



Article

Segment Anything Model Can Not Segment Anything: Assessing AI Foundation Model's Generalizability in Permafrost Mapping

Wenwen Li ^{1,*} , Chia-Yu Hsu ¹ , Sizhe Wang ^{1,2}, Yezhou Yang ², Hyunho Lee ¹, Anna Liljedahl ³, Chandi Witharana ⁴, Yili Yang ³, Brendan M. Rogers ³, Samantha T. Arundel ⁵, Matthew B. Jones ⁶, Kenton McHenry ⁷ and Patricia Solis ¹

¹ School of Geographical Sciences and Urban Planning, Arizona State University, Tempe, AZ 85287, USA; chsu53@asu.edu (C.-Y.H.); wsizhe@asu.edu (S.W.); hlee401@asu.edu (H.L.); patricia.solis@asu.edu (P.S.)

² School of Computing and Augmented Intelligence, Arizona State University, Tempe, AZ 85287, USA; Yz.Yang@asu.edu

³ Woodwell Climate Research Center, Falmouth, MA 02540, USA; aliljedahl@woodwellclimate.org (A.L.); yyang@woodwellclimate.org (Y.Y.); brogers@woodwellclimate.org (B.M.R.)

⁴ Department of Natural Resources and the Environment, University of Connecticut, Storrs, CT 06269, USA; chandi.witharana@uconn.edu

⁵ U.S. Geological Survey, Center of Excellence for Geospatial Information Science, Rolla, MO 65401, USA; sarundel@usgs.gov

⁶ National Center for Ecological Analysis & Synthesis, University of California, Santa Barbara, CA 93106, USA; jones@nceas.ucsb.edu

⁷ National Center for Supercomputing Applications (NCSA), University of Illinois at Urbana, Champaign, IL 61820, USA; mchenry@illinois.edu

* Correspondence: wenwen@asu.edu



Citation: Li, W.; Hsu, C.-Y.; Wang, S.; Yang, Y.; Lee, H.; Liljedahl, A.; Witharana, C.; Yang, Y.; Rogers, B.M.; Arundel, S.T.; et al. Segment Anything Model Can Not Segment Anything: Assessing AI Foundation Model's Generalizability in Permafrost Mapping. *Remote Sens.* **2024**, *16*, 797. <https://doi.org/10.3390/rs16050797>

Academic Editors: Xiaowei Jia, Kuishuang Feng and Yiqun Xie

Received: 16 January 2024

Revised: 17 February 2024

Accepted: 20 February 2024

Published: 24 February 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract: This paper assesses trending AI foundation models, especially emerging computer vision foundation models and their performance in natural landscape feature segmentation. While the term foundation model has quickly garnered interest from the geospatial domain, its definition remains vague. Hence, this paper will first introduce AI foundation models and their defining characteristics. Built upon the tremendous success achieved by Large Language Models (LLMs) as the foundation models for language tasks, this paper discusses the challenges of building foundation models for geospatial artificial intelligence (GeoAI) vision tasks. To evaluate the performance of large AI vision models, especially Meta's Segment Anything Model (SAM), we implemented different instance segmentation pipelines that minimize the changes to SAM to leverage its power as a foundation model. A series of prompt strategies were developed to test SAM's performance regarding its theoretical upper bound of predictive accuracy, zero-shot performance, and domain adaptability through fine-tuning. The analysis used two permafrost feature datasets, ice-wedge polygons and retrogressive thaw slumps because (1) these landform features are more challenging to segment than man-made features due to their complicated formation mechanisms, diverse forms, and vague boundaries; (2) their presence and changes are important indicators for Arctic warming and climate change. The results show that although promising, SAM still has room for improvement to support AI-augmented terrain mapping. The spatial and domain generalizability of this finding is further validated using a more general dataset EuroCrops for agricultural field mapping. Finally, we discuss future research directions that strengthen SAM's applicability in challenging geospatial domains.

Keywords: foundation model; artificial intelligence; mapping; zero-shot; segmentation; GeoAI

1. Introduction

Recent advances in large language models (LLM), such as OpenAI's GPT (Generative Pre-Trained Transformer), have brought a new wave of important changes to AI. LLMs

are trained on millions of web documents (as natural language text) to predict the next or missing words in a sentence, which can be accomplished in a self-supervised manner [1]. These models adopt new architectures and cutting-edge learning strategies (e.g., transformer and multi-head attention) [2] and are capable of extracting language structures and patterns, enabling them to handle multiple new tasks, such as question answering and code debugging, for which they were not originally trained. As a result, LLMs have gained a dramatic capability in generalization and language understanding, and they can serve as the foundation for a wide variety of downstream tasks. This is why the term “foundation model” was coined in 2019.

Embracing foundation models for environmental and social science research has garnered significant interest from the geospatial community. As an emerging concept and exciting new development, foundation models offer the prospect of reducing model development time for individual researchers and gaining the capability needed to analyze the ever-increasing amount of geospatial data. However, in comparison to the research progress in constructing foundation models for natural language processing (NLP), the advances of vision foundation models are catching up in speed. This is because of the different learning goals between language and vision tasks. Many of the NLP tasks can be formulated as text-to-text processing, which involves taking natural language as input and providing natural language response as output [3]. When the training data reach a certain size, the LLMs (e.g., ChatGPT) are able to gain incredible generalizability and can be adapted to a diverse set of NLP tasks even without fine-tuning. The success of ChatGPT provides evidence of the power of such LLMs. In contrast, not all computer vision models can be designed as generative models. The goal for vision tasks varies depending on the required granularity in image analysis, be it for visual question answering, image reconstruction, image-level classification, object detection, or instance segmentation. Due to this variability, it is difficult to derive a generalized model for vision tasks, and the emerging vision foundation models, therefore, remain task-specific.

In the permafrost mapping domain, increasing research has harnessed the power of AI to analyze massive amounts of satellite imagery, achieving high-resolution mapping of this rapidly changing landscape. A Mapping Application for Permafrost Land Environment (MAPLE) has been developed to process big imagery using AI and high-performance computing to map permafrost features, such as ice-wedge polygons, at a pan-Arctic scale [4]. The main AI tool used is the popular deep learning architecture Mask R-CNN. Li et al. developed a real-time deep learning model to segment permafrost features based on a very efficient deep learning architecture - SparseInst [5]. The model achieves predictive performance as good as Mask R-CNN but with a much faster inference speed. Yang et al. developed a semantic segmentation framework based on U-Net to segment thaw slumps, leveraging geospatial data from multiple sources [6]. All these studies have fostered the in-depth integration of AI in permafrost mapping. However, the AI architectures used are all based on supervised learning, requiring substantial computing resources for model training and fine-tuning to achieve satisfactory performance. Model development also necessitates expert knowledge in AI, increasing the barriers to large-scale adoption of AI tools for Arctic researchers.

The emergence of AI foundation models has shown great potential to address the above challenges, as these models are pre-trained on large datasets, achieving a high level of generalizability and domain adaptability. Several models, such as Meta’s Segment Anything Model (SAM), are equipped with the capability for zero-shot learning, allowing the model to directly segment features of interest without additional training effort. Hence, in this research, we aim to assess the strengths and weaknesses of Meta’s SAM in its capacity to support GeoAI vision tasks for permafrost mapping. This model was chosen as it is the first publicly-released large vision model for image segmentation and one of the few that are open-sourced and allow for model adaptation using geo-domain data. Two challenging permafrost datasets, ice-wedge polygons, and retrogressive thaw slumps, were chosen for evaluating SAM’s performance in zero-shot prediction, knowledge-embedded learning, as

well as its prediction accuracy with model integration and fine-tuning on SAM. The results are compared with those from a cutting-edge model, MViTv2 [7], based on supervised learning and multi-scale, transformer-based architecture.

The rest of the paper is organized as follows: Section 2 reviews the growing list of vision foundation models, their model architecture, and supporting vision tasks. Section 3 describes the datasets and a series of strategies we developed to evaluate SAM's instance segmentation performance, as well as the results. Section 4 summarizes our findings and discusses the strengths and limitations of SAM for AI-augmented permafrost mapping. The extension of this workflow to other geospatial problems, such as agricultural field mapping, is also discussed. Section 5 concludes the paper and proposes future research directions.

2. AI Foundation Models for GeoAI Vision Tasks: Recent Progress

As Li and Hsu [8] pointed out, the evolution of GeoAI may be seen in three phases: (I) import—to bring AI technology into geography and apply it to a domain problem; (II) adaptation—to develop problem-specific strategies to improve general AI models to achieve better performance in various geospatial tasks; (III) export—the integration of geospatial principles and knowledge back to AI to develop innovative models to help better solve both geospatial and aspatial problems. Exciting progress has been made in the past few years, and the field is moving quickly beyond Phase I toward Phase II and III [9,10]. For adopting new technology, such as foundation models, we still need to go through these three phases, and the study of GeoAI foundation models is clearly at the initial, exploration phase (Phase I).

Excitingly, there has already been preliminary research toward developing foundation models in the field of remote sensing image analysis. For example, Cha et al. [11] reported a billion-scale foundation model tailored from the Vision Transformer (ViT) based model. It is constructed with 2.4 billion parameters and pre-trained on the MillionAID dataset [12]. MillionAID is annotated for image scene classification, with each training image assigned with a scene label, such as dry land, oil field, and sports land. Next, the model is further fine-tuned to support two downstream tasks: rotated object detection and semantic segmentation. The results show that the proposed large model performs better overall than other models, such as RetinaNet [13] and Masked Auto Encoding (MAE) [14], on the two remote sensing image processing tasks. It also confirms that when a model is trained on a larger, easier-to-retrieve dataset (MillionAID contains 1 million image scenes) and fine-tuned on other smaller datasets, its performance can be improved as the model becomes more “knowledgeable” by learning from bigger data.

In the computer vision field, several big technology organizations, such as Meta, OpenAI, and Microsoft, have put efforts into developing vision foundation models (Table 1). In 2021, OpenAI released CLIP (Contrastive Language-Image Pre-training) [15] as a transferable visual model. The model learned feature representation by matching pairs of text (captions) and images during the pre-training phase, comparing the similarity between the encoded text and image information. This way, if an image is described as “a picture of a dog”, the model will learn the representation of “dog” from the images through natural language supervision. Based on this information, CLIP can achieve zero-shot prediction on image scene classes by finding the textural category with the highest similarity to the image's content. CLIP can be adapted to support traditional vision tasks, such as image classification and video action recognition. The model's performance was tested on over 30 vision tasks and the results demonstrated CLIP's comparable zero-shot performance on large benchmark datasets, such as ImageNet and ObjectNet. However, CLIP's performance was less satisfying when analyzing images not included in its pre-training datasets. Hence, the limitation in the size and distribution of training data will directly affect performance in downstream tasks. Furthermore, CLIP cannot conduct fine-granularity tasks, such as object detection and instance segmentation.

Table 1. Comparison of data and model operation space for foundation models and (pre) foundation models. NLP: Natural language processing. CV: Computer vision. M: million, B: billion, T: trillion.

Organization	Model	Data Size	Model Parameter Space	Released Date	Task	Open Source?
Google	BERT	BooksCorpus (800 M words) and English Wikipedia (2500 M words)	110M (base) 340 M (large)	2018	NLP	✓
OpenAI	GPT-3	570 GB text data from various sources	175 B	2020	NLP	-
Google	GLaM	1.6 T tokens	1.2 T	2021	NLP	-
DeepMind	Gopher	300 B tokens	280 B	2021	NLP	-
Google	PaLM	780 B tokens	540 B	2022	NLP	-
OpenAI	GPT-3.5	300 B tokens	175 B	2022	NLP	-
Meta	LLaMA	1.4 T tokens	65 B	2023	NLP	✓
OpenAI	GPT-4	N/A	1 T	2023	NLP	-
OpenAI	CLIP	400 M image-text pairs	150 M	2021	NLP & CV	✓
OpenAI	DALL·E	250 M text-image pairs	12 B	2021	NLP & CV	-
Microsoft	Florence	900 M image-text pairs	893 M	2021	NLP & CV	-
Meta	SAM	11 M images with 1B masks	636 M	2023	CV	✓
Cha et al. [11]	ViT G12x4	1 M images	2.4 B	2023	CV	-
IBM-NASA	Prithvi	30-meter, 6-band satellite imagery over the continental US	100 M	2023	CV	✓

Microsoft has developed a computer vision foundation model named Florence, which is also trained on a large set of image-text pairs (900 million) collected from the Internet. Florence aims to achieve zero-shot transfer learning by expanding the representation from coarse to fine-granularity, from static images to dynamic scenes, and from RGB images to images with multiple modalities and channels. This way, Florence can be adapted to more vision tasks than CLIP, including object detection. Florence’s pre-training model uses two-tower pipelines, including a CLIP language encoder and a Swin Transformer-based image encoder to encode the image and text data. It uses a contrastive loss function that classifies all image and text description pairs that can be mapped to a unique caption as positive samples and the rest as negative samples. Florence also adopts advanced learning strategies, such as dynamic head, to achieve better adaptation capability for downstream tasks. The model was tested on over 30 datasets on multiple image analysis tasks, including image classification and object detection. The results show that Florence achieves better zero-shot transfer learning than other large visual models, including CLIP. On object detection, it also achieves state-of-the-art performance compared to other cutting-edge models. However, note that object detection with Florence is not zero-shot learning, and the model needs to be fine-tuned for optimal performance. Although promising, Florence is not open-sourced, rendering it difficult to access. Moreover, the model was trained on 512 NVIDIA-A100 GPUs for several days, making training and retraining the model an expensive computational process.

In April 2023, Meta AI released the Segment Anything Model (SAM), which, for the first time, has the power to perform zero-shot image segmentation for more challenging vision tasks. Although AI models such as CLIP and Florence enable vision tasks by associating images with their descriptive text to capture image-scene-level semantics, their ability to identify instance-level semantics within the scene remains weak. SAM enables this capability through a powerful image encoder-decoder-based architecture. A large transformer model is used as the image encoder, and at the decoding phase, the model requires a prompt from the user input to generate the object mask (or object boundary). As shown in Figure 1, the input prompt can be a point, a box, or some text to indicate targets of interest. SAM’s pre-training went through three stages: (1) supervised training on a small set of images with manually annotated instance masks; (2) semi-automated training by incorporating annotated masks that are ambiguous for SAM; and (3) fully

automated training on the entire 11 million image datasets. SAM's advantage is its ability to represent features by learning from a vast number of images, and it can also help accelerate the speed of object annotation for different segmentation tasks through user-machine interaction. Soon after its release, several integrated models [16] were developed to support downstream tasks by incorporating SAM's object segmentation ability. For example, SAM is connected with CLIP to predict each segmented mask's (object's) category, such as water. Figure 1 presents the SAM+CLIP integrated model architecture. The SAM's outputs (segmented masks) are fed to CLIP's input, and by providing CLIP a text prompt (e.g., water), CLIP will output all the masks that belong to "water" through text-image similarity matching based on the resultant embeddings from the two encoders. This way, instance segmentation can be achieved.

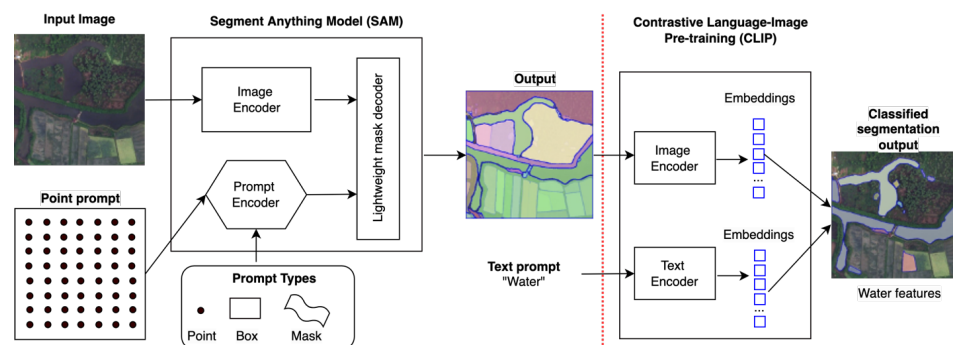


Figure 1. Architecture of SAM (left of the dashed line) and CLIP (right of the dashed line) and their combined workflow for instance segmentation.

Based on the above analysis and considering model maturity, capability, and open-sourceness, in this paper, we select the SAM model to assess its performance in segmenting natural features. Different from other works that apply SAM in remote sensing, our instance segmentation pipelines retain SAM's entire architecture instead of using some of its submodules, such as feature extraction backbones, in the development of the instance segmentation pipeline. This design maximizes the adoption of SAM as a foundation model that emphasizes easy reuse and requires minimal additional effort in model development.

In the next section, we describe the datasets used, our experimental design, and the various prompt strategies developed to comprehensively assess SAM's instance segmentation performance.

3. Materials and Methods

Kirillov et al. [17] evaluated SAM's performance on multiple vision tasks, including edge detection, object proposal generation, and instance segmentation, all in a zero-shot manner. The results show that, even if SAM was not trained for edge detection, it is capable of generating reasonable edge maps on a benchmark dataset BSDS500 [18]. This provides significant evidence for SAM's task adaptation and transfer learning abilities. For mid-level vision tasks, such as object proposal generation, SAM has exhibited remarkable zero-shot performance in segmenting medium and large objects, outperforming a strong benchmark model ViTDet-H [19] based on supervised learning. This experiment was conducted on a challenging vision dataset LVIS (Large Vocabulary Instance Segmentation) [20]. Regarding high-level vision tasks, although SAM's zero-shot segmentation accuracy is lower than the supervised model ViTDet-H, the gap is small (8.82% on the benchmark COCO dataset [21] and 4.08% on the LVIS dataset). Furthermore, its mask quality was rated higher than other supervised models based on human studies.

3.1. Datasets

While these studies demonstrate SAM's satisfying performance in general computer vision benchmark datasets, its domain adaptation capability for segmenting natural

landscape features remains unknown. To address this question, we have selected two natural feature datasets—ice-wedge polygons (IWP) and retrogressive thaw slumps (RTS)—for the assessment.

The first dataset is the AI-ready Ice-Wedge Polygon (IWP) data for instance-level segmentation. IWPs are ambiguous ground surface features found in permafrost-affected landscapes, specifically in regions with ice-rich permafrost. The type of IWP changes when the upper section of an ice-wedge melts, indicating the rate of Arctic warming [22]. The training dataset was selected based on tundra vegetation types within the cold continuous permafrost region, guided by the Circumpolar Arctic Vegetation Map (CAVM) [23]. These types represent major tundra vegetation types, including Wetland, Erect shrub, Prostrate shrub, Graminoid, and Barren, in the ice-rich permafrost regions. The very high-resolution satellite imagery at 0.5m resolution from commercial Maxar sensors is used as the training image. A total of 867 image tiles at the sizes between 226×226 to 507×507 were processed. A total of 34,931 IWP were manually annotated within the study area [4].

The second dataset is the AI-ready Retrogressive Thaw Slumps (RTS) data, also curated for instance-level segmentation. Thaw slumps are active permafrost features that develop rapidly when ice-rich permafrost thaws. They are a type of landslide formed when ground ice begins to melt causing the ground to become unstable and the soil to move, especially on steep slopes. The training data for RTS were selected from seven sites, combining data used in [6,24]. These sites cover a diverse environmental and geomorphological profile. The three Russian sites include the Lena River, Kolguev, and the Yamal and Gydan Peninsulas, with the latter being the sole known area for gas emission craters. The four Canadian sites encompass Herschel Island, Horton Delta, the Tuktoyaktuk Peninsula, and Banks Island, covering areas of shrubby tundra or dwarf shrub tundra with small lakes and streams, as well as coastal areas near ice-rich permafrost. This dataset contains 855 image tiles and a total of 965 RTS using the 4m Maxar imagery as the base map.

These two features are used to evaluate SAM's image segmentation capabilities for two reasons. First, natural feature datasets are generally more challenging to detect than human-made features, which are common targets in general AI computer vision tasks and many remote sensing image analysis tasks [8]. Their forms are driven by underlying complex geospatial processes, leading to large variations across geographical locations, scales, and landscapes. Because research data (e.g., high-resolution satellite imagery) is often managed in databases and not as readily available as images (e.g., street views) containing human-made features, their inclusion in existing large pre-trained models can be very limited. For this reason, they also become an ideal dataset to test the domain adaptation capabilities for general-purpose foundation models for GeoAI vision tasks. Second, both IWP and RTS are important permafrost features, the changes of which provide a strong linkage to Arctic warming and climate change [25]. Therefore, AI-augmented permafrost mapping is becoming increasingly important to provide scientific insights into the pace of permafrost thaw and to support global change research, monitoring, and policy [26,27].

3.2. Experimental Setup

We designed a series of experiments to evaluate SAM's potential usage for natural feature segmentation, particularly for important Arctic permafrost features. In an instance segmentation task, a model needs to provide not only the mask indicating the exact boundary of an object but also a prediction of the object's class. Object localization, object segmentation, and object class prediction are the three factors that affect an instance segmentation model's performance. Because SAM's goal is to segment "anything", its output contains only masks of all objects without their classes. Therefore, SAM, on its own, cannot perform instance segmentation. It must always be combined with other models to create an instance segmentation workflow.

The first approach is to use SAM to segment objects of any class within an image scene and then connect it with a mask classifier to filter interested objects belonging to a specific class, such as IWP. The SAM and CLIP integrated pipeline shown in Figure 1 belongs to

this category. No training is required in the entire pipeline, except for providing a text prompt, for example, ice-wedge polygon, to CLIP. Hence, this process is also known as zero-shot learning (Strategy 1 in Table 2).

Table 2. Strategies for enabling instance segmentation capability of SAM.

No.	Strategies	Prior Knowledge Embedded?	Zero-Shot?	Fine-Tunable?
1	Feed SAM with regular points (32×32)	-	✓	N/A
2	Feed SAM with ground truth bounding box (BBOX)	✓ (strongest)	-	✓
3	Feed SAM with point locations of the features of interest	✓	-	✓
4	Feed SAM with object detector predicted BBOX	✓	-	✓

The second approach is to provide SAM with prior geospatial knowledge; for example, the location (a point or a bounding box) of objects of interest as a prompt, and then ask SAM to generate the segmented masks of objects of that class. This is a surrogate of instance segmentation; Strategies 2 to 4 in Table 2 are such examples. Strategy 2 is to feed the ground truth BBOX to SAM and ask SAM to segment the object within the given region. This strategy provides the strongest geospatial knowledge among all strategies. Strategy 3 involves feeding SAM with the ground truth point locations for the objects of interest and asking the model to segment the object near the point. Strategy 4 involves training an object detector through supervised learning on the training datasets and using its predicted BBOX to feed the SAM model for instance segmentation. As Strategies 2–4 all require the use of training data in the segmentation pipeline, they are not considered zero-shot learning. Since SAM provides code for the model, the pipelines implementing Strategies 2–4 can be fine-tuned using domain-specific datasets.

3.2.1. Zero-Shot Instance Segmentation (Strategy 1)

In this experiment, we investigate the zero-shot capability of SAM to locate and segment natural landscape features. Since SAM generates only masks without any category information, it is incapable of instance segmentation. To address this, we combine SAM with CLIP to perform instance segmentation tasks and use the IWP and RTS datasets to evaluate their performance. CLIP is also a zero-shot model, which evaluates the correlation between a given text prompt and an input image. Taking advantage of this, we can use CLIP to predict the missing category information of each mask produced by SAM. The CLIP model used in this experiment is ViT-B/32, released by OpenAI in 2021 [15]. More specifics of the model can be found in Table 1.

We used a regular 32 by 32 grid of evenly distributed points across the image as a prompt for SAM. The output of SAM includes masks generated for all segmented objects. These masks are then utilized to clip the original image, creating a new image that solely contains the specific object, with black markings outside the mask. We then sent this image to CLIP along with the text prompt to identify masks belonging to the given object class. This SAM+CLIP integrated modeling approach (Figure 1) implements zero-shot instance segmentation. We evaluated the model on both IWP and RTS datasets using different text prompts for CLIP. As shown in Table 3, the best zero-shot prediction accuracy (measured by mAP50) for the IWP datasets was 0.117 when using “ice-wedge polygon” as the prompt for CLIP. Using other text prompts, such as “ice wedge”, resulted in lower mAP values.

As indicated by the mAP value, the results are poor. The model’s predictions for the RTS dataset were even worse, with a maximum prediction accuracy score of only 0.028. When we examined the results visually (Figure 2), we discovered that RTS is a more challenging feature to segment than IWP due to its complex and less-defined shape. Since the large SAM/CLIP models do not have enough knowledge of RTS, they tend to have a low success rate. Figure 2 shows two examples (RTS1 and RTS2) of partially segmented RTS features using the SAM+CLIP model, and for many of the testing images, no predictions were made.

Table 3. Zero-shot performance of the integrated SAM and CLIP model for instance segmentation. mAP: Mean Average Precision. A value of 50 is a threshold to determine the matches between predicted and ground truth masks. When the IOU (Intersection Over Union) of the two is at or above 50%, the predicted mask is considered a true positive.

Model	Dataset	Prompt for CLIP	mAP50 (SAM)	mAP50 (SAM+CLIP)
SAM+CLIP	IWP	ice-wedge polygon ice wedge	0.073	0.117 0.104
	RTS	thaw slump retrogressive thaw slump	0.003	0.028 0.007




Figure 2. Results of zero-shot learning with the integrated SAM+CLIP model. Mask colors in red and blue represent ground-truth labels and model prediction respectively. The last column displays the final result, and the second-to-last column presents the intermediate results from SAM, which are used as input for CLIP. Image source: Maxar.com, accessed on 1 May 2023.

Second, since SAM segments all objects within an image scene, it can generate false positives, such as the lower right portion of IWP1 and the two very large regions segmented in IWP2. These false positives can be filtered out by CLIP (last column in Figure 2). However, CLIP's recall rate cannot be improved beyond SAM's output. Although CLIP can

filter out irrelevant masks, it may also filter out true positives. To evaluate the impact of CLIP on SAM's performance, we calculated the mAP50 for SAM alone and the integrated SAM+CLIP model. As shown in Table 3, adding CLIP improved the model's overall performance compared to using SAM alone. This suggests that CLIP did not negatively affect SAM's performance, and the low overall score is due to SAM's limited segmentation capabilities for natural landscape features.

3.2.2. Knowledge-Embedded Instance Segmentation with SAM (Strategies 2–4)

Our second set of experiments evaluates SAM's instance segmentation capability with knowledge-embedded learning. Here, the knowledge specifically refers to spatial knowledge, indicating the approximate locations (represented by a point or a BBOX) where a permafrost feature may exist in the image. Ground truth BBOX provides highly accurate location information about the target objects. Thus, the embedded knowledge in these experiments (Strategies 2a and 2b in Table 4) is also the strongest. Masks are not used in these strategies as they provide the answers for SAM. It is important to note that ground truth information (such as ground truth BBOX) should be used only during the training/fine-tuning phase, and the model should not receive any kind of ground truth information during the testing phase, whether for a supervised learning model or a foundation model. Since Strategies 2 and 3 (Table 4) require such ground-truth information to be provided (in the form of a prompt) to SAM during the testing phase, experiments using these strategies indicate only SAM's theoretical upper-bound segmentation performance and cannot be used in real-world scenarios.

In addition, we also assess SAM's performance using less accurate location prompts. This includes utilizing an object detector to provide BBOX information to SAM for further segmentation, which is a more practical approach (Strategy 4). It involves training the object detector on domain-specific datasets and using it during the testing phase to predict the BBOX of the objects of interest before feeding this information to SAM for object segmentation.

Table 4. Strategies used to enable SAM's instance segmentation capability through embedded knowledge, and the model's predictive results. Each strategy number corresponds to the definitions provided in Table 2.

Model	Learning Type	Strategy No.	Prompt	mAP50 (IWP)	mAP50 (RTS)
SAM without fine-tuning	Ground truth knowledge plus zero-shot	2a	Ground truth BBOX	0.844	0.804
		3a	Ground truth point	0.233	0.085
	supervised learning	4a	Object detector predicted BBOX	0.521	0.290
SAM with fine-tuning	Ground truth knowledge plus fine-tuning	2b	Ground truth BBOX	0.989	0.926
		3b	Ground truth point	0.609	0.298
	supervised learning	4b	Object detector predicted BBOX	0.595	0.303
Benchmark segmentation model (MViTv2) [7]	Supervised learning	0	N/A	0.605	0.354

Because SAM provides the source code for training models, the weights can be further fine-tuned using domain datasets to achieve better results. Our main focus in this experiment was to evaluate the impact of fine-tuning on SAM's performance in segmenting natural features (Strategies 2b, 3b, and 4b in Table 4). SAM comprises three primary components: an image encoder, a prompt encoder, and a mask-decoding transformer, all of which are transformer-based architectures. In our experiment, we froze the model parameters in both the image and prompt encoders and focused our attention on fine-tuning the mask-decoding transformer to reduce computation cost and increase result sensitivity. To optimize the model's performance, we used Dice Loss, which calculates the overlap

between the predicted segmentation masks and the ground truth masks. By comparing the results before and after fine-tuning, we can obtain a more comprehensive understanding of SAM's domain adaptation ability for out-of-distribution datasets.

Table 4 shows SAM's strong segmentation performance when provided with ground truth BBOX, achieving mAP50 scores of 0.844 for IWP and 0.804 for RTS (Strategy 2a). However, when given a ground truth point, the accuracy drops significantly to 0.233 for IWP and 0.085 for RTS (Strategy 3a). Interestingly, fine-tuning SAM on domain datasets substantially improves predictive accuracy, with detection accuracy reaching 0.609 for IWP and 0.298 for RTS (Strategy 3b). Despite its poor zero-shot performance on new datasets, SAM demonstrates strong domain adaptation capabilities through fine-tuning (Table 3). This highlights SAM's generalizability in learning and extracting common feature representations from large image datasets. Proper fine-tuning can significantly enhance its performance on new datasets. However, it is important to note that feeding the model with ground truth information during testing (Strategies 2 and 3) is not practical in real-world scenarios.

Strategy 4, which involves training an object detector and feeding SAM with the predicted BBOX at the testing phase, is a practical solution. Here, we selected the Mask region-based convolutional neural network (Mask R-CNN) architecture as our primary object detector because of its dual capability in generating both object BBOX and masks. This allows us to use its predicted BBOX as the input for SAM and also use Mask R-CNN's instance segmentation results as a baseline to evaluate the segmentation quality of SAM. Specifically, we adopted MViTv2 [7], a Mask R-CNN type of model that achieves cutting-edge instance segmentation performance in our study. This model uses a multi-scale vision transformer (MViT) as the feature extraction backbone to replace the traditional CNN-based backbone (e.g., ResNet 50) in a Mask R-CNN model. The MViT achieves cutting-edge instance segmentation performance by taking advantage of both the transformer models in capturing long-range data dependencies and the classic CNN models in hierarchical feature learning and strong information flow enabled by residual connections. In our experiments, the MViTv2, which has 103M parameters, was adopted as it was tested to yield the best performance on our training datasets. The MViTv2 result also offers a practical upper bound in segmentation accuracy for IWP and RTS.

The results in Table 4 indicate that before fine-tuning, Strategy 4a (feeding SAM with object-detector-predicted BBOX) performs better than feeding SAM with ground truth points (Strategy 3a) but worse than when feeding SAM with ground truth BBOX (Strategy 2a). However, SAM's performance is lower on both datasets when comparing the mAPs of SAM using Strategy 4 to the benchmark segmentation model (MViTv2). After fine-tuning, SAM's predictive accuracy improves from 0.521 to 0.595 (mAP50) for IWP segmentation, which is close to the benchmark model's mAP of 0.605 (Strategy 0). However, for the RTS dataset, SAM's performance does not significantly improve after fine-tuning (0.303 vs. 0.290 for Strategy 4), and a gap remains compared to the benchmark segmentation model's result of 0.354. This gap can be attributed to several factors. First, the RTS dataset is more challenging than the IWP dataset, making the learning of representative features crucial for effective segmentation. The MViT model incorporates innovative strategies to integrate the advantages of CNN, originally designed for image analysis, into transformer models primarily designed for processing sequential data. Consequently, it can capture stronger image feature representations than regular transformer or CNN-based models. Second, our fine-tuning focuses on the decoder part, as fine-tuning the encoder part is very computationally expensive and not effectively feasible.

These experiments show that with and without providing ground-truth BBOX (Strategy 2 vs. Strategy 4), there can be over a 20% performance gap in the instance segmentation results. Under model fine-tuning and when SAM is fed with ground truth BBOX (Strategy 2b), the model achieves the best segmentation results closest to the ground truth masks. However, in practice, because SAM does not have the ability to perform instance segmentation, it needs to rely on another object detector's predicted results to achieve the goal. Consequently, its performance is limited by the upstream object detector. Therefore, the results for Strategy

4 for both datasets using SAM remain lower than those of the benchmark model MViTv2 (Strategy 0). Additionally, MViTv2's performance can come very close to or be better than the pipeline implementing Strategy 3b when SAM is provided a ground-truth point and fine-tuned using the domain datasets. It is important to mention that Strategies 2 and 3 are not applicable in real-world application scenarios. Hence, the results for Strategy 4 and Strategy 0 have more practical significance. Figure 3 displays segmentation results for the same images shown in Figure 2. These are the results from the fine-tuned SAM and the benchmark segmentation model (MViTv2).

There are several interesting observations from Figure 3. First, when SAM is fed with ground truth points (Strategy 3b), the resulting masks tend to be smaller than when SAM is fed a ground truth BBOX (Strategy 2b). This is likely because when the input prompt is smaller, SAM's predictions tend to favor labeling smaller masks with higher confidence. Second, some of SAM's segmentation results exhibit holes (IWP1, Strategy 2b, and Strategy 4b), whereas the MViTv2 results (Strategy 0) do not display any holes. This is likely because SAM is trained on datasets containing holes, whereas MViTv2 is trained exclusively on domain datasets that do not include holes. Third, the results of Strategy 4b (feeding SAM with MViTv2-predicted BBOX) and Strategy 0 (MViTv2) are similar since they both segment masks based on the same BBOX information. Consequently, when MViTv2 fails to make certain predictions (e.g., RTS1), SAM (Strategy 4b) will also be unable to produce the corresponding prediction at the respective location.

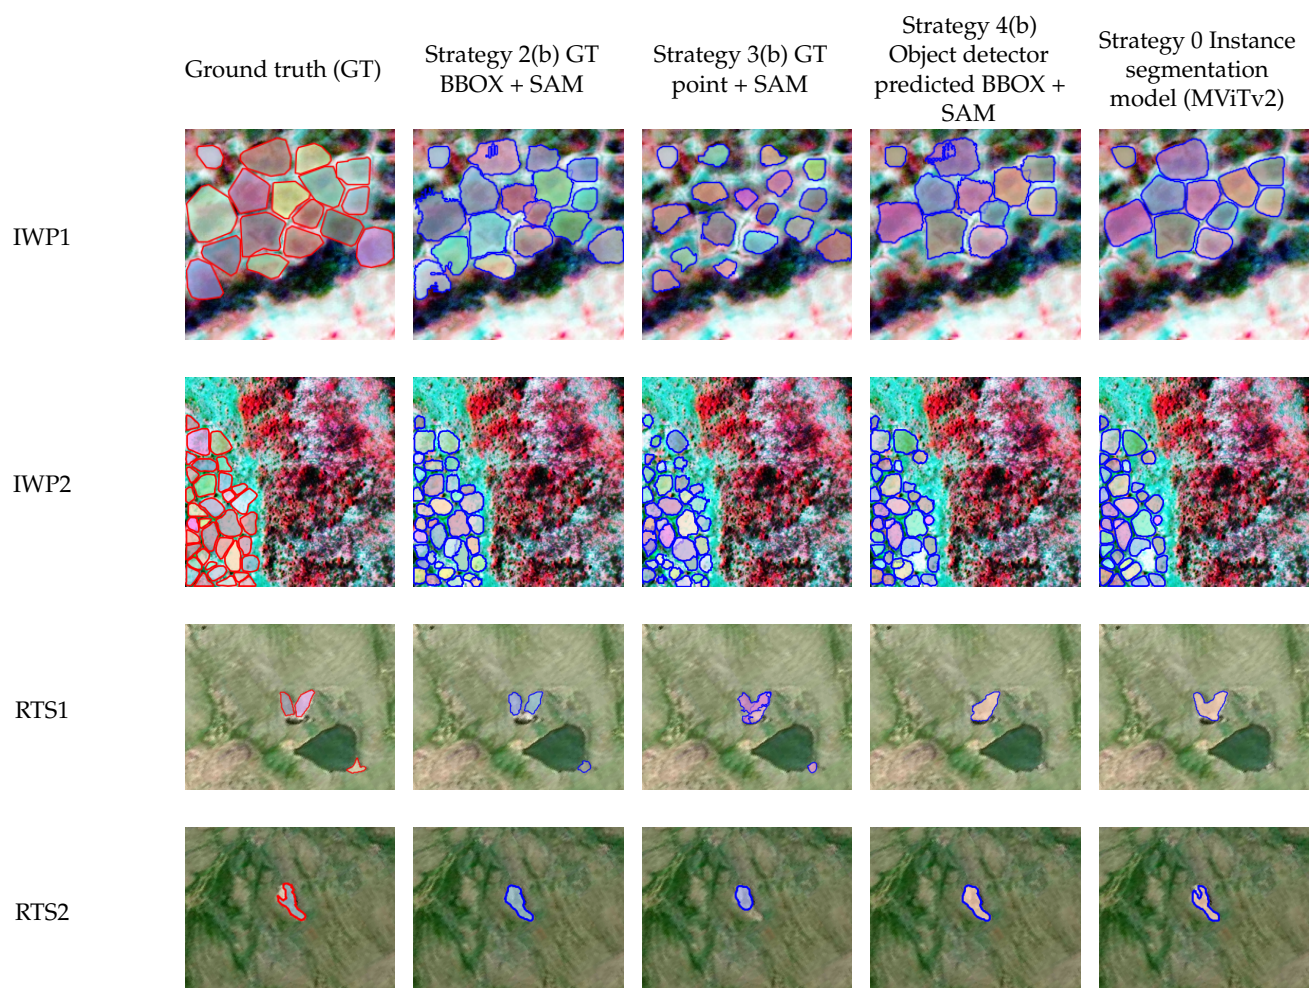


Figure 3. Results from knowledge-embedded learning with SAM. The results are those after fine-tuning. The images are the same as those in Figure 2. Image source: Maxar.com, accessed on 1 May 2023.

4. Discussion

4.1. Summary Analysis

Based on the above analysis, we can see that when directly adopting SAM with no prior knowledge provided (as referenced in the experiment in Section 3.2.1), its segmentation capability for out-of-distribution natural features is not good because these features are likely not considered in the model's pre-training processes. Consequently, SAM's zero-shot learning performance for AI-augmented terrain mapping remains poor. Experiments in Section 3.2.2 show that when providing SAM with strong knowledge—the ground truth BBOX—it can segment the objects inside the target area very well. While ground-truth BBOX cannot be used in a fully automated model prediction scenario, it can be applied in an interactive AI-augmented mapping with human annotators, accelerating the process of data labeling and training data preparation. Some tools have already been developed to integrate SAM into the crowdsourced mapping procedure, such as DS-Annotate [28], representing a major step forward in AI and foundation-model-assisted mapping.

When SAM is used as a downstream model to an object detector that trains on the domain dataset and feeds SAM with its predicted BBOX, the model shows much better performance than using it for zero-shot instance segmentation. This is a strategy that can be used in practical applications. However, its performance may be limited by the upstream object detector, which feeds it with the predicted BBOX. Comparing SAM's performance on these natural feature datasets and general-purpose computer vision datasets such as COCO and LVIS (discussed at the beginning of Section 3), we found that the margins between SAM and a supervised learning model are considerably larger when SAM is adapted for segmenting natural features. Specifically, the difference is nearly 14% (0.521 for SAM vs. 0.605 for MViTv2 in Table 4) for the IWP dataset and 22% (0.290 for SAM vs. 0.354 for MViTv2 in Table 4) for the RTS dataset. This suggests a focused area for improving SAM's domain adaptability in processing geospatial data, especially the challenging permafrost features.

One positive aspect of SAM is that its performance can be improved after fine-tuning. This is likely due to SAM's strategy of pre-training on a huge number of images, giving it the generalizability to extract common characteristics of diverse objects, even though it might not have learned the intrinsic representation of natural features. Compared to other emerging geospatial vision foundation models, such as IBM's Prithvi, which requires 6-band input [29], SAM's input is the most commonly seen RGB bands; hence, it has potentially higher adaptability to diverse datasets.

4.2. Spatial and Domain Generalizability Test

To further validate our findings, additional experiments were conducted using SAM for agriculture field mapping. The EuroCrops dataset is employed in this analysis. EuroCrops [30] is the largest harmonized open crop dataset across the European Union; it consists of 944 image scenes captured at the beginning of the growing period in April 2019. The image mosaic is created from two cloud-free Top of Atmosphere (TOA) Sentinel-2 images, each with a spatial resolution of 10 meters. The study area encompasses central Denmark, characterized by dominant agricultural land use and relatively flat terrains. Each image scene measures 128 by 128 in size. 80% of randomly selected samples are used for training, with the remaining 20% reserved for testing.

The same set of experiments was conducted on SAM using this dataset. Table 5 illustrates the zero-shot performance by SAM and the combined SAM+CLIP model. It is evident that using SAM alone results in a relatively low model performance. As depicted in Figure 4, SAM may mistakenly segment two crop fields as one large field or fail to identify certain crop fields. Upon incorporating CLIP with SAM and utilizing different prompts for semantic filtering, increased prediction accuracy is observed, especially when "Crop land" is used as a prompt. The last column in Figure 4 showcases results using this approach. Notably, CLIP helps to remove some irrelevant urban regions (as seen in image AGR2), contributing to the enhancement of the final results. In contrast, when the prompt "Agricultural field" is utilized, the model yields slightly lower accuracy than when using

SAM alone (Table 5). This suggests that the choice of prompt text is a factor influencing the CLIP model's results.

Table 5. Zero-shot performance of the integrated SAM and CLIP model for instance segmentation of agricultural fields. The same experimental settings were used as in Table 3.

Model	Dataset	Prompt for CLIP	mAP50 (SAM)	mAP50 (SAM+CLIP)
SAM+CLIP	EuroCrops	Crop land Agricultural field	0.118	0.161 0.109




Figure 4. Agricultural field mapping results using zero-shot learning with the integrated SAM+CLIP model. The second column shows ground truth labels (in red). The model prediction results are in blue. The last column displays the final result, and the second-to-last column presents the intermediate results from SAM, which are used as input for CLIP. Image source: Sentinel.

Experiments were also conducted to assess SAM's performance on additional learning strategies, and the statistical results are shown in Table 6. When SAM is provided with the ground truth BBOX of the agricultural lands, its segmentation accuracy can reach up to 0.907 with zero-shot learning and 0.922 after model fine-tuning. This result is much better than that obtained using a ground-truth point as the input prompt. When used in real-world scenarios and provided with another model's predicted BBOX as input, SAM's prediction accuracy (0.694) can get very close to using a cutting-edge supervised learning model MViTv2 (0.717). Figure 5 illustrates some example results for visual inspection. Overall, the findings on this more general dataset are consistent with what was found in the two permafrost datasets. This further verifies our findings about the strengths and limitations of using SAM in geospatial applications.

Table 6. Strategies used to enable SAM’s instance segmentation capability for agricultural field mapping through embedded knowledge, and model predictive results. Each strategy number corresponds to the definitions provided in Table 2.

Model	Learning Type	Strategy No.	Prompt	mAP50 (IWP)
SAM without fine-tuning	Ground truth knowledge plus zero-shot	2a	Ground truth BBOX	0.844
		3a	Ground truth point	0.233
	supervised learning	4a	Object detector predicted BBOX	0.521
SAM with fine-tuning	Ground truth knowledge plus fine-tuning	2b	Ground truth BBOX	0.989
		3b	Ground truth point	0.609
	supervised learning	4b	Object detector predicted BBOX	0.595
Benchmark segmentation model (MViTv2) [7]	Supervised learning	0	N/A	0.605

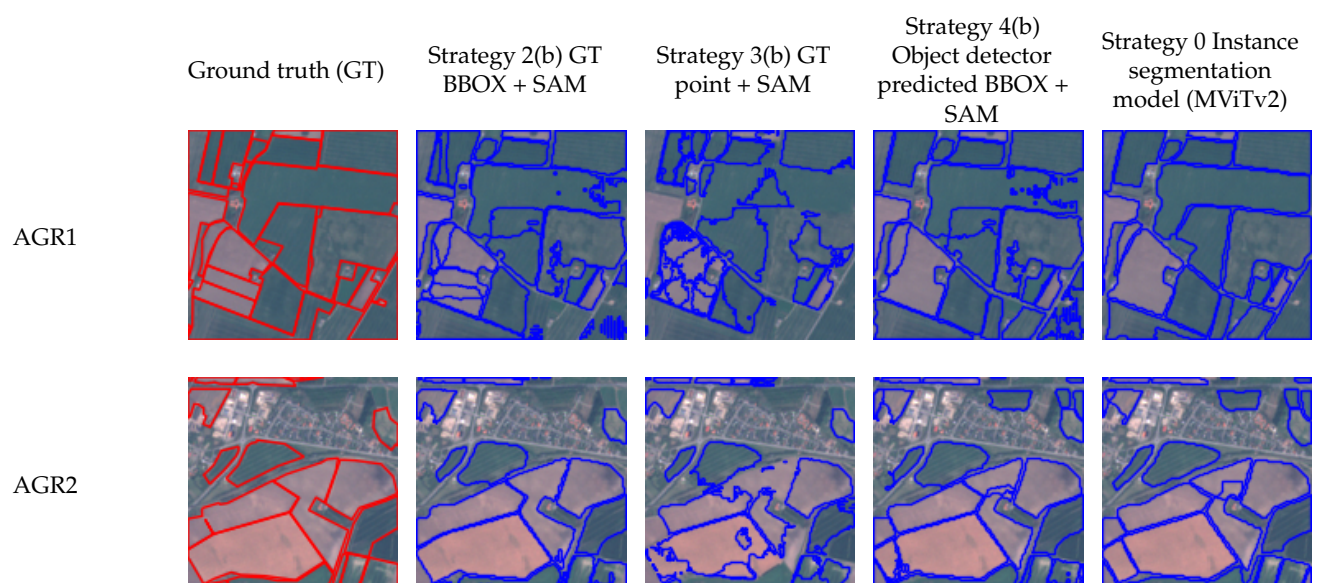


Figure 5. Agricultural field mapping results from knowledge-embedded learning with SAM. The results are those after fine-tuning. The images are the same as those in Figure 4. Image source: Sentinel, accessed on 1 December 2023.

5. Conclusions

This paper describes a methodological framework for enabling and assessing vision foundation models for GeoAI vision tasks. Through developing a series of enabling prompt strategies and operational image-analysis pipelines, we have demonstrated multiple pathways to utilize and adopt a new vision foundation model—SAM for GeoAI instance segmentation. The systematic evaluations help us gain an in-depth view of SAM’s behavior, as well as its strengths and weaknesses in segmenting challenging permafrost landscapes. To the best of our knowledge, this is one of the very first works examining the effectiveness of an AI foundation model in permafrost mapping. The set of strategies, including zero-shot learning, model integration, knowledge-embedded learning, and model fine-tuning, can be easily reused and adopted to evaluate SAM’s performance across different datasets and support diverse, real-world AI-augmented mapping applications (e.g., agricultural land mapping).

To emphasize, the instance segmentation pipeline developed in this paper has maximized the use of SAM’s entire data processing pipeline instead of using only its submodules, thereby fully utilizing its function as a foundation model and reducing model development costs. Our results show varying performance when providing SAM with different prompts. We have also found a bigger performance gap between SAM and cutting-edge supervised learning-based models in segmenting challenging environmental features (i.e., permafrost),

compared to other benchmark computer vision datasets such as COCO and LVIS. This considerable difference emphasizes that SAM's underperformance is more prominent when dealing with natural features (e.g., permafrost) data. However, we have also observed strengths in SAM's model for domain adaptation through the learning of low-level image features shared among different kinds of objects. This strength is evidenced by substantial performance improvement after fine-tuning with permafrost datasets. In the future, more metrics can be further applied to assess SAM's performance in segmenting environmental features, from geometric, spectral, and visual perspectives.

To close the gap between SAM and other cutting-edge supervised learning-based models for instance segmentation of permafrost features (and natural features in general), we suggest several areas of improvement. Data-wise, it is not surprising that SAM's pre-training did not include certain amounts of natural features, as they are often considered as "background" in an image. Hence, expanding SAM's model with more benchmark natural feature datasets will enhance its representation learning ability towards such unique features. Fortunately, SAM's open-source nature enables this expanded capability. Second, these natural features, such as RTS, demonstrate distinct characteristics in different data modalities, such as digital elevation model (DEM) data and Normalized Difference Vegetation Index (NDVI), in addition to optical imagery [31]. Expanding SAM's base model and allowing it to learn the spectral and statistical properties of natural features will further enhance the model's predictive performance. Third, this exploratory analysis of SAM in permafrost mapping provides valuable guidance for the geospatial community to develop the next generation of GeoAI models for large-scale permafrost mapping. By integrating SAM's powerful segmentation head with the multi-scale feature extraction capability of cutting-edge AI models, we can build a more robust model with higher accuracy, addressing crucial data and knowledge gaps in Arctic permafrost thaw, its connection to carbon emissions, and global climate change. We also hope that this paper will spark more discussions about adapting SAM and other visual foundation models to support important domain applications within and beyond spatial sciences.

Author Contributions: Conceptualization, W.L.; methodology, C.-Y.H. and S.W.; validation, W.L., C.-Y.H., S.W., Y.Y. (Yezhou Yang), H.L., A.L., C.W., Y.Y. (Yili Yang), B.M.R., S.T.A., M.B.J., K.M. and P.S.; formal analysis, W.L., C.-Y.H. and S.W.; investigation, W.L., C.-Y.H. and S.W.; resources, W.L.; data curation, Y.Y. (Yili Yang) and C.W.; writing—original draft preparation, W.L., C.-Y.H. and S.W.; writing—review and editing, W.L., C.-Y.H., S.W., Y.Y. (Yezhou Yang), H.L., A.L., C.W., Y.Y. (Yili Yang), B.M.R., S.T.A., M.B.J., K.M. and P.S.; visualization, C.-Y.H., W.L.; supervision, W.L.; project administration, W.L.; funding acquisition, W.L., A.L., K.M., M.J. B.M.R. and P.S. All authors have read and agreed to the published version of the manuscript.

Funding: This work is supported in part by the National Science Foundation under awards 1853864, 2120943, 2230034, 2230035, and 2033521, as well as Google.org's Impact Challenge for Climate Innovation Program.

Data Availability Statement: Code used for the experiments can be accessed through the GitHub repo: <https://github.com/ASUcicilab/SAMPermafrostAssessment>, accessed on 20 February 2024

Acknowledgments: Image data access support for this work was provided in part by the Polar Geospatial Center under NSF-OPP awards 1043681 and 1559691. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, 2–7 June 2019; Volume 1, pp. 4171–4186. <https://doi.org/10.18653/v1/n19-1423>.
2. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is All You Need. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, <https://doi.org/10.48550/arXiv.1706.03762>.

3. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 1877–1901.
4. Udawalpola, M.R.; Hasan, A.; Liljedahl, A.; Soliman, A.; Terstriep, J.; Witharana, C. An Optimal GeoAI Workflow for Pan-Arctic Permafrost Feature Detection from High-Resolution Satellite Imagery. *Photogramm. Eng. Remote Sens.* **2022**, *88*, 181–188.
5. Li, W.; Hsu, C.Y.; Wang, S.; Witharana, C.; Liljedahl, A. Real-time GeoAI for high-resolution mapping and segmentation of arctic permafrost features: The case of ice-wedge polygons. In Proceedings of the 5th ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery, Seattle, WA, USA, 1 November 2022; pp. 62–65.
6. Yang, Y.; Rogers, B.M.; Fiske, G.; Watts, J.; Potter, S.; Windholz, T.; Mullen, A.; Nitze, I.; Natali, S.M. Mapping retrogressive thaw slumps using deep neural networks. *Remote Sens. Environ.* **2023**, *288*, 113495.
7. Li, Y.; Wu, C.Y.; Fan, H.; Mangalam, K.; Xiong, B.; Malik, J.; Feichtenhofer, C. Mvitv2: Improved multiscale vision transformers for classification and detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 4804–4814.
8. Li, W.; Hsu, C.Y. GeoAI for Large-Scale Image Analysis and Machine Vision: Recent Progress of Artificial Intelligence in Geography. *ISPRS Int. J. Geo-Inf.* **2022**, *11*, 385.
9. Li, W. GeoAI: Where Machine Learning and Big Data Converge in GIScience. *J. Spat. Inf. Sci.* **20**, 2020, 71–77.
10. Goodchild, M.F.; Li, W. Replication across Space and Time Must Be Weak in the Social and Environmental Sciences. *Proc. Natl. Acad. Sci. USA* **2021**, *118*, e2015759118.
11. Cha, K.; Seo, J.; Lee, T. A Billion-scale Foundation Model for Remote Sensing Images. *arXiv* **2023**, arXiv:2304.05215.
12. Long, Y.; Xia, G.S.; Li, S.; Yang, W.; Yang, M.Y.; Zhu, X.X.; Zhang, L.; Li, D. On Creating Benchmark Dataset for Aerial Image Interpretation: Reviews, Guidances, and Million-Aid. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2021**, *14*, 4205–4230.
13. Lin, T.Y.; Goyal, P.; Girshick, R.; He, K.; Dollár, P. Focal Loss for Dense Object Detection. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 2980–2988.
14. Wang, D.; Zhang, Q.; Xu, Y.; Zhang, J.; Du, B.; Tao, D.; Zhang, L. Advancing Plain Vision Transformer towards Remote Sensing Foundation Model. *IEEE Trans. Geosci. Remote Sens.* **2022**, *61*, 1–15.
15. Radford, A.; Kim, J.W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J. Learning Transferable Visual Models from Natural Language Supervision. In Proceedings of the International Conference on Machine Learning, PMLR, Virtual, 13 December 2021; pp. 8748–8763.
16. Jiang, P. Conect Segment-Anything with CLIP. 2023. Available online: <https://github.com/PengtaoJiang/Segment-Anything-CLIP> (accessed on 5 May 2023).
17. Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A.C.; Lo, W.Y. Segment Anything. *arXiv* **2023**, arXiv:2304.02643.
18. Martin, D.; Fowlkes, C.; Tal, D.; Malik, J. A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In Proceedings of the 8th IEEE International Conference on Computer Vision, Vancouver, BC, Canada, 7–14 July 2001; Volume 2, pp. 416–423.
19. Li, J.; Yang, T.; Ji, W.; Wang, J.; Cheng, L. Exploring denoised cross-video contrast for weakly-supervised temporal action localization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 19914–19924.
20. Gupta, A.; Dollar, P.; Girshick, R. Lvis: A dataset for large vocabulary instance segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 5356–5364.
21. Lin, T.Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; Zitnick, C.L. Microsoft coco: Common objects in context. In Proceedings of the Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, 6–12 September 2014; Proceedings, Part V 13; Springer: Berlin/Heidelberg, Germany, 2014; pp. 740–755.
22. Liljedahl, A.K.; Boike, J.; Daanen, R.P.; Fedorov, A.N.; Frost, G.V.; Grosse, G.; Hinzman, L.D.; Iijima, Y.; Jorgenson, J.C.; Matveyeva, N. Pan-Arctic Ice-Wedge Degradation in Warming Permafrost and Its Influence on Tundra Hydrology. *Nat. Geosci.* **2016**, *9*, 312–318.
23. Reynolds, M.K.; Walker, D.A.; Balser, A.; Bay, C.; Campbell, M.; Cherosov, M.M.; Daniëls, F.J.; Eidesen, P.B.; Ermokhina, K.A.; Frost, G.V.; et al. A raster version of the Circumpolar Arctic Vegetation Map (CAVM). *Remote Sens. Environ.* **2019**, *232*, 111297.
24. Nitze, I.; Heidler, K.; Barth, S.; Grosse, G. Developing and testing a deep learning approach for mapping retrogressive thaw slumps. *Remote Sens.* **2021**, *13*, 4294.
25. Nitze, I.; Heidler, K.; Barth, S.; Grosse, G.; Bernhard, P. Towards a Deep-Learning based Inventory of Retrogressive Thaw Slumps across the Arctic. In Proceedings of the AGU Fall Meeting Abstracts, Chicago, IL, USA, 12–16 December 2022; Volume 2022, p. C52B–01.
26. Meredith, M.; Sommerkorn, M.; Cassotta, S.; Derksen, C.; Ekaykin, A.; Hollowed, A.; Kofinas, G.; Mackintosh, A.; Melbourne-Thomas, J.; Muelbert, M.; et al. *Polar Regions. Chapter 3, IPCC Special Report on the Ocean and Cryosphere in a Changing Climate*; National Oceanic and Atmospheric Administration: Washington, DC, USA, 2019.
27. Natali, S.M.; Bronen, R.; Cochran, P.; Holdren, J.P.; Rogers, B.M.; Treharne, R. Incorporating permafrost into climate mitigation and adaptation policy. *Environ. Res. Lett.* **2022**, *17*, 091001.
28. Avery. Our ML Tooling 2022—Part 2: Data Operations, 2022. Available online: <https://developmentseed.org/blog/2022-12-19-ml-tooling-2> (accessed on 17 September 2023).

29. Jakubik, J.; Chu, L.; Fraccaro, P.; Gomes, C.; Nyirjesy, G.; Bangalore, R.; Lambhate, D.; Das, K.; Oliveira Borges, D.; Kimura, D.; et al. Prithvi-100M. *arXiv* **2023**, arXiv:2310.18660. <https://doi.org/10.57967/hf/0952>.
30. Schneider, M.; Schelte, T.; Schmitz, F.; Körner, M. EuroCrops: The Largest Harmonized Open Crop Dataset Across the European Union. *Sci. Data* **2023**, *10*, 612.
31. Wang, S.; Li, W. GeoAI in Terrain Analysis: Enabling Multi-Source Deep Learning and Data Fusion for Natural Feature Detection. *Comput. Environ. Urban Syst.* **2021**, *90*, 101715.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.