling across the dual-scale feature representation, we pre-train **TerraMindv1-B-single** as a single-scale model on primarily token-level data with optical S-2 L2A data as only pixel-level input (compared to pixel-level S-2 L1C, S-2 RGB, S-1 GRD, S-1 RTC, and DEM in TerraMindv1-B and -L). TerraMindv1-B-single is pretrained on 500B tokens from over one million samples for 6 days on 32 NVIDIA A100 GPUs. We summarize the scaling behavior in model size, compute, and data in Figure 9 of the supplementary material. We additionally provide final validation losses in Table 9 comparing v1-B and v1-L with the theoretical random loss.

4.3. Generation

After pretraining, TerraMind can generate tokens for any modality, conditioned on any subset of input modalities. The generative capabilities further unlock various zero-shot tasks, such as water body segmentation, land use segmentation, image captioning, and geolocation prediction. For the generation of image-like modalities, the decoder receives mask tokens for the modality to be generated and predicts the corresponding tokens based on the encoded input. For sequence-like modalities, the decoder generates the output autoregressively. After generating tokens from the target modality, the corresponding tokenizer decoder allows to map from token-space to image or text space. TerraMind further supports chained generation which ensures consistency across generated modalities. The chained generation represents a conditional probability distribution where the prior probability distribution $P(\mathbf{M}_1)$ is given by the input modality \mathbf{M}_1 , and all subsequent modalities \mathbf{M}_2 to \mathbf{M}_K are generated conditioned on the input modality and potentially other generated modalities:

$$P(\mathbf{M}_1, \dots, \mathbf{M}_K) = P(\mathbf{M}_1) \prod_{i=2}^K P(\mathbf{M}_i | \mathbf{M}_1, \dots, \mathbf{M}_{i-1}). \quad (4)$$

4.4. Thinking-in-Modalities

Thinking in Modalities (TiM) is a recursive fine-tuning and inference technique designed to enhance multimodal learning by leveraging the generative capabilities of the model itself. Given an input $x \in \mathcal{X}$ (e.g., an optical satellite image), the model first generates additional synthetic modalities $\tilde{x} = f_{\text{gen}}(x)$ on a token-level using a learned generative function f_{gen} . These generated tokens are then concatenated with the original input and jointly processed by the downstream model f (e.g., TerraMind encoder with a segmentation head), yielding the final output $y = f(x, f_{\text{gen}}(x))$. This formulation allows the model to reason over both observed and inferred modalities, effectively enriching the input space. TiM can leverage multiple generated modalities which are then generated in a chained approach (see Eq. 4). For example, for k modalities, the input is augmented with newly generated modalities:

$$\tilde{x}^{(k+1)} = \tilde{x}^{(k)} \cup f_{\text{gen}}(x^{(k)}), \quad (5)$$

and the final model output is described by:

$$y = f(\tilde{x}^{(K)}). \quad (6)$$

This recursive augmentation mimics a chain-of-thought process, enabling the model to iteratively refine its internal representation and improve performance on downstream tasks, particularly in scenarios with missing modalities.

5. Experiments

In this section, we describe the performance gains resulting from TerraMind and experiment with the unlocked capabilities through the any-to-any generation of TerraMind. We especially investigate the quality of the artificially generated data as an input to subsequent finetuning tasks similar to chain-of-thought in LLMs.

5.1. Foundational experiments

Multimodality vs. unimodality. As a first motivational experiment, we outline the benefit of using multimodal data in Earth observation at the example of water body mapping. Specifically, we leverage the ViT-B encoders from the unimodal tokenizer models for S-1, S-2, and LULC, concatenate their embeddings, and train a segmentation head with four ConvNeXt [43] blocks as a late fusion approach. The results in Table 2 (left) suggest that regardless of which modalities we combine, the combination of two modalities always outperforms each unimodal model. Combining all three modalities achieves the best overall performance.

Token-level fusion vs. late fusion. In Table 2 (right), we investigate the effects of fusing the inputs on a token level through masked token reconstruction compared to the late fusion from the previous paragraph. We observe that token-level fusion outperforms late fusion in water body

Input	Late fusion	Token-level fusion
S-1	61.01	63.94 (2.93pp \uparrow)
S-2	72.70	76.32 (3.62pp \uparrow)
LULC	71.77	70.96 (0.81pp \downarrow)
S-1 + S-2	73.83	76.74 (2.91pp \uparrow)
S-1 + LULC	73.86	73.76 (0.10pp \downarrow)
S-2 + LULC	75.65	77.04 (1.39pp \uparrow)
S-1 + S-2 + LULC	76.00	76.88 (0.88pp\uparrow)

Table 2. Task: Water body mapping on Sen1Floods11 [9], performance measured in IoU on water class. Model sizes and architectures are comparable. **Left column:** Late fusion of tokenizers. Effect of leveraging multimodal data in late-fusion fine-tuning of frozen modality-specific ViT-B models. The average improvement of full multimodality over the individual unimodal performance is 7.5pp IoU. **Right column:** Finetuning results of TerraMind.v1-B-single as a mid fusion approach based on masked correlation learning. Performance gains over late fusion are depicted as percentage points increases/decreases in parentheses.

mapping and the performance gains are particularly high when LULC data is not available. This suggests that early fusion captures an internal representation of the multimodal state—especially pronounced for LULC—that benefits fine-tuning. With those findings in mind, we will explore the effects of using additional multi-modal pixel-level input in a dual-scale pretraining in Section 5.5.

5.2. Generation experiments


TerraMind supports the generation of any modality from any other modality. In the following, we provide examples of the generation performance starting from: (i) an information-rich modality, like optical S-2 L2A data, and (ii) minimal information based on the geolocation. In Figure 3, we observe that TerraMind performs very well in generating image-like modalities like radar data, land use maps, and digital elevation maps from optical S-2 L2A data. We provide a quantitative overview on the quality of the generations on unseen

Modalities	MAE \downarrow	RMSE \downarrow	SSIM \uparrow	PSNR \uparrow
S-1 GRD \rightarrow S-2 L2A	0.074	0.116	0.750	26.210
S-1 GRD \rightarrow DEM	163.0	320.8	0.878	20.694
S-1 GRD \rightarrow NDVI	0.180	0.225	0.438	18.990
S-1 RTC \rightarrow S-2 L2A	0.113	0.194	0.695	24.251
S-1 RTC \rightarrow DEM	298.8	799.2	0.873	20.009
S-1 RTC \rightarrow NDVI	0.172	0.211	0.465	19.529
S-2 L2A \rightarrow S-1 GRD	2.942	3.877	0.531	28.678
S-2 L2A \rightarrow S-1 RTC	2.636	3.391	0.430	28.993
S-2 L2A \rightarrow DEM	215.8	745.5	0.942	20.616

Table 3. Quantitative evaluation of generations on unseen global validation dataset using 10 diffusion steps. MAE and RMSE metrics are in physical units: meter (DEM), reflectance (S-2), and db (S-1).




(a) **Input:** S-2 L2A data capturing Singapore in January 2025.



(b) **Generation:** S-1 RTC composition generated by TerraMind.



(c) **Input:** S-2 L2A data capturing Northern Spain in January 2025.



(d) **Generation:** S-1 GRD composition generated by TerraMind.

Figure 6. Generated Sentinel-1 data using TerraMind over Singapore and Northern Spain.

validation data in Table 3. Overall, we observe an interesting asymmetry in the generative performance of TerraMind where (a) radar-to-optical generation achieves reasonable generation quality in terms of SSIM and PSNR – indicating structural and visual fidelity with some perceptual degradation – and (b) optical-to-radar generation yields higher PSNR values but lower SSIM, suggesting visually plausible outputs that lack strong structural alignment. The quality of generated DEM suggest to be structurally very strong, but noisy. The errors for DEM generations suggest that the generations structurally make sense but the level of altitude is difficult to infer for the model. We compare these scores with the reconstruction quality of the auto-encoding tokenizers in the supplementary material that can serve as upper bounds. Additionally, we provide experiments on the generation quality using token-level instead of pixel-level inputs in the supplementary material. Finally, we demonstrate the quality of large-scale generations across Singapore and Northern Spain in Figures 19 and 20.

5.3. Zero-shot experiments

Based on its generative capabilities, TerraMind unlocks several zero-shot applications, like land-use segmentation, water body mapping, geo-localization, and vegetation mapping. In the following, we focus on water body mapping and geo-localization as image- and sequence-level zero-shot tasks.

Water body mapping. In Table 4, we compare the zero-shot performance of TerraMind with its fine-tuned performance and other finetuned benchmarks for water body mapping. Overall, TerraMindv1-B-single achieves a zero-shot IoU of 45.4% compared to SOTA-level fine-tuning performance of 84.4% of DeCUR. In ablations with TerraMindv1-single trained on DynamicWorld LULC data, we boost this to up to 69.8% suggesting that TerraMind harnesses up to over 80% of the SOTA performance in zero-shot setting. Additionally, it’s notable that none of the benchmarking model can be applied in a zero-shot context, highlighting the relevance of TerraMind’s capabilities.

Model	Input	Type	IoU _{Water}
TerraMindv1-B	S-2	zero-shot	45.40
TerraMindv1-B-single	S-2	zero-shot	69.75
Prithvi 2.0 / DeCUR / ...		zero-shot	N/A
Baseline [9]	S-2	finetune	31.25
Prithvi 2.0 300M	S-2	finetune	80.97
DeCUR	S-2	finetune	82.17

Table 4. Zero-shot results of TerraMind on water body mapping compared to fine-tuned performance of benchmarks.

Geo-localization. TerraMind is able to predict the geolocation of a specific data instance. To better visualize the geo-localization capabilities, we prompt the model for the most

likely locations of the land use class “bare land” (deserts etc.) in a Monte-Carlo-sampling in Figure 7. The probability distribution of the model fits the expectation of where to find bare land, highlighting the Sahara region and middle-east, as well as Mexico and Southern California.




Figure 7. Prediction distribution of the land use class “bare land” with a sampling temperature of $T = 1.0$. TerraMind has an accurate internal representation of the geolocation of specific contexts, like land use classes.

5.4. Few-shot experiments

TerraMind is pretrained on a cross-modal patch classification objective. Therefore, we expect a well-structured latent space that clusters different concepts accurately. To investigate our hypothesis, we apply 1-Nearest-Neighbor (1-NN) classification experiments in the community-standard setting of 1-shot 5-way on two datasets: EuroSAT and METER-ML. In those experiments, there are no weight updates of any kind, so that we can assess the quality of the embedding space structure.

Model	Input	EuroSAT	METER-ML
CLIP-ViT-B/16	S-2 RGB	57.00	29.15
CLIP-ViT-B/16	NAIP	–	32.01
DeCUR	S-2 L1C	50.54	27.87
Prithvi 1.0 100M	S-2 L1C	60.11	26.08
Prithvi 2.0 300M	S-2 L1C	<u>61.06</u>	28.26
TerraMindv1-B	S-2 L1C	70.83	33.90
TerraMindv1-B	NAIP	–	<u>32.23</u>

Table 5. 1-shot 5-way classification results on EuroSAT and METER-ML measured in mean accuracy \uparrow , averaged over 200 runs. TerraMind outperforms benchmarks from CV and EO domain, suggesting a well-structured latent space.

In Table 5, we observe that TerraMind outperforms several other benchmarks from both the CV and EO domain on the EuroSAT dataset by at least 10pp in accuracy. Our results further show that for methane source classification on METER-ML, TerraMind outperforms benchmark models and generalizes to high-resolution NAIP data with one order of magnitude higher resolution than the pre-training data. We present additional experiments with other few-shot settings in the supplementary material.

Model	BurnSr*	MADOS*	PASTIS	Sen1FI11	FBP*	DEN*	CTM-SS	SN7*	AI4Farms*	Avg. mIoU	Avg. Rank
CROMA	82.42	67.55	32.32	<u>90.89</u>	51.83	38.29	49.38	59.28	25.65	55.29	6.61
DOFA	80.63	59.58	30.02	89.37	43.18	<u>39.29</u>	51.33	61.84	27.07	53.59	8.22
GFM-Swin	76.90	64.71	21.24	72.60	67.18	34.09	46.98	60.89	27.19	52.42	10.00
Prithvi 1.0 100M	83.62	49.98	33.93	90.37	46.81	27.86	43.07	56.54	26.86	51.00	11.00
RemoteCLIP	76.59	60.00	18.23	74.26	69.19	31.78	52.05	57.76	25.12	51.66	11.22
SatlasNet	79.96	55.86	17.51	90.30	50.97	36.31	46.97	61.88	25.13	51.65	10.67
Scale-MAE	76.68	57.32	24.55	74.13	<u>67.19</u>	35.11	25.42	62.96	21.47	49.43	11.44
SpectralGPT	80.47	57.99	35.44	89.07	33.42	37.85	46.95	58.86	26.75	51.87	10.11
S.-S12-MoCo	81.58	51.76	34.49	89.26	53.02	35.44	48.58	57.64	25.38	53.02	10.06
S.-S12-DINO	81.72	49.37	36.18	88.61	51.15	34.81	48.66	56.47	25.62	52.51	10.89
S.-S12-MAE	81.91	49.90	32.03	87.79	51.92	34.08	45.80	57.13	24.69	51.69	12.39
S.-S12-Data2Vec	81.91	44.36	34.32	88.15	48.82	35.90	54.03	58.23	24.23	52.22	10.72
UNet Baseline	84.51	54.79	31.60	91.42	60.47	39.46	47.57	<u>62.09</u>	46.34	57.58	4.89
ViT Baseline	81.58	48.19	38.53	87.66	59.32	36.83	44.08	52.57	<u>38.37</u>	54.13	10.28
TerraMindv1-B-single	<u>84.00</u>	65.01	<u>40.80</u>	90.32	–	–	52.66	59.71	27.71	–	–
TerraMindv1-B	82.42	<u>69.52</u>	40.51	90.62	59.72	37.87	55.80	60.61	28.12	<u>58.35</u>	<u>3.94</u>
TerraMindv1-L	82.93	75.57	43.13	90.78	63.38	37.89	<u>55.04</u>	59.98	27.47	59.57	3.44

Table 6. Performance evaluation of TerraMind across nine benchmark datasets using the PANGAEA evaluation protocol. Higher mIoU values (\uparrow) and lower rank values (\downarrow) are reported. The best model per column is highlighted in bold and the second best is underscored. We indicate unimodal datasets with *, demonstrating that TerraMind outperforms other benchmarks in both unimodal and multimodal settings. All foundation model encoders are frozen, while U-Net and ViT baselines are trained from scratch for each specific task.

5.5. Fine-tuning experiments

Besides the novel capabilities that TerraMind introduces, we benchmark the fine-tuning performance of TerraMind in both unimodal and multimodal settings following the community-standard PANGAEA benchmark [49]. We summarize the results in Table 6. Overall, TerraMindv1-B outperforms all other GeoFMs by at least 3pp avg. mIoU. Importantly, we observe that TerraMind is the only foundation model approach in EO that across the PANGAEA benchmark outperforms task-specific U-Net models. Performance is approximately 2pp avg. mIoU higher for TerraMindv1-L, with a peak of 5pp in multimodal datasets. Furthermore, TerraMindv1-L outperforms also specialised ViT baselines by 5pp avg. mIoU. Finally, TerraMindv1-B outperforms TerraMindv1-B-single by 1pp. avg. mIoU. The latter result demonstrates the benefit of using dual-scale features when comparing v1-B-single with v1-B. Note that per suggestion of the PANGAEA authors, we exclude the xView2 and BioMassters task as we could not reproduce the reported performances. Finally, we assess the impact of leveraging multimodal data as input to TerraMindv1-B compared to utilizing either optical or radar data as unimodal input to better understand the effect of leveraging multimodal data in finetuning. We observe that across all three multimodal tasks, TerraMindv1-B performs best with access to both optical and radar data with a significant improvement for CTM-SS.

5.6. Thinking in modalities

Based on the generative capabilities of TerraMind and the fact that multi-modal data improved downstream task performances in many scenarios, we additionally investigate the

	PASTIS	Sen1FI11	CTM-SS
S-1	20.04	80.39	24.45
S-2	40.20	89.57	50.90
Multimodal	40.51	90.62	55.80

Table 7. Benefit of using multimodal input in TerraMind.v1-B on multimodal datasets of PANGAEA. We report the mIoU (%) \uparrow .

value of artificially generated data for finetuning tasks. Borrowing on the idea of chain-of-thought finetuning of large language models, we use TerraMind to generate artificial data as an intermediate step to improve finetuning. We provide an overview on this new method in the beginning of this paper in Figure 1.

Fine-Tuning	Input	IoU _{Water}	mIoU
TerraMindv1-B	S-1	68.00	81.06
TerraMindv1-B	S-2	82.26	89.70
TerraMindv1-B TiM	S-1 + <i>gen. LULC</i>	72.25	83.65
TerraMindv1-B TiM	S-2 + <i>gen. LULC</i>	84.75	91.14

Table 8. Thinking-in-modalities (TiM) tuning compared with standard full fine-tuning approaches on the Sen1Floods11 dataset.

We employ the TiM tuning approach on water body mapping and *during finetuning* use radar or optical data to generate additional artificial LULC data which we leverage along the input. Our results in Table 8 indicate a superior performance of the TiM finetuning compared to leveraging

uni-modal data by up to 2pp mIoU. We also note that TiM tuning strongly benefits from a trainable encoder. This finding points us in the direction of TerraMind being able to generate data that improve finetuning to downstream applications.

6. Conclusion

TerraMind’s pre-training approach of combining token-level and pixel-level information has unlocked a range of new model capabilities in Earth Observation. TerraMind demonstrates not only beyond state-of-the-art performance when tested against 12 other geospatial models on the PANGEA benchmark, it also represents the first truly generative multi-modal model in the domain. Because of the ability of integrating heterogeneous and diverse data sources demonstrated by the TerraMind’s approach, we expect that TerraMind-like models will expand soon to multi-temporal, multi-resolution, and hyperspectral data. This will allow to maximise impact in real world applications and fully leverage the data rich ecosystem available in the Earth Observation domain.

References

- [1] A. Hore and D. Ziou. Image quality metrics: PSNR vs. SSIM. In *Proc. 20th International Conference on Pattern Recognition (ICPR)*, pp. 2366-2369, 2010. 17
- [2] European Space Agency. Copernicus dem. <http://dx.doi.org/10.5270/ESA-c5d3d65>, 2022. 4
- [3] Guillaume Astruc, Nicolas Gonthier, Clement Mallet, and Loic Landrieu. Anysat: An earth observation model for any resolutions, scales, and modalities. *arXiv preprint arXiv:2412.14123*, 2024. 3
- [4] Guillaume Astruc, Nicolas Gonthier, Clement Mallet, and Loic Landrieu. Omnisat: Self-supervised modality fusion for earth observation, 2024. 3
- [5] Nicolas Audebert, Bertrand Le Saux, and Sébastien Lefèvre. Joint Learning from Earth Observation and OpenStreetMap Data to Get Faster Better Semantic Maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1552–1560, 2017. 3
- [6] Benedikt Blumenstiel, Nassim Ait Ali Braham, Conrad M Albrecht, Stefano Maurogiovanni, and Paolo Fraccaro. SSL4EOS12 v1.1 – A Multimodal, Multiseasonal Dataset for Pretraining. *arXiv preprint arXiv:2503.00168*, 2025. 3, 14
- [7] Benedikt Blumenstiel, Paolo Fraccaro, Valerio Marsocci, Johannes Jakubik, Stefano Maurogiovanni, Mikolaj Czerkawski, Rocco Sedona, Gabriele Cavallaro, Thomas Brunswiler, Juan Bernabe-Moreno, and Nicolas Longépé. Terramesh: A planetary mosaic of multimodal earth observation data. *arXiv preprint arXiv:2504.11172*, 2025. 2
- [8] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 3
- [9] Derrick Bonafilia, Beth Tellman, Tyler Anderson, and Erica Issenberg. Sen1floods11: A georeferenced dataset to train and test deep learning flood algorithms for sentinel-1. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2020. 7, 8
- [10] Florian Bordes, Richard Yuanzhe Pang, Anurag Ajay, Alexander C Li, Adrien Bardes, Suzanne Petryk, Oscar Mañas, Zhiqiu Lin, Anas Mahmoud, Bargav Jayaraman, et al. An introduction to vision-language modeling. *arXiv preprint arXiv:2405.17247*, 2024. 3
- [11] Jize Cao, Zhe Gan, Yu Cheng, Licheng Yu, Yen-Chun Chen, and Jingjing Liu. Behind the scene: Revealing the secrets of pre-trained vision-and-language models. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, pages 565–580. Springer, 2020. 3
- [12] Xu Cao, Tong Zhou, Yunsheng Ma, Wenqian Ye, Can Cui, Kun Tang, Zhipeng Cao, Kaizhao Liang, Ziran Wang, James M Rehg, et al. Maplm: A real-world large-scale vision-language benchmark for map and traffic scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21819–21830, 2024. 3
- [13] Yuxing Chen and Lorenzo Bruzzone. Self-supervised change detection in multi-view remote sensing images. *arXiv preprint arXiv:2103.05969*, 2021. 3
- [14] Chenwei Wang, et al. SAR Target Image Generation Method Using Azimuth-Controllable Generative Adversarial Network. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (JSTARS)*, Vol. 15, 2022. Online: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9933645&tag=1>. 17
- [15] Fabian Deuser, Konrad Habel, and Norbert Oswald. Sample4geo: Hard negative sampling for cross-view geolocalisation. *arXiv preprint arXiv:2303.11851*, 2023. 3
- [16] Ivica Dimitrovski, Ivan Kitanovski, Dragi Koccev, and Nikola Simidjievski. Current trends in deep learning for earth observation: An open-source benchmark arena for image classification. *ISPRS Journal of Photogrammetry and Remote Sensing*, 197:18–35, 2023. 2
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. 2, 5
- [18] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Ayaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, et al. Palm-e: An embodied multimodal language model. 2023. 3
- [19] Victor Durnov. xvview2 1st place solution. 2
- [20] Adam Van Etten, Dave Lindenbaum, and Todd M. Bacastow. Spacenet: A remote sensing dataset and challenge series, 2019. 2
- [21] Casper Fibaek, Luke Camilleri, Andreas Luyts, Nikolaos Dionelis, and Bertrand Le Saux. PhilEO Bench: Evaluat-

- ing Geo-Spatial Foundation Models, In Proc. Int Geoscience and Remote Sensing Symposium (IGARSS), 2024. 2
- [22] Alistair Francis. Sensor independent cloud and shadow masking with partial labels and multimodal inputs. *IEEE Transactions on Geoscience and Remote Sensing*, 2024. 4, 14
- [23] Alistair Francis and Mikolaj Czerkawski. Major tom: Expandable datasets for earth observation. *arXiv preprint arXiv:2402.12095*, 2024. 3, 14
- [24] Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. *arXiv preprint arXiv:2405.21075*, 2024. 3
- [25] Anthony Fuller, Koreen Millard, and James R. Green. Croma: Remote sensing representations with contrastive radar-optical masked autoencoders, 2023. 3
- [26] Anatol Garioud, Nicolas Gonthier, Loic Landrieu, Apolline De Wit, Marion Valette, Marc Poupée, Sebastien Giordano, and Boris Watrelos. FLAIR : a country-scale land cover semantic segmentation dataset from multi-source optical imagery. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 3
- [27] Carlos Gomes, Isabelle Wittmann, Damien Robert, Johannes Jakubik, Tim Reichelt, Michele Martone, Stefano Maurogiovanni, Rikard Vinge, Jonas Hurst, Erik Scheurer, et al. Lossy neural compression for geospatial analytics: A review. *arXiv preprint arXiv:2503.01505*, 2025. 4
- [28] Sebastian Hafner, Yifang Ban, and Andrea Nascetti. Unsupervised domain adaptation for global urban extraction using sentinel-1 sar and sentinel-2 msi data. *Remote Sensing of Environment*, 280:113192, 2022. 3
- [29] Boran Han, Shuai Zhang, Xingjian Shi, and Markus Reichstein. Bridging remote sensors with multisensor geospatial foundation models, 2024. 3
- [30] Soyeon Caren Han, Feiqi Cao, Josiah Poon, and Roberto Navigli. Multimodal large language models and tunings: Vision, language, sensors, audio, and beyond. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 11294–11295, 2024. 3
- [31] Jitesh Jain, Jianwei Yang, and Humphrey Shi. Vcoder: Versatile vision encoders for multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27992–28002, 2024. 3
- [32] Johannes Jakubik, Sujit Roy, C. E. Phillips, Paolo Fraccaro, Denys Godwin, Bianca Zadrozny, Daniela Szwarzman, Carlos Gomes, Gabby Nyirjesy, Blair Edwards, Daiki Kimura, Naomi Simumba, Linsong Chu, S. Karthik Mukkavilli, Devyani Lambhate, Kamal Das, Ranjini Bangalore, Dario Oliveira, Michal Muszynski, Kumar Ankur, Muthukumaran Ramasubramanian, Iksha Gurung, Sam Khallaghi, Hanxi, Li, Michael Cecil, Maryam Ahmadi, Fatemeh Kordi, Hamed Alemohammad, Manil Maskey, Raghu Ganti, Kommy Wellemariam, and Rahul Ramachandran. Foundation models for generalist geospatial artificial intelligence, 2023. 3
- [33] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, page 2. Minneapolis, Minnesota, 2019. 5
- [34] Samar Khanna, Patrick Liu, Linqi Zhou, Chenlin Meng, Robin Rombach, Marshall Burke, David Lobell, and Stefano Ermon. Diffusionsat: A generative foundation model for satellite imagery, 2023. 3
- [35] Kohei Arai, Michihiro Mikamo, and Shunsuke Onishi. Method for Image Quality Evaluation of Satellite-based SAR Data. *International Journal of Advanced Computer Science and Applications (IJACSA)*, Vol. 14, No. 7, 2023. Online: http://thesai.org/Downloads/Volume14No7/Paper_13-Method_for_Image_Quality_Evaluation_of_Satellite_based_SAR_Data.pdf. 17
- [36] Saad Lahrichi, Zion Sheng, Shufan Xia, Kyle Bradbury, and Jordan Malof. Is self-supervised pre-training on satellite imagery better than imagenet? a systematic study with sentinel-2. *arXiv preprint arXiv:2502.10669*, 2025. 3
- [37] Bo Li, Kaichen Zhang, Hao Zhang, Dong Guo, Renrui Zhang, Feng Li, Yuanhan Zhang, Ziwei Liu, and Chunyuan Li. Llavavnext: Stronger llms supercharge multimodal capabilities in the wild, 2024. 4, 14
- [38] Jiaxin Li, Danfeng Hong, Lianru Gao, Jing Yao, Ke Zheng, Bing Zhang, and Jocelyn Chanussot. Deep learning in multimodal remote sensing data fusion: A comprehensive review. *International Journal of Applied Earth Observation and Geoinformation*, 112:102926, 2022. 3
- [39] Ke Li, Gang Wan, Gong Cheng, Liqiu Meng, and Junwei Han. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS journal of photogrammetry and remote sensing*, 159:296–307, 2020. 2
- [40] Xiang Li, Congcong Wen, Yuan Hu, Zhenghang Yuan, and Xiao Xiang Zhu. Vision-language models in remote sensing: Current progress and future trends, 2024. 3
- [41] Zhiqiu Lin, Samuel Yu, Zhiyi Kuang, Deepak Pathak, and Deva Ramanan. Multimodality helps unimodality: Cross-modal few-shot learning with multimodal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19325–19337, 2023. 3
- [42] Fan Liu, Delong Chen, Zhangqingyun Guan, Xiaocong Zhou, Jiale Zhu, Qiaolin Ye, Liyong Fu, and Jun Zhou. Remoteclip: A vision language foundation model for remote sensing, 2024. 3
- [43] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s, 2022. 6
- [44] Gabriel Machado, Edemir Ferreira, Keiller Nogueira, Hugo Oliveira, Matheus Brito, Pedro Henrique Targino Gama, and Jefersson Alex dos Santos. Airound and cv-brct: Novel multiview datasets for scene classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:488–503, 2020. 3
- [45] Gengchen Mai, Chris Cundy, Kristy Choi, Yingjie Hu, Ni Lao, and Stefano Ermon. Towards a foundation model for geospatial artificial intelligence (vision paper). In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, New York, NY, USA, 2022. Association for Computing Machinery. 3

- [46] Oscar Manas, Alexandre Lacoste, Xavier Giró-i Nieto, David Vazquez, and Pau Rodriguez. Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9414–9423, 2021. 3
- [47] Clive Tinashe Marimo, Benedikt Blumenstiel, Maximilian Nitsche, Johannes Jakubik, and Thomas Brunschweiler. Beyond the visible: Multispectral vision-language learning for earth observation. *arXiv preprint arXiv:2503.15969*, 2025. 3
- [48] Valerio Marsocci and Nicolas Audebert. Cross-sensor self-supervised training and alignment for remote sensing, 2024. 3
- [49] Valerio Marsocci, Yuru Jia, Georges Le Bellier, David Kerekes, Liang Zeng, Sebastian Hafner, Sebastian Gerard, Eric Brune, Ritu Yadav, Ali Shibli, et al. Pangaea: A global and inclusive benchmark for geospatial foundation models. *arXiv preprint arXiv:2412.04204*, 2024. 2, 9, 19
- [50] Matias Mendieta, Boran Han, Xingjian Shi, Yi Zhu, Chen Chen, and Mu Li. Gfm: Building geospatial foundation models via continual pretraining. *arXiv preprint arXiv:2302.04476*, 2023. 3
- [51] Fabian Mentzer, David Minnen, Eirikur Agustsson, and Michael Tschannen. Finite scalar quantization: Vq-vae made simple. *arXiv preprint arXiv:2309.15505*, 2023. 4, 5, 16
- [52] David Mizrahi, Roman Bachmann, Oğuzhan Fatih Kar, Teresa Yeo, Mingfei Gao, Afshin Dehghan, and Amir Zamir. 4m: Massively multimodal masked modeling, 2023. 2, 5
- [53] Andrea Nascetti, RITU YADAV, Kirill Brodt, Qixun Qu, Hongwei Fan, Yuri Shendryk, Isha Shah, and Christine Chung. Biomasters: A benchmark dataset for forest biomass estimation using multi-modal satellite time-series. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 2
- [54] Vishal Nedungadi, Ankit Kariryaa, Stefan Oehmcke, Serge Belongie, Christian Igel, and Nico Lang. Mmearth: Exploring multi-modal pretext tasks for geospatial representation learning. *arXiv preprint arXiv:2405.02771*, 2024. 2, 3
- [55] Fernando Paolo, Tsu ting Tim Lin, Ritwik Gupta, Bryce Goodman, Nirav Patel, Daniel Kuster, David Kroodsmma, and Jared Dunnmon. xview3-sar: Detecting dark fishing activity using synthetic aperture radar imagery, 2022. 2
- [56] Prabhishkek Singh and Raj Shree. Analysis and effects of speckle noise in SAR images. In *Proc. International Conference on Advances in Computing, Communication, & Automation (ICACCA)*, 2016. DOI: 10.1109/ICACCA.2016.7748978. Online: <http://ieeexplore.ieee.org/document/7748978>. 17
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PmLR, 2021. 3, 18
- [58] Shuhuai Ren, Linli Yao, Shicheng Li, Xu Sun, and Lu Hou. Timechat: A time-sensitive multimodal large language model for long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14313–14323, 2024. 3
- [59] Ayesha Shafique, Guo Cao, Zia Khan, Muhammad Asad, and Muhammad Aslam. Deep learning-based change detection in remote sensing images: A review. *Remote Sensing*, 14(4): 871, 2022. 2
- [60] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017. 18
- [61] Aidan M Swope, Xander H Rudelis, and Kyle T Story. Representation learning for remote sensing: An unsupervised sensor fusion approach. *arXiv preprint arXiv:2108.05094*, 2021. 3
- [62] Devis Tuia, Konrad Schindler, Begüm Demir, Gustau Camps-Valls, Xiao Xiang Zhu, Mrinalini Kochupillai, Sašo Džeroski, Jan N. van Rijn, Holger H. Hoos, Fabio Del Frate, Mihai Datcu, Jorge-Arnulfo Quiané-Ruiz, Volker Markl, Bertrand Le Saux, and Rochelle Schneider. Artificial intelligence to advance earth observation: a perspective, 2023. 2
- [63] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 4
- [64] Yi Wang, Conrad M Albrecht, Nassim Ait Ali Braham, Lichao Mou, and Xiao Xiang Zhu. Self-supervised learning in remote sensing: A review. *arXiv preprint arXiv:2206.13188*, 2022. 3
- [65] Yi Wang, Nassim Ait Ali Braham, Zhitong Xiong, Chenying Liu, Conrad M Albrecht, and Xiao Xiang Zhu. Ssl4eos12: A large-scale multimodal, multitemporal dataset for self-supervised learning in earth observation [software and data sets]. *IEEE Geoscience and Remote Sensing Magazine*, 11(3):98–106, 2023. 3
- [66] Jiannan Wu, Muyan Zhong, Sen Xing, Zeqiang Lai, Zhaoyang Liu, Zhe Chen, Wenhai Wang, Xizhou Zhu, Lewei Lu, Tong Lu, et al. Visionllm v2: An end-to-end generalist multimodal large language model for hundreds of vision-language tasks. *Advances in Neural Information Processing Systems*, 37:69925–69975, 2025. 3
- [67] Xinyu Bai and Feng Xu. Accelerating Diffusion for SAR-to-Optical Image Translation via Adversarial Consistency Distillation, 2024. Online: <http://arxiv.org/pdf/2407.06095>. 17
- [68] Zhitong Xiong, Yi Wang, Fahong Zhang, Adam J. Stewart, Joëlle Hanna, Damian Borth, Ioannis Papoutsis, Bertrand Le Saux, Gustau Camps-Valls, and Xiao Xiang Zhu. Neural plasticity-inspired foundation model for observing the earth crossing modalities, 2024. 3
- [69] Lingxiao Yang, Ru-Yuan Zhang, Yanchen Wang, and Xiaohua Xie. Mma: Multi-modal adapter for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23826–23837, 2024. 3
- [70] Qidong Yang, Jonathan Giezendanner, Daniel Salles Civitarese, Johannes Jakubik, Eric Schmitt, Anirban Chandra, Jeremy Vila, Detlef Hohl, Chris Hill, Campbell Watson, et al. Multi-modal graph neural networks for localized off-grid weather forecasting. *arXiv preprint arXiv:2410.12938*, 2024. 3
- [71] Zhiping Yu, Chenyang Liu, Liqin Liu, Zhenwei Shi, and Zhengxia Zou. Metaearth: A generative foundation model for global-scale remote sensing image generation. *IEEE Trans-*

actions on Pattern Analysis and Machine Intelligence, 2024. [3](#)

- [72] Xiaohui Yuan, Jianfang Shi, and Lichuan Gu. A review of deep learning methods for semantic segmentation of remote sensing imagery. *Expert Systems with Applications*, 169: 114417, 2021. [2](#)
- [73] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600-612, 2004. [17](#)
- [74] Jingyi Zhang, Jiaying Huang, Sheng Jin, and Shijian Lu. Vision-language models for vision tasks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. [3](#)
- [75] Linying Zhao and Shunping Ji. Cnn, rnn, or vit? an evaluation of different deep learning architectures for spatio-temporal representation of sentinel time series. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16:44–56, 2022. [2](#)
- [76] Xiao Xiang Zhu, Devis Tuia, Lichao Mou, Gui-Song Xia, Liangpei Zhang, Feng Xu, and Friedrich Fraundorfer. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE geoscience and remote sensing magazine*, 5(4):8–36, 2017. [2](#)

TerraMind: Large-Scale Generative Multimodality for Earth Observation

Supplementary Material

In the following, we provide additional information on our data, the pretraining of TerraMind and its tokenizers, the quality of the tokenization, any-to-any generation matrices, and comparisons of TerraMind in unimodal and multimodal finetuning against specialized U-Net and ViT models.

7. TerraMesh Dataset

All versions of TerraMind have been pretrained on TerraMesh or a subset of it. TerraMesh is a comprehensive multimodal Earth observation dataset designed for large-scale model pre-training. It will be made publicly available under a permissive license in a preprint during the review process of this paper. The dataset includes nine modalities and we visualize examples of the dataset in Figure 8.

The dataset contains over 9 million globally distributed, spatiotemporally aligned samples across nine core modalities. Each modality is precisely co-registered at a 10-meter resolution, primarily based on Sentinel-2 grids. The S-1 and S-2 samples are sourced from MajorTOM-Core [23] and SSL4EO-S12 v1.1 [6]. It integrates Sentinel-1 SAR data with Sentinel-2 optical data (L1C top-of-atmosphere and L2A bottom-of-atmosphere reflectance), ensuring versatility for various downstream tasks. Because the source datasets contain only one S-1 product, each sample has either S-1 GRD or S-1 RTC data. Additionally, TerraMesh includes normalized difference vegetation index (NDVI) maps derived from Sentinel-2, Copernicus digital elevation model (DEM) data providing topographic context, and land-use/land-cover (LULC) maps from ESRI, enhanced with accurate cloud masks generated by the SEnSeI v2 model[22].

To ensure broad geographic and thematic diversity, TerraMesh employs subsampling techniques, selectively including representative samples from each global ecoregion and land-cover class, while downsampling highly homogeneous regions such as deserts and tundra. Another critical aspect is the data preprocessing pipeline, which includes reprojection, temporal alignment, and filtering to minimize missing data and artifacts, ensuring high-quality, analysis-ready samples.

TerraMind.v1-B-single was pre-trained on a subset of TerraMesh with one million samples, specifically the SSL4EOS12 v1.1 locations, using only four image modalities: S-2 L2A, S-1 GRD, DEM, and LULC. Additionally, we performed continuous pre-training with image captions. These captions were created using LLaVA-Next [37] and Overture Maps data. The automated captioning pipeline includes a prompt with a chain-of-thought process to generate diverse captions. The captioning model is asked to generate three question-answer pairs and describe the full image later.

We use the S-2 RGB bands and Overture base layer tags as inputs. Domain experts evaluated a subset of 1.3k captions, resulting in 69% of the captions without any hallucinations while the average completeness scores were 3.87 on a scale from 0 to 5. The captions will be released as an additional preprint during the review process of this paper.

8. Pretraining details

In this section, we give additional details on the pretraining of both TerraMind and its tokenizers.

8.1. Tokenizer models

The tokenizer models are pretrained using a Vision Transformer (ViT) encoder and a patched UNet decoder, with input images ranging from 224x224 to 256x256 in size. The model was trained with patch sizes of 16x16 for the ViT encoder and 4x4 for the UNet decoder. A tanh MLP was used before the quantizer, as outlined in the ViT-VQGAN paper, to enhance tokenization quality.

The model utilized a Finite-Scalar Quantization (FSQ) approach with a codebook size of 8-8-8-6-5, aiming to learn consistent and abstract representations across image patches. The latent dimension was set to 5. We leverage the normalization of codebook entries to the unit sphere during training. This concept is borrowed from the ViT-VQGAN approach, which applies a specific form of normalization to improve the quality and efficiency of learned representations. Additionally, an EMA-based quantizer was used with a decay rate of 0.99 to track and improve quantization over time.

During diffusion-based pretraining, the model was trained for 1000 timesteps using a linear beta schedule, with MSE loss as the objective. The training leveraged half-precision (fp16) and used an AdamW optimizer with specific learning rate scheduling and warmup strategies. The model also incorporated model EMA for stable training and set a batch size of 1 per GPU with various regularization techniques like grad clipping and random horizontal flips.

We pretrained the TerraMind tokenizers for image-like modalities with DDP on 4 GPUs for a total of 100 epochs on the respective modality of TerraMesh. We use a base learning rate of 1e-4, an effective batch size of 64 samples per GPU, i.e. the global batch size is 256. We reach a GPU utilization of 99% for single channel modalities like LULC and NDVI, and over 80% for all multi-channel modalities.

8.2. TerraMind

We pretrained both TerraMindv1-B and TerraMindv1-L with DDP on 32 GPUs. We determine the global batch size based

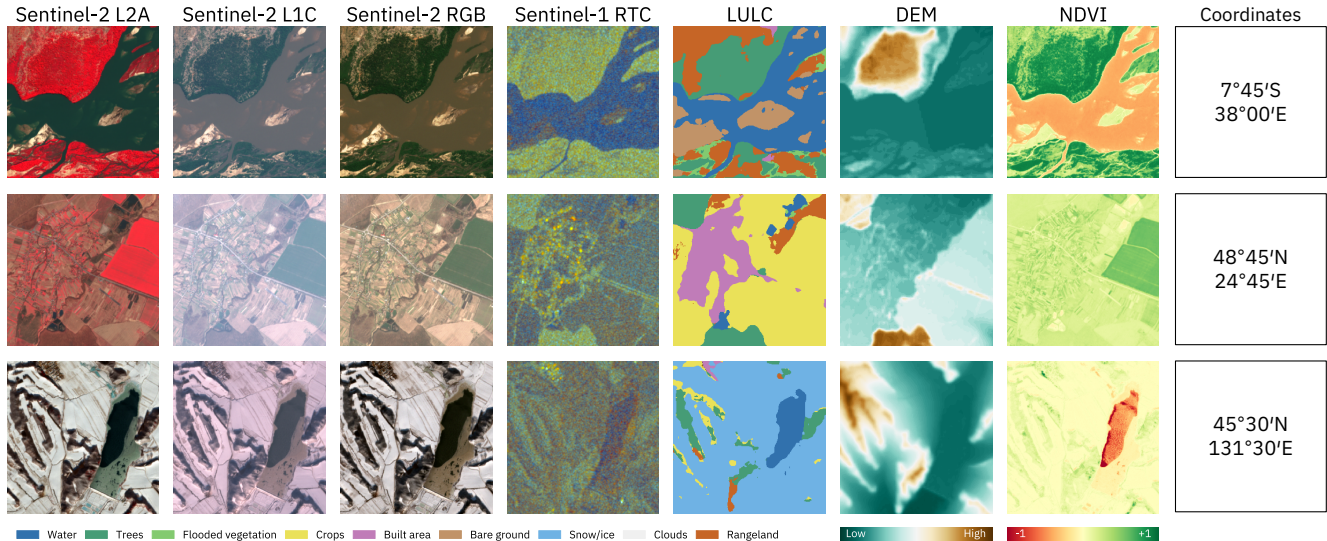


Figure 8. Visualization of the spatial-temporal alignment across modalities in TerraMesh. S-2 L2A uses IRRG pseudo-coloring and S-1 RTC is visualized in db scale as VH-VV-VV/VH. Copernicus DEM is scaled based on the image value range.

on initial experimental runs comparing a global batch size of 2K, 4K, and 8K. In addition, we determine the base learning rate starting from $1e-4$ and iteratively experimented with half and double learning rates. Ultimately, we end up with a base learning rate of $2e-4$ for a cosine annealing scheduler set to run for 500B tokens. For the v1-L model, we reach a GPU utilization of 85+%. Overall, the training of TerraMindv1-B took 12 days on 32 A100 GPUs, i.e., 9’216 GPU hours. Over the course of the pretraining, we also experiment with different configurations of the Dirichlet sampling distribution. In total, the pretraining experiments have been approximately three times larger than the final runs resulting in approximately 30K GPU hours allocated for pretraining.

We provide an overview on the scaling dynamics when going from TerraMindv1-B to TerraMind v1-L in Figure 9 with identical hyperparameters and compute. Overall, as expected, we observe a significant gap in the validation losses across modalities. We finally provide the validation losses per modality after pretraining of TerraMindv1-B and TerraMindv1-L in Table 9.

Model	S-2 L2A	S-1 GRD	S-1 RTC	DEM	NDVI
Random	9.68	9.68	9.68	9.68	9.68
V1-B	5.67	7.84	7.64	2.19	6.42
V1-L	5.34	7.69	7.53	2.14	6.25

Table 9. Validation losses of full pre-training of TerraMindv1-B and v1-L.




Figure 9. Example of the scaling behavior of TerraMind comparing v1-B and v1-L models for the first 350B tokens on the validation loss of optical S-2 L2A data. Overall, TerraMind-L outperforms TerraMind-B after approximately 10% of the training schedule of the large model.

9. Tokenizer performance and general learnings

In the following, we provide details on the tokenizations of TerraMind. At least for image-like modalities, the tokenizations represent an important and computationally heavy phase of the pretraining, which is why we highlight important learnings in the following.

Learnings. Overall, we learned that the tokenizer performance can be quite sensitive, which is especially related

to the significant bottleneck compression of up to 3000x after the encoder. When leveraging finite-scalar quantization (FSQ) instead of vector quantization (VQ), we observed exactly what the original FSQ paper [51] claims: FSQ makes quantization easier – yet in our experiments it did not improve the reconstruction performance in terms of MSE losses. We leverage FSQ as the training was more stable and less sensitive to the learning rate, which is likely related to the fact that, unlike VQ, FSQ does not require an additional codebook loss. We still observed that all tokenizer models were sensitive to the learning rate, with higher learning rates resulting in non-differentiability (NaN losses), and low learning rates caused blurry results.

In addition, we experimented with the codebook size. In our experiments, we observed that the level of detail in the reconstructions was significantly higher for single channel input compared to multi channel input (e.g., 12 band S2-L2A data). Naturally, with less channels, the compression bottleneck for equal-sized codebooks is lower. Therefore, we hypothesized whether multi-spectral data requires larger codebook sizes to obtain higher level of detail in the reconstructions. In contrast to our expectation, when increasing the codebook size over 16K for modalities with more than three input channels, the reconstructions had significant artefacts. This suggests that even though the compression bottleneck is lower, higher codebook sizes are more difficult for the model to use, which is in line with previous literature. However, we were surprised to see more artefacts in the reconstructions of models with a codebook size 32K compared to 16K.

Finally, we experimented with exponential moving average (EMA) updates for the tokenizer models. As expected, the models were less responsive to gradient updates. The resulting reconstructions smoothed out more of finegrained features. Together with the generative diffusion process in the tokenizer decoder, the resulting reconstructions often looked like hallucinations, e.g. bridges over rivers were not existing anymore in the reconstruction images. We therefore decided to omit exponential moving average in our tokenizer models.

9.1. FSQ vs. VQ

Generally, our pretraining experiments comparing FSQ with vector quantization suggest that both approaches can achieve the same level of performance, yet reaching optimal levels of performance with VQ is regarded to be more challenging than using FSQ. We visualize this through (a) the reconstruction loss and (b) the gradient norms of the tokenizer pretraining on S-2 L2A data in Figures 10 and 11, respectively. Overall, we observe that both approaches reach the same level of convergence, however FSQ requires less tuning and is generally more stable than VQ. This especially also applies for the grad norms.

Performance. In the following, we assess the accuracy of




Figure 10. Pretraining reconstruction losses of S-2 L2A modality comparing finite-scalar quantization (FSQ) and vector quantization (VQ) approaches. Overall, both approaches reach the same level of performance. The FSQ approach converges smoother than VQ, while requiring less tuning.




Figure 11. Gradient norms for pretraining of S-2 L2A tokenizers comparing finite-scalar quantization (FSQ) and vector quantization (VQ) approaches. The FSQ approach converges smoother than VQ, while requiring less tuning.

our tokenizer models. Besides visual quality assessments and quantitative assessments with MSE metrics, we were particularly interested in whether our tokenizers exhibit geospatial biases. Understanding this is crucial to ensure TerraMind has a uniform level of performance across the globe. In addition, we investigate the reconstructions of radar data in more detail, as radar data by nature includes significant noise in the amplitude data. This could interfere with the noise generation in the diffusion process of the decoder, which is why we assess the structure of the reconstructions using SSIM and PSNR metrics.




Figure 12. Spatial distribution of mean squared errors of the S-1 tokenizer on the validation set of the pretraining data.




Figure 14. Spatial distribution of mean squared errors of the DEM tokenizer on the validation set of the pretraining data.

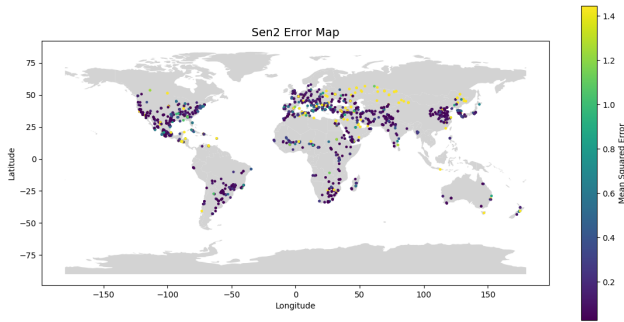


Figure 13. Spatial distribution of mean squared errors of the S-2 tokenizer on the validation set of the pretraining data.

In Figures 12 to 14, we provide an overview on the spatial distributions of the S-1 GRD, S-2 L2A, and DEM tokenizer on the validation data of the SSL4EO-S12 subset which is focused on urban areas and therefore relevant for many downstream applications. Overall, we observe low MSE errors and particularly low deviation across geographic regions. For optical S-2 data, we observe minor difficulties in reconstructing images from Northern Asia, which we manually investigated. Overall, the vast majority of those samples are depicting snowy/icy conditions that have very high reflectance values of up to 12,000 compared to a normal range of [0, 255] in RGB data. On those long tail distribution samples, the S-2 tokenizer naturally has more difficulties.

S1-tokenizer quantitative analyses. In the following, we pay particular attention to the performance of the radar S-1 tokenizer, which might be more challenging to train on a reconstruction task due to the inherent speckle noise in radar satellite data. We therefore evaluate the reconstructions of the S-1 tokenizer using the structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR). Both input and reconstruction for S-1 are in a dB scale. In addition to S-1 evaluation metrics being computed in the dB space in Table 10, they also are calculated in the denormalized space. On the contrary, the S-2 evaluation metrics are computed in the normalized space.

We give a more extensive background on radar data in the following for interested readers and non-EO experts. Reconstructing realistic and accurate synthetic aperture radar (SAR) S-1 VV and VH data is challenging due to factors inherent in the specific characteristics of SAR and the S-1 mission. SAR data is affected by complex interactions between the radar signal and Earth’s surface. SAR is based on radar backscatter, which is influenced by surface roughness and moisture content. The interaction of radar waves with different surfaces, including vegetation structure and urban environments, can produce complex backscatter patterns. The two polarizations, VV and VH, capture different scattering mechanisms: VV is sensitive to surface roughness and vegetation, while VH captures cross-polarized interactions that are influenced by surface and volumetric features [14, 35, 56]. In addition, SAR inherently contains speckle noise, which obscures fine details, making it difficult to extract accurate information. To evaluate the SAR data tokenizers of TerraMind, we employ various evaluation metrics to assess quality and accuracy. We compute the MAE and RMSE for quantifying pixel-level differences, the SSIM to compare image structural content, and the PSNR [1, 67, 73].

Table 10 presents the quantitative evaluation of the TerraMind tokenizer reconstructions across multiple modalities. The results show a reasonable reconstruction performance for optical data, indicating both structural and perceptual fidelity. For radar modalities, S-1 GRD and S-1 RTC achieve comparable PSNR values, though SSIM scores are lower, suggesting that while the reconstructions are visually plausible, they exhibit moderate structural deviations. In addition to these quantitative metrics, we also conducted qualitative assessments through visual inspection to identify artifacts and inconsistencies not captured by numerical scores alone.

10. Additional experiments

In the following, we provide additional experiments, especially with regard to the quality of the latent space and the full finetuning performance. To understand the quality of the

Modality	MAE	RMSE	SSIM	PSNR
S-1 GRD	2.403	3.220	0.565	30.291
S-1 RTC	2.216	2.888	0.466	30.389
S-2 L2A	0.055	0.134	0.851	27.439
DEM	170.7	737.2	0.974	20.712
NDVI	0.091	0.168	0.647	21.517

Table 10. Evaluation of SAR VV and VH and S-2 reconstructions by the TerraMind tokenizers using MSE ↓, SSIM ↑ and PSNR ↑ on the validation dataset of the SSL4EO-S12 subset (8.5k samples).

latent space, we compute performances of nearest neighbor approaches for image classification tasks or using prototypical neural networks. We assess the performance of full fine-tuning by comparing with end-to-end trained, task-specific models like U-Nets and ViTs. We additionally compare the quality of the generations with the pseudo-labels used to pretrain TerraMind in an ablation experiment in a zero-shot setup.

10.1. Geolocation prediction

To better understand how TerraMind assigns geolocations, we further employ a Monte-Carlo sampling on the latitude-longitude grid for an optical tile from the validation data in Figure 15. We observe that while TerraMind is not predicting the correct geolocation (●), there is a very high likelihood that the predicted geolocation is one of the adjacent grid points that have been seen during pretraining (●). This result suggests that even for data from unseen geolocations, TerraMind remembers similar samples from the pretraining data (●) and returns the geolocation of the samples with high similarity. This capability paired with the global pretraining of TerraMind suggests that geo-localization of data from unseen locations is possible but determined by the similarity to images from adjacent locations.




Figure 15. Distribution of predicted geo-locations for an optical S-2 L2A sample from the validation set. ● is the correct location, ● are Monte-Carlo sampled locations from TerraMind, ● represents the distribution of training locations. TerraMind’s geo-localization seems to be based on similar optical samples in the training dataset for which TerraMind then outputs the geolocation.

We further extend the analysis of Figure 7 by additionally prompting the model for likely locations of urban areas.

Overall, we observe that the model correctly identifies many densely populated areas across the globe. We also note over-predictions in, for example, North Africa and middle-east. This observation suggests that the model might confuse bare land and urban areas in these regions.




Figure 16. Prediction distribution of the land use class “urban” with a sampling temperature of $T = 1.0$. TerraMind has a reasonable internal representation of the geolocation of specific contexts, like land use classes.

10.2. Few-shot experiments

We present additional few-shot experiments with the EuroSAT and METER-ML dataset in Table 11. We use the embeddings of the pre-trained encoders without any additional fine-tuning. The patch embeddings of each image are averaged for image-level classification tasks.

The experiments include four different few-shot settings with varying numbers of examples and classes. 5-way refers to sampling five classes per run, while full-way describes experiments with all dataset classes per run. 1-shot and 5-shot indicate that one or five images are sampled for each class per run. 5-shot experiments with five support samples per class are using Prototypical Networks [60] for classification. This approach averages the embeddings of the selected labeled images (support set) and classifies the target images (query set) based on the class prototype with the lowest Euclidean distance from each sample. In the 1-shot setting, Prototypical Networks are mathematically equal to 1-Nearest-Neighbor classification. We refer to the original paper for details [60]. Different from literature, we evaluate each run on the full test set instead of subsampling query images.

TerraMind performs best on both datasets, outperforming all other geospatial foundation models as well as the CLIP vision encoder [57]. Interestingly, the base version leads to overall better results than the large model. Similarly, Prithvi’s smaller 1.0 version has comparable results to its larger 2.0 300M version, indicating that model size has only a limited effect on few-shot performance.

In addition to S-2 L1C, the METER-ML dataset provides high resolution RGB images from NAIP with 1 m resolution. Only CLIP and TerraMind can process RGB images without any fine-tuning. While CLIP profits largely from the higher resolution inputs, TerraMind only performs marginally better

Model	Input	EuroSAT				METER-ML			
		5-way 1-shot	5-way 5-shot	full-way 1-shot	full-way 5-shot	5-way 1-shot	5-way 5-shot	full-way 1-shot	full-way 5-shot
CLIP-ViT-B/16	S-2 RGB	57.00	70.72	43.92	58.30	29.15	37.44	23.13	30.53
CLIP-ViT-B/16	NAIP	–	–	–	–	32.01	42.35	25.66	35.81
DeCUR	S-2 L1C	50.54	64.35	37.53	50.82	27.87	33.64	20.95	27.21
Prithvi 1.0 100M	S-2 L1C	60.11	73.29	46.86	60.66	26.08	35.81	22.33	29.21
Prithvi 2.0 300M	S-2 L1C	61.06	73.21	47.47	60.47	28.26	36.13	22.52	29.59
TerraMindv1-B	S-2 L1C	70.83	87.94	57.48	79.66	33.90	43.89	26.85	<u>37.41</u>
TerraMindv1-B	NAIP	–	–	–	–	32.23	<u>44.75</u>	25.53	37.85
TerraMindv1-L	S-2 L1C	<u>70.07</u>	<u>86.29</u>	<u>56.58</u>	<u>77.39</u>	<u>33.09</u>	<u>42.72</u>	<u>26.02</u>	36.34
TerraMindv1-L	NAIP	–	–	–	–	32.59	44.99	25.94	38.29

Table 11. Few-shot classification results on EuroSAT and METER-ML measured in mean accuracy \uparrow averaged over 200 runs. 5-way refers to five randomly sampled classes per run, which is a default setting used in few-shot learning. Full-way refers to sampling all dataset classes, i.e., ten EuroSAT classes and seven METER-ML classes. We highlight the best two models in bold and underlined.

and sometimes worse than with multispectral S-2 data. Notice that TerraMind shows similar performance gaps as CLIP when comparing NAIP data to S-2 RGB. This indicates that additional multispectral channels have a comparable effect on few-shot performance as high-resolution images.

10.3. Finetuning comparisons with baseline models

Since the first approaches to foundation models for Earth observations, experts in the field discuss on the usability of such models compared to task-specific models that are trained for each application individually. Recent benchmark results suggested that task-specific models, like U-Nets, often outperform finetuned GFM [49]. We therefore additionally investigate how TerraMind compares with task-specific U-Nets and ViT models following the PANGAEA evaluation protocol in Table 6. As advised by the authors of PANGAEA, we again report results on nine of the eleven datasets as we could not reproduce the performance on the remaining two datasets. The task-specific models are trained from scratch for each individual task, while all GFMs including TerraMind are finetuned with a frozen encoder and an UperNet head. For reasons of consistency, we report TerraMind results based on full finetuning instead of TiM tuning. Yet, for several tasks, we expect slightly higher results when applying TiM tuning for TerraMind. Overall, our results demonstrate that TerraMindv1-B outperforms task-specific UNet and ViT models across the PANGAEA benchmark in both unimodal and multimodal settings by 1pp avg. mIoU and 4pp avg. mIoU respectively. In multimodal settings, the improvement peaks to 4.5pp improvement of TerraMindv1-B over task-specific U-Nets. To the best of our knowledge, this is the first time a GFM model outperforms task-specific models on a global benchmark.

In addition, we observe that for most datasets, TerraMindv1-B outperforms TerraMindv1-B-single. This demonstrates the benefit from scaling in the data and feature

dimension—i.e., leveraging dual-scale feature representations on a pixel level and a token level.

10.4. Comparing generations and pseudo-labels

We evaluate the model generations for modalities where we used pseudo-labels as input data. For example, in initial experiments with TerraMind.v0.1, we leverage Google’s DynamicWorld model to pseudo-label LULC maps which we use as input to the model. In the following experiment in Table 12, we test the performance of the DynamicWorld model against the generations of TerraMind. Overall, we observe that while finetuned TerraMind.v0.1-B outperforms DynamicWorld, the generation of TerraMind does not surpass the inference results of DynamicWorld.

Approach	Input	IoU _{Water}
TerraMindv1-B-single	S-2 L1C	69.87
Dynamic World pseudo-labeling	S-2 L1C	71.98
TerraMindv1-B-single finetuning	S-2 L1C	76.32

Table 12. Results on the Sen1Floods11 test set comparing flood maps derived from TerraMind’s out-of-the-box LULC generations to those derived from LULC pseudo-labeling with Dynamic World. The results are inferior to those obtained by fine-tuning a specialized model for this downstream task, which is expected.

10.5. TiM tuning for crop mapping

We further investigate the relevance of TiM tuning for crop type mapping in order to understand the relevance of generating artificial data for more finegrained segmentation tasks. That means, we generate artificial LULC data which includes agricultural crop as a *single class* and investigate whether this additional information helps to segment nine different types of crops in satellite images. We experiment

with the South Africa Crop Type Mapping dataset (<https://source.coop/esa/fusion-competition>) and present the results in Table 13. Overall, we observe that TiM tuning improves the performance by around 1pp. That means that even though the generated artificial data does not include further information on the location and shape of certain crops, the information on where to expect crop land in general helps to guide the model to an improved performance.

	Input	mIoU
TerraMindv1-B	S-2	41.87
TerraMindv1-B TiM	S-2 + <i>gen. LULC</i>	42.74

Table 13. Thinking-in-modalities (TiM) tuning compared with standard full fine-tuning approaches on the SA Crop dataset.

11. Any-to-any generation

In Figure 18, we provide an example of any-to-any generation on four image-like modalities and two sequence-like modalities. Overall, we observe that when we start from modalities with high information content (e.g., fine-grained image-like modalities), the reconstructions are particularly good. Even with less information content, the model is able to generate consistent artificial data. However, we can clearly observe that the quality compared to the ground truth (represented by the input in the left of the figure) is decreasing. Finally, it is interesting to see how artefacts are introduced by the model when starting from lower information content in the input. For example, when prompting TerraMind to generate data from DEM input, we observe that the model pays significant attention to the darker streams in the DEM image, which are later generated as a river in LULC.

While we expect to see accurate generations from information-rich modalities like optical data, it is particularly interesting to understand how TerraMind deals with low information content. Therefore, we prompt TerraMind to generate a subset of modalities starting from the geolocation in Figure 17. Interestingly, for a geolocation from the middle-east, the model generates an optical image that resembles a desert. While the generated optical image is based on the right context, the actual structure is unsurprisingly different from the ground truth. Based on the chained generation, this difference ripples down across all other modalities as well causing consistent but inaccurate generations. This example emphasizes the relevance of access to information-rich, fine-grained features to facilitate accurate generations.

Next to the evaluation of raw, pixel-level input in Table 3, we further evaluate the generation quality using tokenized input in Table 14. Interestingly, we observe only minor reduction in performance compared to pixel-level input even




Figure 17. Randomly selected chained generation example with uni-modal geo-location input data. Top row is artificially generated data by TerraMind, bottom row represents a ground truth sample at this grid location, respectively.

though the tokenized representations are compressed significantly (up to 3000x for S-2 L2A). Overall, our results suggest that leveraging tokenized inputs can be a reasonable alternative to leveraging pixel-level data for the generation of artificial data with TerraMind.

11.1. Large-scale generations

In Figures 19 and 20, we provide additional qualitative results for large-tile generations at the example of Singapore. Specifically, we leverage a 35.5km \times 69.5km optical S-2 L2A tile as input and iteratively generate overlapping 224x224 pixel generations for S-1 RTC, S-1 GRD, NDVI, and LULC. In the overlapping areas, we apply the mean of all generations in order to enhance the spatial conciseness of the generations. TerraMind consistently removes the clouds in the S-1 generations. It makes assumptions for hidden areas, which are look accurate for large features like water bodies or the shore line. Other features like airports or ships are also clearly visible in the S-1 and NDVI generations.

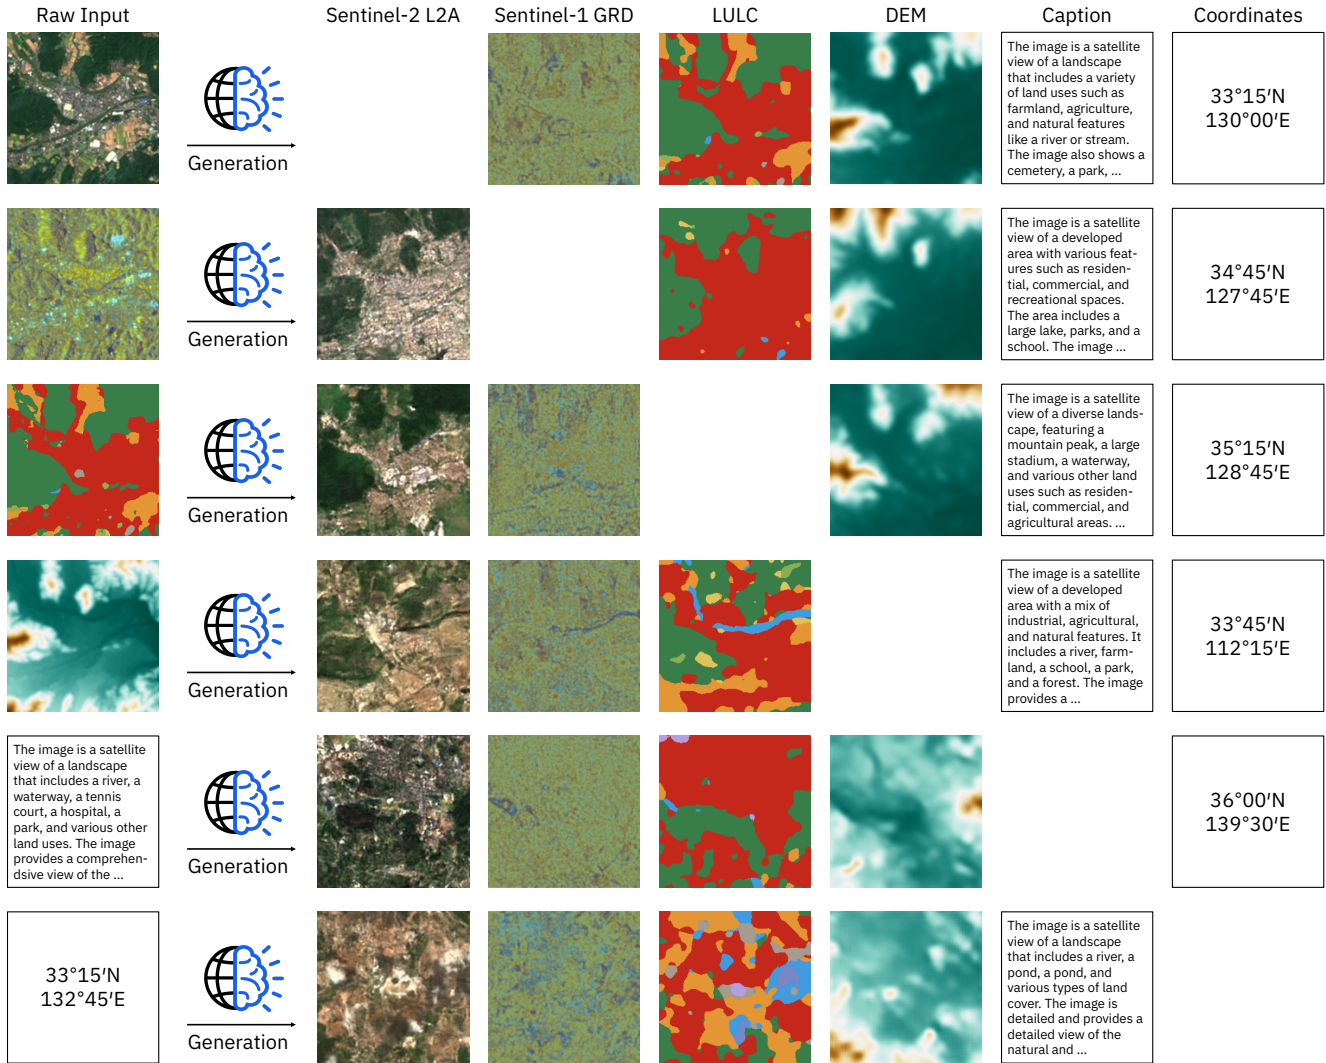



Figure 18. Any-to-any generation example of TerraMindv1-B-single. Fine-grained input like optical and radar achieve particularly good performances.

Modalities	MAE	RMSE	SSIM	PSNR
Tokenized S-2 L2A → S-1 GRD	3.3180	4.3309	0.5131	27.715
Tokenized S-2 L2A → S-1 RTC	3.0544	3.9178	0.4131	27.739
Tokenized S-2 L2A → DEM	572.5	1040.6	0.5728	17.718
Tokenized S-1 GRD → S-2 L2A	0.0820	0.1238	0.7182	25.630
Tokenized S-1 GRD → NDVI	0.1949	0.2425	0.4124	18.324
Tokenized S-1 GRD → DEM	327.4	550.3	0.7271	16.008
Tokenized S-1 RTC → S-2 L2A	0.1195	0.1935	0.6638	24.266
Tokenized S-1 RTC → NDVI	0.1895	0.2348	0.4500	18.606
Tokenized S-1 RTC → DEM	457.9	851.6	0.7095	19.457

Table 14. Performance of TerraMind on tokenized inputs using 10 diffusion steps. Metrics include MAE ↓, RMSE ↓, PSNR ↑, and SSIM ↑.




(a) Input: S-2 L2A data from Singapore captured January 9th, 2025.



(b) Generation: TerraMind output for S-1 composition

Figure 19. Large-tile generations of TerraMind for Singapore (1/1)




(c) Generation: TerraMind output for LULC

Figure 19. Large-tile generations of TerraMind for Singapore (2/2)




(a) Input: S-2 L2A data from Santiago de Compostela.




(b) Generation: TerraMind output for S-1 GRD composition

Figure 20. Large-tile generations of TerraMind for Santiago de Compostela (1/3)




(c) TerraMind generation for S-1 RTC composition



(d) Generation: TerraMind output for vegetation

Figure 20. Large-tile generations of TerraMind for Santiago de Compostela (2/3)



(e) Generation: TerraMind output for digital elevation

Figure 20. Large-tile generations of TerraMind for Santiago de Compostela (3/3)