SegFormer and SegFormer-UNet for anthropogenic geomorphic feature extraction from land surface parameters

Sarah Farhadpour [1*], Aaron E. Maxwell [1], and Behnam Solouki [1]

[1] West Virginia University Department of Geology and Geography, Morgantown, WV, USA

*Correspondence: sf00039@mix.wvu.edu

Paper is a non-peer reviewed preprint submitted to *EarthArXiv*.

The paper has been submitted for peer review to the *International Journal of Remote Sensing*.

Abstract

Accurate, scalable mapping of anthropogenic geomorphic features from high spatial resolution terrain data remains challenging. While convolutional neural networks (CNNs) excel at characterizing local texture and patterns, their limited receptive fields may fail to capture broader spatial context. Transformer-based architectures, such as SegFormer, support stronger long-range dependency modelling, but their value for geomorphic mapping has not been systematically tested. We compare three architectures for pixel-level extraction of features from light detection and ranging (lidar)-derived land surface parameters (LSPs): (1) SegFormer, (2) SegUNet (a hybrid SegFormer-UNet model), and (3) a fully CNN-based UNet with a ResNet-34 encoder. Models are trained on 512×512 LSP chips for three tasks: agricultural terraces, historic mine benches, and valley fill faces. We evaluate random vs. ImageNet initialization and, when pretraining is used, frozen vs. unfrozen encoders. Training uses focal Tversky loss to address class imbalance; performance is assessed with F1-score, precision, recall, and overall accuracy, with robustness summarized via resampling of test chips. Results suggest that end-to-end training is the most consistent driver of gains across backbones and datasets, particularly for SegFormer. UNet-style architectures with CNN-based decoders generalize most reliably: SegUNet and the UNet typically achieve the best precision and recall balance and the smoothest optimization. Contrary to common practice, random initialization often matches or exceeds ImageNet pretraining for LSP inputs, suggesting a domain gap between natural-image features and lidar-based terrain representations. Across tasks, difficulty tracks object morphology: agricultural terraces are hardest, mine benches intermediate, and valley fill faces easiest. Practically, we recommend unfrozen encoders with a CNN-based decoder as a strong default, and we caution against assuming ImageNet pretraining will help when using LSP predictor variables. These findings provide clear guidance for selecting and optimizing deep learning models for anthropogenic geomorphic feature extraction from high spatial resolution digital terrain data.

1. Introduction

Remotely sensed data combined with feature extraction and classification techniques are widely applied to geomorphic mapping tasks such as landslide susceptibility modelling (Azarafza et al. 2021; Thi Ngo et al. 2021), archaeological mapping (Albrecht et al. 2019; Guyot et al. 2021), landform classification (Xu et al. 2022; Li et al. 2020), and geomorphic feature extraction (van der Meij et al. 2022; Maxwell et al. 2023). Deep learning-based methods, particularly convolutional neural networks (CNNs), are particularly applicable to such task due to their ability to characterize texture, local context, and spatial patterns (Zhang, Zhang, and Du 2016; Zhu et al. 2017; Ma et al. 2019). More specifically, CNN-based pixel-level classification methods (i.e., semantic segmentation), introduced through fully convolutional networks in 2014–2015 (Long, Shelhamer, and Darrell 2015), are widely used. However, one limitation of CNNs is their relatively small receptive field, which can hinder their ability to capture longer-range spatial dependencies that are often important in geomorphic mapping tasks. This structural constraint may reduce their effectiveness for extracting elongated or spatially dispersed landforms or capturing broader spatial context.

More recently, transformer-based architectures have gained traction for semantic segmentation tasks. Unlike CNNs, they can capture long-range dependencies and patterns, offering strong generalization across diverse data types. A notable example is SegFormer (Xie et al. 2021), which combines hierarchical feature extraction with lightweight transformer modules. While promising, its comparative performance against CNN-based methods, such as UNet-like

architectures (Ronneberger, Fischer, and Brox 2015), in geomorphic mapping remains underexplored.

This study addresses this gap by comparing three model architectures: (1) a SegFormer model with five encoder variants based on the Mix Transformer (MiT-B0 to MiT-B4) module, (2) a hybrid SegFormer-UNet model (termed SegUNet in this study) that combines the same transformer-based encoder with a UNet-style, CNN-based decoder, and (3) a standard UNet model with a ResNet-34 (Residual Network-34), CNN-based encoder (He et al. 2016). All models are trained using land surface parameters (LSPs) derived from high spatial resolution, light detection and ranging (lidar)-based digital terrain models (DTMs). We apply these models to three anthropogenic geomorphic feature extraction tasks: (1) mapping agricultural terraces in Iowa, USA (terraceDL dataset (Maxwell 2023a)); (2) delineating valley fill faces from surface coal mine reclamation in southern West Virginia, eastern Kentucky, and southwestern Virginia, USA (vfillDL dataset (Maxwell 2023b)); and (3) extracting historic mine benches in northern West Virginia, USA (mineBenchDL dataset (Maxwell, Farhadpour, and Ali 2024b)). We compare models trained with random initialization and ImageNet (Deng et al. 2009) pretraining, further evaluating the effect of freezing vs. unfreezing encoder parameters during training.

This comparative analysis is designed to evaluate how architecture type, parameter initialization, and training strategy influence segmentation performance for geomorphic features with varying geometry, spatial context, and representation in the landscape, and to assess whether, given comparable training sizes, the increased receptive field of transformers can offer improvements for geomorphic feature extraction.

2. Background

2.1. Light detection and ranging and digital terrain data

Lidar is an active remote sensing technology capable of acquiring highly accurate and detailed elevation data at high spatial resolutions (i.e., less than a 2 m cell size for derived DTMs). Lidar functions by emitting laser pulses from a laser range finder, which can reflect back to the sensor after interacting with the ground or above-ground features. The two-way travel time of these laser pulses, combined with global navigation satellite systems (GNSS) for spatial referencing and an inertial measurement unit (IMU) for platform orientation characterization, enables the generation of precise, high spatial resolution, three-dimensional point clouds of the ground surface and landscape features. The ability of lidar to capture multiple returns from a single laser pulse facilitates partial penetration of tree canopies, enabling terrain mapping even in forested areas, particularly when data are collected during leaf-off periods. Non-ground returns can be filtered using algorithms and manual adjustments to isolate ground points, which are then interpolated to create raster-based DTMs, providing a bare-earth surface representation (Lillesand, Kiefer, and Chipman 2015).

Lidar-derived DTMs allow for the extraction of various LSPs, which characterize terrain properties such as slope or steepness, orientation (aspect), terrain roughness, topographic position, and surface curvature. These parameters correlate with critical environmental factors for modelling and mapping, including moisture distribution, solar radiation exposure, and erosion susceptibility (Franklin 2020; Maxwell and Shobe 2022). Beyond topographic representation, DTMs are crucial for monitoring and modelling landscape changes over time. For example, they have been used to analyse sediment connectivity and assess the impacts of anthropogenic structures and landscape modification on natural processes (Graf et al. 2018; Cazorzi et al. 2013). The precision of DTMs is important in predicting dynamic processes, such as soil erosion and slope failures, since the

accuracy of elevation data directly influences model reliability (Chen et al. 2014; Bossi et al. 2015). High spatial grain DTMs derived from lidar data also enable the identification and analysis of fine-scale features, crucial for detailed geomorphic analyses and environmental assessments (Chen, Gao, and Devereux 2017; Sampson et al. 2016). Lidar data have transformed fields such as landform and surficial geologic mapping (Sofia, Fontana, and Tarolli 2014; Tarolli 2014), geohazard assessment and modelling (Jaboyedoff et al. 2012; van Westen, Castellanos, and Kuriakose 2008), and archaeological exploration (Chase et al. 2012; Fernandez-Diaz et al. 2014; Hesse 2010).

In geomorphic mapping and feature extraction tasks, using LSPs as predictor variables is advantageous compared to image bands because LSPs explicitly represent physical attributes of the terrain surface, such as steepness, curvature, and topographic position. Unlike spectral image bands that primarily capture surface reflectance, LSPs can directly correlate with geomorphic processes and terrain characteristics critical for accurate feature identification. Maxwell et al. (Maxwell et al. 2023) documented that using multiple LSP-based predictor variables supports more accurate extraction of specific geomorphic features using deep learning-based methods, outperforming traditional terrain visualization methods such as hillshades, multidimensional hillshades, and slopeshades.

2.2. SegFormer and transformers

The recent emergence of transformer-based architectures has reshaped deep learning well beyond the original context of natural language processing. The transformer introduced by Vaswani et al. (Vaswani et al. 2017) replaces recurrence and convolution with stacked self-attention and feed-forward layers, allowing the model to learn global dependencies across entire sequences. Multi-head self-attention decomposes the attention operation into multiple parallel heads so that different types of relationships can be modelled simultaneously, while positional encodings inject information about order that is not inherent to the architecture itself (Vaswani et al. 2017). Although originally proposed for sequence transduction, these properties make transformers attractive for other modalities, including imagery, where long-range spatial interactions are important.

Dosovitskiy et al. (Dosovitskiy et al. 2020) adapted this framework to computer vision with the Vision Transformer (ViT), which treats an image as a sequence of fixed-size patches and applies a standard transformer encoder to those tokens. ViT demonstrated that convolution is not strictly necessary to achieve strong performance on image classification, provided sufficiently large pre-training datasets are available (Dosovitskiy et al. 2020). However, its relatively weak inductive bias compared to convolutional neural networks (CNNs) and its reliance on large-scale data motivated subsequent work that hybridized transformer concepts with architectures tailored to dense prediction tasks. For semantic segmentation, models such as Segmenter (Strudel et al. 2021) and TrSeg (Jin, Han, and Ko 2021) leverage transformer encoders with additional mechanisms for multi-scale context aggregation, enabling more effective segmentation of objects spanning a range of sizes.

SegFormer (Xie et al. 2021) is a transformer-based segmentation architecture designed explicitly to balance accuracy, efficiency, and robustness. As illustrated in Figure 1, it consists of two main components: a hierarchical encoder built from Mix Transformer (MiT) stages and a lightweight all-MLP decoder. The encoder progressively downsamples the input while maintaining overlapping patch partitions, which helps preserve local continuity and yields multi-scale feature maps that capture both fine and coarse spatial structure (Xie et al. 2021). In contrast to standard ViT implementations, SegFormer does not rely on fixed positional embeddings; this design choice

helps maintain performance when the model is applied at spatial resolutions that differ from those seen during training (Xie et al. 2021). The MiT backbone is available in multiple sizes (MiT-B0 through MiT-B5), spanning compact real-time variants to larger, more accurate configurations.

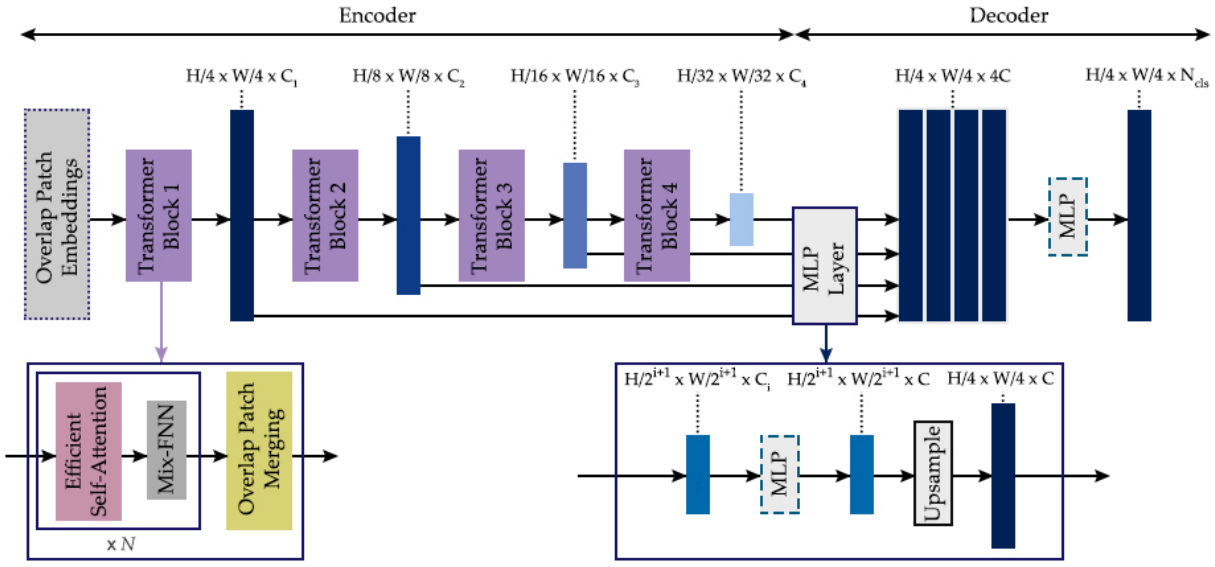

Figure 1. The SegFormer framework (Xie et al. 2021) comprises two primary components: a hierarchical transformer-based encoder designed to capture both coarse and fine features, and a lightweight multi-layer perceptron (MLP)-based decoder that efficiently fuses these multi-scale features to generate a semantic segmentation mask. FFN = feed-forward network.

The decoder operates purely with MLP layers, avoiding computationally expensive modules while still fusing information across scales. It first projects encoder feature maps from each stage to a common channel dimensionality, upsamples these representations to a shared intermediate spatial resolution (typically one quarter of the input size), concatenates them, and applies another MLP to produce a fused feature representation (Xie et al. 2021). A final upsampling step yields the full-resolution segmentation map. This design allows SegFormer to exploit the long-range context captured by the encoder while keeping the decoding process simple and efficient. Xie et al. (Xie et al. 2021) reported that SegFormer outperforms several conventional CNN-based segmentation models in terms of accuracy and computational cost and exhibits improved robustness to various image corruptions.

Beyond these general claims of efficiency and robustness, empirical evidence from remote sensing studies further demonstrates SegFormer's strengths. SegFormer has already been applied across a range of remote sensing tasks, where it has often demonstrated advantages over conventional CNN-based architectures. For example, for urban land cover semantic segmentation of two specific datasets, SegFormer and its variants achieved higher mIoU (mean intersection-over-union) and F1-scores than widely used CNN models such as UNet, ResUNet, and DeepLabV3, indicating improved capacity to capture complex spatial patterns (Chen, Li, et al. 2023). Similarly, in UAV-based crop type segmentation, SegFormer outperformed UNet, PSPNet, and DeepLabV3+ as measured by mIoU and pixel-level overall accuracy, underscoring its effectiveness for agricultural mapping tasks (Xiang et al. 2023). SegFormer has also been successfully employed for building detection in high spatial resolution aerial and satellite imagery, where it surpassed CNN-based models including UNet, PSPNet, HRNet, and DeepLabV3+ in

terms of boundary integrity and overall detection accuracy (Li, Rui, Yang, Liu, et al. 2023). Collectively, these studies suggest that SegFormer frequently performs better than traditional CNN-based architectures in remote sensing applications, particularly for tasks that require capturing longer-range spatial dependencies alongside fine-grained local detail.

These characteristics are particularly relevant for geomorphic mapping. Self-attention enables transformer-based models to relate distant pixels directly, which may help characterize landforms and landscape features whose diagnostic patterns extend over tens to hundreds of meters (Vaswani et al. 2017). At the same time, this family of models has limitations that are especially pertinent for geomorphic applications. Transformers generally require larger training datasets to generalize well because they lack the strong spatial inductive biases provided by convolution (Dosovitskiy et al. 2020), yet geomorphic mapping often relies on relatively small, labour-intensive labelled datasets. In addition, the quadratic memory and computational cost of full self-attention can constrain the use of transformers with high spatial resolution inputs that are common in lidar-derived products. Given this mixture of potential benefits and drawbacks, SegFormer is a natural candidate for systematic evaluation in lidar-based geomorphic feature extraction, where long-range spatial context is important but training data and computational budgets are limited.

2.3. UNet architecture

UNet (Ronneberger, Fischer, and Brox 2015), illustrated in Figure 2a, is a widely used encoder–decoder architecture for semantic segmentation. The encoder applies a sequence of double-convolution blocks, each consisting of two 3×3 convolutions followed by batch normalization and a rectified linear unit (ReLU) activation. Batch normalization stabilizes intermediate feature distributions and supports faster, more stable training (Bjorck et al. 2018; Ioffe 2015; Luo et al. 2018), while ReLU introduces nonlinearity and sparsity (Agarap 2018; Sharma, Sharma, and Athaiya 2020). Spatial resolution is reduced by 2×2 max-pooling with stride 2, so that successive encoder stages operate over increasingly large receptive fields but coarser spatial grids (Ronneberger, Fischer, and Brox 2015).

The decoder mirrors this structure and progressively restores spatial resolution using transpose convolutions or other upsampling operations (e.g., bilinear interpolation). At each stage, feature maps from the encoder are concatenated with the corresponding decoder activations via skip connections, allowing high-resolution spatial detail to be reintroduced into deeper, semantically rich representations. A final 1×1 convolutional layer serves as the pixel-wise classification head, mapping decoder features to class logits. For multiclass segmentation, the number of output channels matches the number of classes; for binary segmentation, either a single logit or a two-channel (foreground/background) output is used, followed by sigmoid or softmax normalization as appropriate (Ronneberger, Fischer, and Brox 2015).

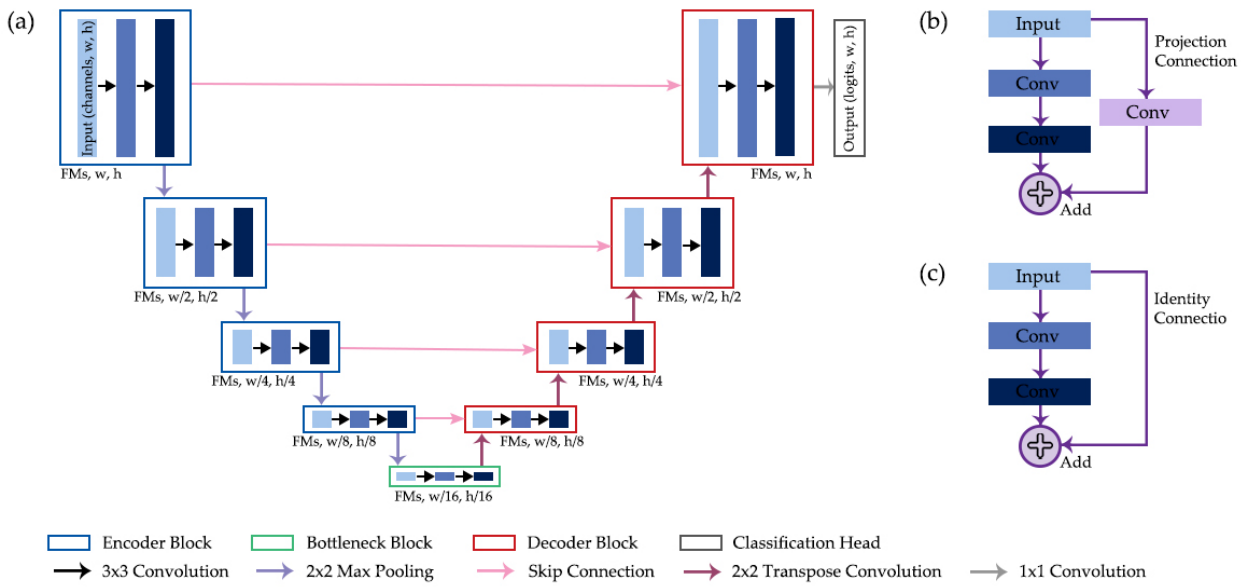

Figure 2. (a) UNet architecture (Ronneberger, Fischer, and Brox 2015), where FM represents feature maps, and *w* and *h* denote the number of columns and rows of cells in the arrays, respectively. Blocks represent data or tensors at various stages of processing. (b) Residual projection connection (He et al. 2016) integrated into a double-convolution block, used when the number of feature maps and/or spatial dimension sizes change. (c) Residual identity connection (He et al. 2016) within a double-convolution block, applied when the number of feature maps and spatial dimension sizes remain unchanged.

Beyond the original biomedical formulation, UNet has become a flexible template that is readily adapted for diverse tasks (Farhadpour and Maxwell 2025). The encoder can be replaced with standard CNN backbones such as ResNet (He et al. 2016) or DenseNet (Huang et al. 2017), while additional modules can be inserted in the bottleneck or along skip connections to enhance representation capacity (Kugelman et al. 2022; Krithika alias AnbuDevi and Suganthi 2022). Recent variants also integrate transformer-based components into either the encoder or the skip pathways (Sha et al. 2021; Cao et al. 2022; Wang et al. 2022). This modularity makes UNet a general framework for experimenting with different feature extractors and context mechanisms within a consistent encoder–decoder design.

2.4. ResNet-34 encoder

Residual connections are a key innovation that make it possible to train very deep neural networks without severe degradation in performance. In the ResNet family (He et al. 2016), these shortcuts bypass one or more convolutional layers and encourage the network to learn residual corrections to an identity mapping, rather than a full transformation from scratch. As illustrated in Figures 2b–c, residual links are typically implemented in two ways: identity connections, where the input is added directly to the output when the number of feature maps and spatial dimensions are unchanged, and projection connections, which use 1×1 (or occasionally 3×3) convolutions to align channel counts and spatial sizes before addition. Both variants promote stable gradient flow, maintain information across depth, and make deeper architectures easier to optimize.

Within a UNet framework, a ResNet-34 encoder provides a natural replacement for the original double-convolution encoder blocks. Residual blocks allow the model to increase depth

and representational capacity while controlling computational cost, and ImageNet-pretrained weights (Deng et al. 2009) offer a strong starting point for transfer learning in domains with limited labelled data. Integrating ResNet into UNet has demonstrated significant performance improvements across various domains. For example, a UNet with a ResNet-based encoder, referred to as Res-UNet, was used for tree species classification in high spatial resolution airborne imagery, achieving superior accuracy compared to base UNet (Cao and Zhang 2020). Similarly, a hybrid architecture combining residual connections and additional enhancements, ResUNet-a, performed well for urban object segmentation using high spatial resolution aerial imagery (Diakogiannis et al. 2020). For geological fault mapping from seismic data, a ResUNet model demonstrated improved performance over base UNet (Liu et al. 2020). In medical imaging, a model incorporating residual connections and dilated convolution for liver and lesion segmentation, RMS-UNet, outperformed traditional UNet models (Khan, Luo, and Wu 2022).

These examples underscore the versatility and effectiveness of integrating residual connections into UNet. Residual connections enhance gradient flow, preserve features across layers, and improve the model's ability to capture both fine details and contextual information.

2.5. Transfer learning

Transfer learning has become a central strategy in deep learning, allowing models trained on large, generic datasets to be adapted to new tasks with far less labelled data and reduced computational demand (Zhu et al. 2017). In this paradigm, representations learned from extensive image collections such as ImageNet are reused as a starting point for downstream problems, rather than training a network entirely from random initialization (Ma et al. 2019; Yuan et al. 2020). This reuse of pre-learned features often accelerates convergence and improves generalization, especially in settings where annotated data are scarce or expensive to obtain (Zhu et al. 2017).

Several studies have successfully implemented transfer learning techniques in remote sensing and/or geomorphic mapping. Pires de Lima & Marfurt (Pires de Lima and Marfurt 2020) investigated remote sensing scene classification problems and documented that ImageNet-based transfer learning improved accuracy for distinguishing diverse land use and land cover types, highlighting its value for tasks requiring high feature extraction accuracy. Cao et al. (Cao, Wu, and Zou 2024), who applied transfer learning for sustainable ecological management, demonstrated that models trained on established datasets can be adapted to new environments. While their study primarily focused on ecological sustainability, the approach highlights the versatility of transfer learning in analysing landscape changes, an essential aspect of geomorphic mapping. Similarly, Chen & Zheng (Chen and Zheng 2023) explored transfer learning for surface information extraction, leveraging diversity and uncertainty in the training process to filter samples effectively. Additionally, Zhao and Feng (Zhao and Feng 2022) applied transfer learning to water body recognition from remotely sensed images, highlighting its efficiency in tasks with scarce labelled data. Building on this broader literature, our prior work also compared transfer learning strategies for anthropogenic geomorphic feature extraction from lidar-derived LSPs using UNet architectures (Maxwell, Farhadpour, and Ali 2024a). We documented that ImageNet-based transfer learning generally improved performance, even with frozen encoders, and provided gains over random initialization. By contrast, geomorphon-based transfer learning and transfer learning between the different geomorphic datasets used in the study offered minimal benefits despite being derived from terrain-focused features. These findings suggest that the abstractions captured by ImageNet pretraining, such as edges and textures, can be broadly transferable and highly effective for geomorphic mapping, strengthening the case for transfer learning as a powerful strategy in this domain.

These studies underscore the role of transfer learning in enabling deep learning models to generalize across heterogeneous landscapes, improving feature recognition even in areas with limited annotated training data. The success of transfer learning is further enhanced by its dual strategies of feature extraction and fine-tuning. Feature extraction involves utilizing the learned representations from a pre-trained model to classify new data, while fine-tuning adjusts the model's parameters to better fit the target task (Cheng 2023; Novakovsky et al. 2021). These strategies have been instrumental not only in remote sensing applications but also in diverse areas, such as plant species identification to medical prognosis prediction, where the availability of labelled data is often limited (Chai et al. 2021; Tariku et al. 2023).

3. Methods

3.1. Study areas and datasets

This study draws upon three existing lidar-based datasets that target different types of anthropogenic landforms in the United States: terraceDL, minebenchDL, and vfillDL (Figure 3). Together, they span agricultural engineering structures, legacy surface mining benches, and valley fills, and therefore sample a range of morphologies, spatial scales, and geographic settings relevant to geomorphic mapping. Terraced agricultural landscapes in Iowa, USA form the basis of the terraceDL dataset (Maxwell 2023a). These features were mapped as line segments during the Iowa Best Management Practices (BMP) project (McNeely et al. 2017), where interpreters used lidar-derived surfaces and aerial imagery to identify terrace structures designed to slow runoff and reduce soil loss. For the present study, terrace features located within three Hydrologic Unit Code (HUC)-8 watersheds were extracted and divided into non-overlapping subsets for training (66,000 features), validation (15,505), and testing (26,453), as shown in Figure 3a. Example terrain visualizations associated with terraceDL are illustrated in Figure 3d and 3g.

The minebenchDL dataset (Maxwell, Farhadpour, and Ali 2024b) focuses on bench-like platforms created during contour surface coal mining in northern West Virginia, USA. These features typically formed when coal seams were cut along slope contours, leaving narrow benches adjacent to steep highwalls. Although many of these benches are now vegetated and difficult to distinguish in optical imagery, they remain apparent in high spatial resolution DTMs. To construct minebenchDL, 169 USGS 1:24,000 quarter-quadrangles containing mine benches were manually interpreted using lidar, aerial imagery, and supplemental mine disturbance layers. The quadrangles were then randomly assigned to training (69), validation (50), and testing (50) groups, yielding a mapped inventory of 2,416 bench features (Figure 3b, 3e and 3h).

The vfillDL dataset (Maxwell 2023b) represents valley fills associated with mountaintop removal mining in southern West Virginia and adjacent areas of eastern Kentucky and southwestern Virginia. In these steep landscapes, mine overburden is commonly placed into nearby valleys, creating large artificial landforms that differ from the surrounding natural topography (Miller and Zégre 2014; Palmer et al. 2010; Ross, McGlynn, and Bernhardt 2016; Wickham et al. 2013). Mapping for vfillDL focused on the exposed faces of valley fills, the portions most clearly expressed in lidar-derived terrain data, rather than attempting to infer their entire volumetric footprint. A total of 1,105 fills from southern West Virginia were used for training, while an additional 1,178 fills distributed across the broader region were split evenly into validation and test subsets (Figure 3c, 3f and 3i).

Together, these three datasets cover different regions, scales, and geometries of anthropogenic modification, from narrow, closely spaced terraces to sparse, elongated mine benches and broad

valley fills. This diversity provides a rigorous test of our models' ability to generalize across landforms that vary substantially in size, planform shape, and surrounding topographic context.

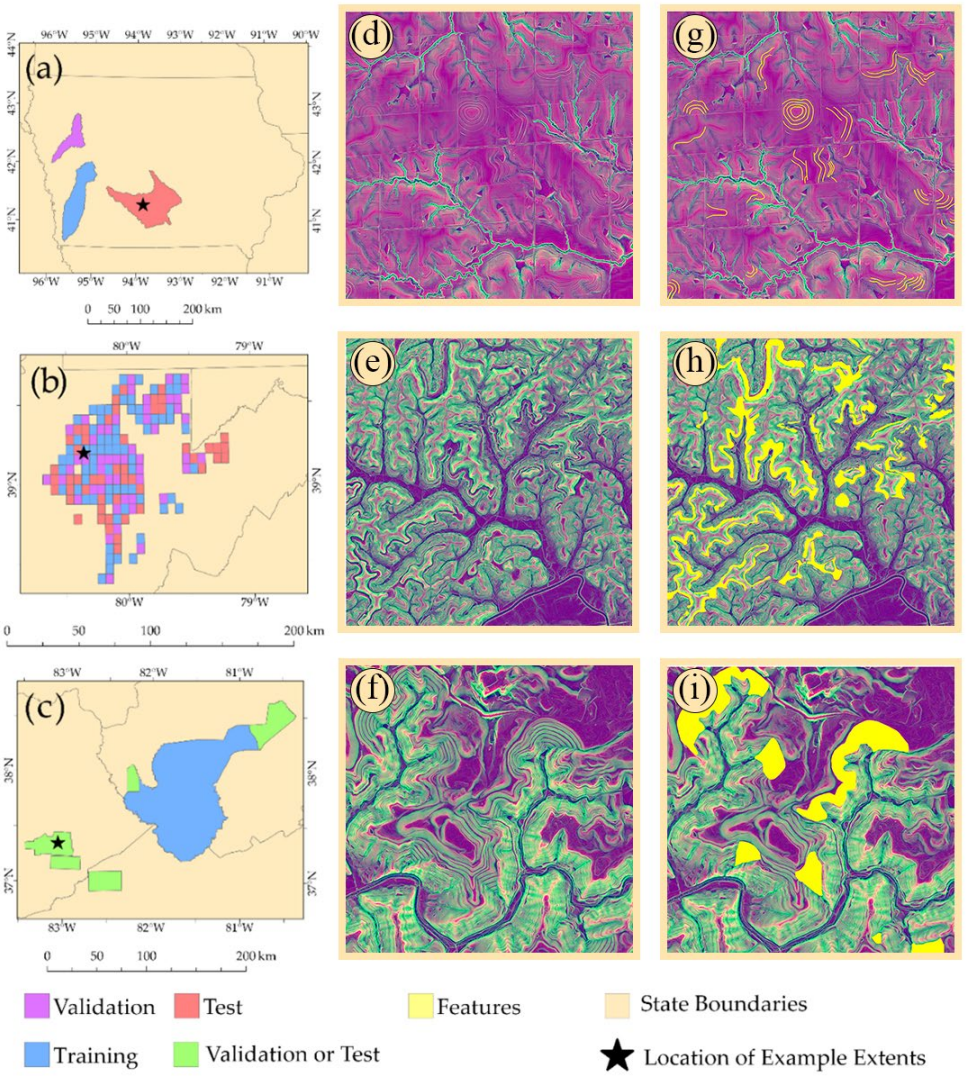

Figure 3. Locations of the three study datasets: (a) the agricultural terraces dataset (terraceDL) (Maxwell 2023a) located in Iowa, USA; (b) the historic mine benches dataset (minebenchDL) (Maxwell, Farhadpour, and Ali 2024b) from northern West Virginia, USA; and (c) the valley fills dataset (vfillDL) (Maxwell 2023b) derived from surface coal mining areas in southern West Virginia, eastern Kentucky, and southwestern Virginia, USA. Panels (d) through (i) illustrate example terrain surfaces and their corresponding geomorphic features.

3.2. Land surface parameters (LSPs)

High spatial resolution terrain information for this study was obtained from lidar-derived digital terrain models produced through the USGS 3D Elevation Program (3DEP) (Sugarbaker et al. 2014). The 3DEP initiative provides nationwide elevation data collected under uniform quality standards (Arundel et al. 2015), ensuring that the terrain inputs used here are consistent across all three study regions. To construct the model input space, we used a compact set of terrain descriptors designed to emphasize landform position, local relief, and slope characteristics at

multiple spatial scales. This three-layer representation has been shown to be effective for identifying anthropogenic landforms in previous geomorphology-focused research (Maxwell et al. 2023), although the parameters were recalculated specifically for the areas considered in this study.

The first layer, a broad-scale topographic position index (TPI), characterizes how elevated or depressed a location is relative to its surroundings. This measure is derived using a 50-m circular neighbourhood, where each cell's elevation is compared to the average elevation within that window. Cells situated higher than their surroundings receive positive values, whereas cells positioned lower than their neighbours receive negative values. This layer highlights ridges, drainage bottoms, and broad hillslope positions. The second layer represents terrain steepness using the square root of slope. Slope was calculated in degrees and then transformed using a square root operation to dampen the influence of extreme gradients while retaining meaningful contrast among moderate slopes. This transformation has been found useful for feature extraction because it compresses long-tailed slope distributions without removing important variation. The third layer is a fine-scale landform position metric, computed as a second TPI but using an annulus window with inner and outer radii of 2 m and 5 m, respectively. This filtering isolates subtle breaks in slope and small ridges or depressions that may be difficult to capture with coarser metrics, making it particularly helpful for distinguishing engineered surfaces from surrounding natural terrain.

To place these layers on comparable scales, TPI and square root slope values were first truncated to the interval [−10, 10] and then linearly rescaled to [0, 1]. Slope in degrees was also transformed so that flatter areas map to higher values: specifically, we applied (90 − slope) / 90, yielding normalized values between 0 and 1 with low slopes near 1 and steep slopes near 0. All processing was performed in R (R Core Team. 2022) using the terra package (Hijmans 2024) for raster operations. Multi-scale filters, including the circular and annulus neighbourhoods, were implemented using functions provided in the MultiscaleDTM package (Ilich et al. 2024).

3.3. Chips and data partitions

Semantic segmentation models require inputs of consistent spatial dimensions, so each dataset was divided into fixed-size image sections, or chips, prior to training. For this study, the three-layer LSP stacks and their corresponding pixel-level labels were cut into non-overlapping tiles measuring 512 × 512 cells, providing standardized inputs for both CNN- and transformer-based architectures. Because the geomorphic features of interest occupy a small fraction of the landscape, a targeted sampling strategy was used to reduce class imbalance during training. For each dataset, 1,000 training chips were generated: 800 chips containing at least one positive-class pixel and 200 chips consisting entirely of background. This ensured that the models encountered sufficient examples of the target features while still learning from varied background terrain extents.

Validation and test sets were handled differently. To obtain an unbiased assessment of model performance, all chips produced from the validation and test regions were retained, regardless of class composition. This approach exposes each model to the full range of spatial heterogeneity present in the study areas and yields performance estimates that reflect realistic mapping conditions and landscape class proportions. Chip generation was carried out using a custom R workflow that segmented the predictor raster grids and labels into aligned tiles while preserving their spatial referencing. The resulting counts of training, validation, and test chips for each dataset are summarized in Table 1. All partitions were defined to ensure spatial independence between training, validation, and testing regions for the terrace, mine bench, and valley fill datasets. This standardized preprocessing pipeline provides a consistent basis for comparing model behaviour across landform types that differ substantially in geometry, scale, and contextual complexity.

Table 1. Summary of the number of training, validation, and test chips used for each dataset included in the study.

Dataset	Training	Validation	Test
terraceDL	1,000	1,488	4,183
minebenchDL	1,000	2,334	2,304
vfillDL	1,000	423	423

3.4. Training process

All deep learning models were developed and trained using the PyTorch library ("PyTorch [WWW Document], n.d.") within the Python programming environment ("Welcome to Python.org [WWW Document], n.d.") on a Linux-based workstation. The workstation is equipped with an AMD Ryzen Threadripper Pro 3955WX 16-core CPU, 128 GB of RAM, and three NVIDIA RTX A5000 GPUs, providing a total of 72 GB of VRAM. GPU acceleration for computations was implemented through NVIDIA's Compute Unified Device Architecture (CUDA) toolkit("CUDA Deep Neural Network (cuDNN) | NVIDIA Developer [WWW Document], n.d. "; "CUDA Toolkit - Free Tools and Training [WWW Document], n.d. "). The Segmentation Models library for PyTorch (Yakubovskiy) was used to implement the UNet model (Ronneberger, Fischer, and Brox 2015) with a ResNet-34 encoder (He et al. 2016), UNet models with SegFormer-based backbones (Xie et al. 2021), and the SegFormer models (Xie et al. 2021). Training was performed using a mini-batch size of 25 image chips over 50 epochs. Model weights were updated after each mini-batch using backpropagation and the AdamW optimizer (Loshchilov and Hutter 2017).

Different learning rates were applied depending on the initialization method, either with randomly initialized parameters or pre-trained ImageNet parameters (Deng et al. 2009). When ImageNet weights were used for model initialization, input data were normalized based on the published mean and standard deviation values of the ImageNet dataset (means: 0.485, 0.456, 0.406; standard deviations: 0.229, 0.224, 0.225). For models with pre-trained encoders that were unfrozen and updatable during training, previous experiments have shown unstable learning dynamics when a constant learning rate is applied across all model components. This instability, characterized by significant fluctuations in validation metrics, has been reported in previous studies, as noted by Ma et al. (Ma et al. 2024). To address this issue, a differentiated learning rate strategy was implemented for stable training. For the base UNet and SegUNet models with unfrozen, pre-trained encoders, learning rates were adjusted such that the first two encoder blocks used a learning rate of 1e-6, the last two encoder blocks used 1e-5, and the decoder blocks along with the classifier were trained with a learning rate of 1e-4. For SegFormer models, the learning rates were higher, with 1e-4 assigned to the first two encoder blocks, 1e-3 to the last two encoder blocks, and 1e-2 to the decoder blocks and classifier. This decision was informed by preliminary experiments where using the same learning rate configuration as the base UNet and SegUNet models led to poor convergence. Specifically, training loss plateaued while validation loss exhibited high variability, and assessment metrics for the validation dataset widely fluctuated without improvement. Increasing the learning rates for SegFormer components mitigated these issues by promoting more effective weight updates, resulting in smoother loss curves and more stable progression. This empirically tuned strategy therefore ensured more reliable convergence for SegFormer architectures. In cases where the encoder was frozen (non-trainable) or the model was initialized with random weights, a uniform learning rate of 1e-4 was applied to all trainable

components. This adaptive learning rate strategy improved training stability and ensured consistent performance across different architectures and initialization methods.

To mitigate overfitting, data augmentation techniques were employed, with random horizontal and vertical flips applied to training chips at a probability of 0.3. Due to substantial class imbalances in the datasets, where positive samples representing geomorphic features were sparse, a focal Tversky loss (Abraham and Khan 2019) was used with hyperparameters $\alpha = 0.7$, $\beta = 0.3$, and $\gamma = 0.75$. Parameters α and β were adjusted to prioritize recall for underrepresented classes, while γ was tuned to balance predictions between easy-to-detect background classes and challenging positive classes (Abraham and Khan 2019). Table 2 outlines the experiments performed for the terraceDL, minebenchDL, and vfillDL datasets. Pre-trained model experiments were conducted with both frozen and unfrozen encoders, where frozen encoders were not updatable and unfrozen encoders were updatable during training. The MiT-B5 SegFormer configuration was excluded from evaluation because preliminary attempts led to out-of-memory errors, reflecting the model's large size relative to available computational resources. As opposed to using the model state after the final training epoch, the final model was selected as the model state after the epoch that yielded the highest F1-score for the validation dataset.

Table 2. Overview of the experiments conducted. An "X" denotes the combinations tested.

Model Name	Initiation	Encoder	Frozen	Unfrozen
SgeFormer	Random	MiT-b0		X
		MiT -b1		X
		MiT -b2		X
		MiT -b3		X
		MiT -b4		X
	ImageNet	MiT -b0	X	X
		MiT -b1	X	X
		MiT -b2	X	X
		MiT -b3	X	X
		MiT -b4	X	X
SegUNet	Random	MiT -b0		X
		MiT -b1		X
		MiT -b2		X
		MiT -b3		X
		MiT -b4		X
	ImageNet	MiT -b0	X	X
		MiT -b1	X	X
		MiT -b2	X	X
		MiT -b3	X	X
		MiT -b4	X	X
UNet	Random	ResNet-34		X
	ImageNet	ResNet-34	X	X

3.5. Model assessment

Model performance was quantified using the held-out test partitions for the terraceDL, minebenchDL, and vfillDL datasets. For each model–dataset combination, pixel-wise confusion matrices were constructed from predictions on the test chips, and four standard metrics were derived: overall accuracy (OA), precision, recall, and F1-score (Table 3). OA expresses the fraction of correctly labelled pixels among all pixels in the test set and provides a coarse indication of performance. However, because the background class greatly outnumbers the anthropogenic feature class, OA can remain high even when detection of the positive class is poor. To better characterize how well each model identifies geomorphic features, we therefore place greater emphasis on precision, recall, and F1-score.

Precision describes the reliability of positive predictions by quantifying the proportion of predicted feature pixels that are truly features (i.e., low precision indicates many false positives). Recall measures how completely the feature class is recovered and is sensitive to missed detections (i.e., low recall indicates many false negatives). The F1-score combines precision and recall into a single indicator via their harmonic mean, and is particularly informative when evaluating models under class imbalance (Maxwell, Warner, and Guillén 2021a, 2021b; Farhadpour, Warner, and Maxwell 2024; Tharwat 2021). Together, these metrics provide a complementary view of performance, allowing us to compare architectures not only in terms of overall correctness but also in how they trade off false positives and false negatives when segmenting sparse anthropogenic landforms.

Table 3. Pixel-wise classification metrics derived from the confusion matrix, where TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively. OA = overall accuracy.

Metric	OA	Recall	Precision	F1-score
Equation	$\frac{TP + TN}{TP + TN + FP + FN}$	$\frac{TP}{TP + FN}$	$\frac{TP}{TP + FP}$	$\frac{2 \times Precision \times Recall}{Precision + Recall}$

4. Results

4.1. Performance

In this section, we report results for three architectures trained with different encoders, parameter initializations, and when ImageNet pretraining is used, with the encoder either frozen or unfrozen. Figure 4 shows the training-loss curves for the mineBenchDL, vfillDL, and terraceDL experiments, while Figure 5 shows the validation F1-scores computed at the end of each training epoch. Randomly initialized models are always unfrozen. On the optimization side (Figure 4), all configurations reduce loss over time, but both the rates and endpoints vary. For ImageNet initialized runs, unfreezing the encoder consistently lowers training loss relative to freezing, and this advantage persists through the final epochs for all MiT scales and tasks. Within the MiT-0 to MiT-4 columns, SegUNet (Random) reaches the lowest losses followed by SegUNet (ImageNet, unfrozen) then SegFormer (ImageNet, unfrozen). SegFormer (Random) converged slowest and finished with the highest loss. Several SegFormer–ImageNet curves exhibit mid-training instability (spikes around epochs ~30–35), whereas SegUNet curves are generally smoother. Increasing backbone depth from MiT-B0 → B4 yields modest convergence gains, most visible for ImageNet-unfrozen runs, without altering this ranking. The ResNet-34 UNet (both ImageNet and

Random) converges rapidly to low losses, often competitive with, and occasionally below, the best SegUNet settings.

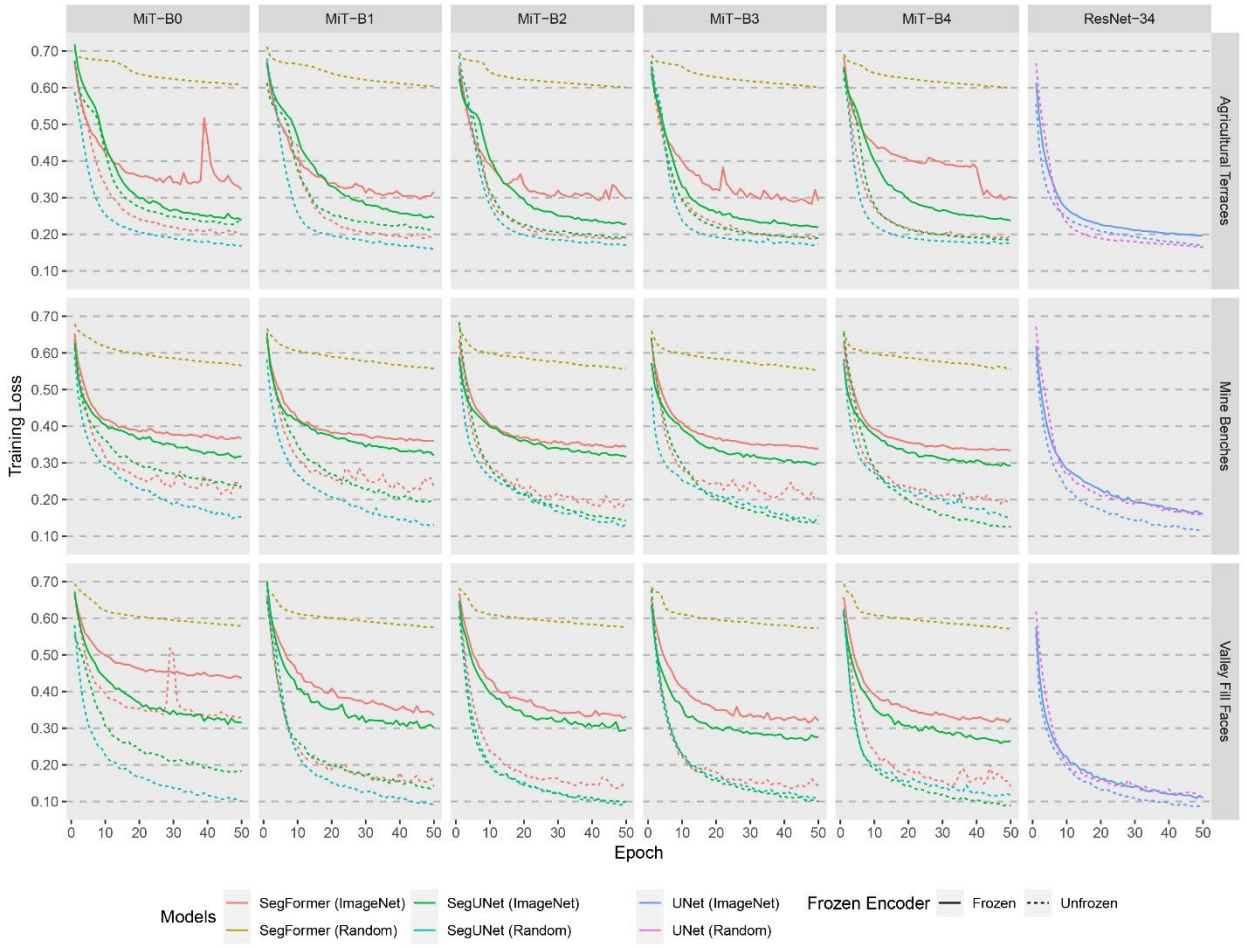

Figure 4. Training loss for terraceDL (a), minebenchDL (b), and vfillDL (c) datasets using different models, different encoders, different weight initiations, and with the encoder frozen or unfrozen across all 50 training epochs.

Generalization patterns in validation F1-scores (Figure 5) align with these optimization trends. For ImageNet initialized models, unfrozen configurations generally outperform frozen ones almost consistently; allowing encoder updates yields earlier and higher F1-score plateaus. Across MiT backbones, SegUNet (Random) and SegUNet (ImageNet, unfrozen) cluster at the top, while SegFormer (ImageNet) exhibits pronounced volatility with sharp dips, especially for the agricultural terraces problem, despite partial recovery later; SegFormer (Random) is steadier than its ImageNet counterpart but typically trails the SegUNet variants. The ResNet-34 UNet (Random and ImageNet) rises quickly and remains competitive with the best SegUNet runs, with the random variant often slightly higher and notably more stable. Most models approach asymptotic F1-scores by ~10–15 epochs, after which gains are incremental; the large oscillations are largely confined to SegFormer (ImageNet). Across tasks, both metrics indicate a consistent difficulty gradient: agricultural terraces are the hardest (higher residual losses and lower, more volatile F1-score), mine benches are intermediate, and valley fill faces are the easiest (lowest losses and highest, most stable F1-scores).

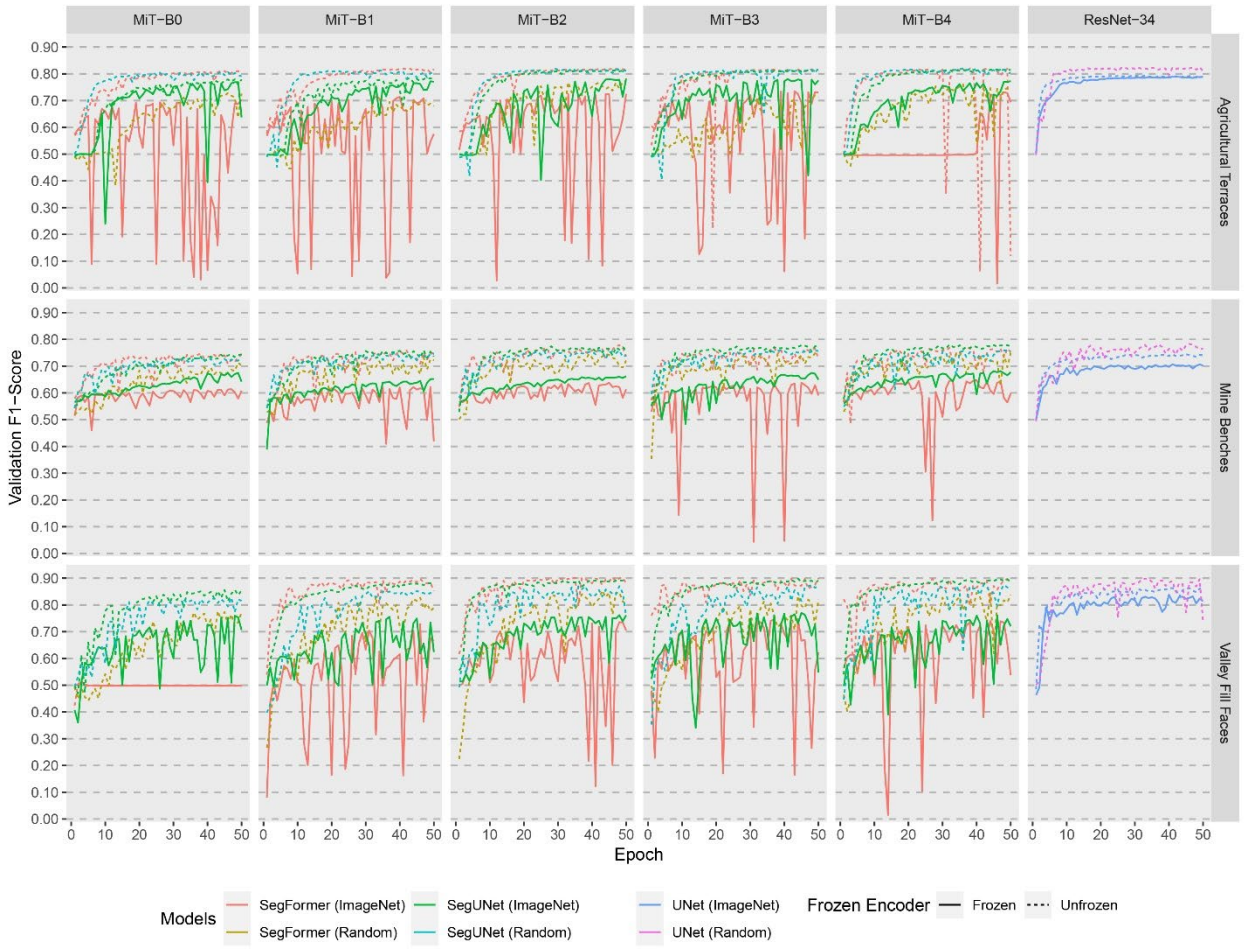

Figure 5. Validation F1-score for terraceDL (a), minebenchDL (b), and vfillDL (c) datasets using different models, different encoders, different weight initiations, and with the encoder frozen or unfrozen across all 50 training epochs. Dashed lines represent unfrozen models while solid lines represent models with a frozen encoder.

4.2. Comparison of test set predictions

Figure 6 compiles the F1-scores obtained for the withheld test sets for all configurations, with green and pink bands indicating ImageNet and random initialization, respectively, and the hatched block highlighting the UNet with a ResNet-34 encoder. The accompanying Tables 4 to 6 (agricultural terraces, mine benches, valley fill faces) report OA, F1-scores, recall, and precision for the different models, encoder configurations, and parameter initializations, with either frozen or unfrozen encoders. OA is uniformly high and less discriminative due to the large number of background pixels relative to positive case pixels for all cases (i.e., extreme class imbalance), so we emphasize F1-score together with precision and recall. Across tasks and backbones, unfreezing the encoder yields a consistent improvement on the test set, mirroring the training/validation behaviour. This effect is strongest for SegFormer, where frozen ImageNet models occupy the lower end of the distribution, while unfrozen counterparts move into the competitive range. UNet-style decoders (SegUNet and the base UNet) deliver the most reliable top-end performance and the smoothest trade-offs between precision and recall. Notably, random initialization often matches or

exceeds ImageNet pretraining, suggesting that pretraining on natural images is not universally advantageous when using lidar-derived LSP predictor variables.

Agricultural terrace extraction generally had the lowest metrics, with F1-scores spanning ~0.56–0.76. The strongest scores are achieved by SegFormer (MiT-B4, ImageNet, unfrozen) at F1-scores = 0.7589, followed closely by UNet (ResNet-34, Random) at 0.7484 and SegUNet (MiT-B2, Random) at 0.7408. Frozen ImageNet SegFormers cluster near 0.56–0.59, and gains from unfreezing can exceed +0.15–0.20 F1-score. Precision and recall are generally balanced for the best models. For mine benches, F1-scores tighten upward (~0.64–0.78). The top configuration is SegUNet (MiT-B3, ImageNet, unfrozen) at F1-scores = 0.7793, with UNet (ResNet-34, Random) close behind at 0.7711. SegFormer (MiT-B2/B3, ImageNet, unfrozen) also performs well (F1-scores = 0.772). Again, frozen ImageNet variants lag by ~0.12–0.16 F1-scores. Valley fill faces extraction generally yielded the lowest test metrics, with several models exceeding 0.80 F1-score. The UNet (ResNet-34, Random) is best at F1-scores = 0.8415, and both SegFormer (MiT-B1, ImageNet, unfrozen) and SegUNet (MiT-B2, ImageNet, unfrozen) are close (F1-scores = 0.8361 and 0.8343, respectively). Unfreezing benefits are evident in both precision and recall, while frozen ImageNet models remain near 0.61–0.63 F1-scores.

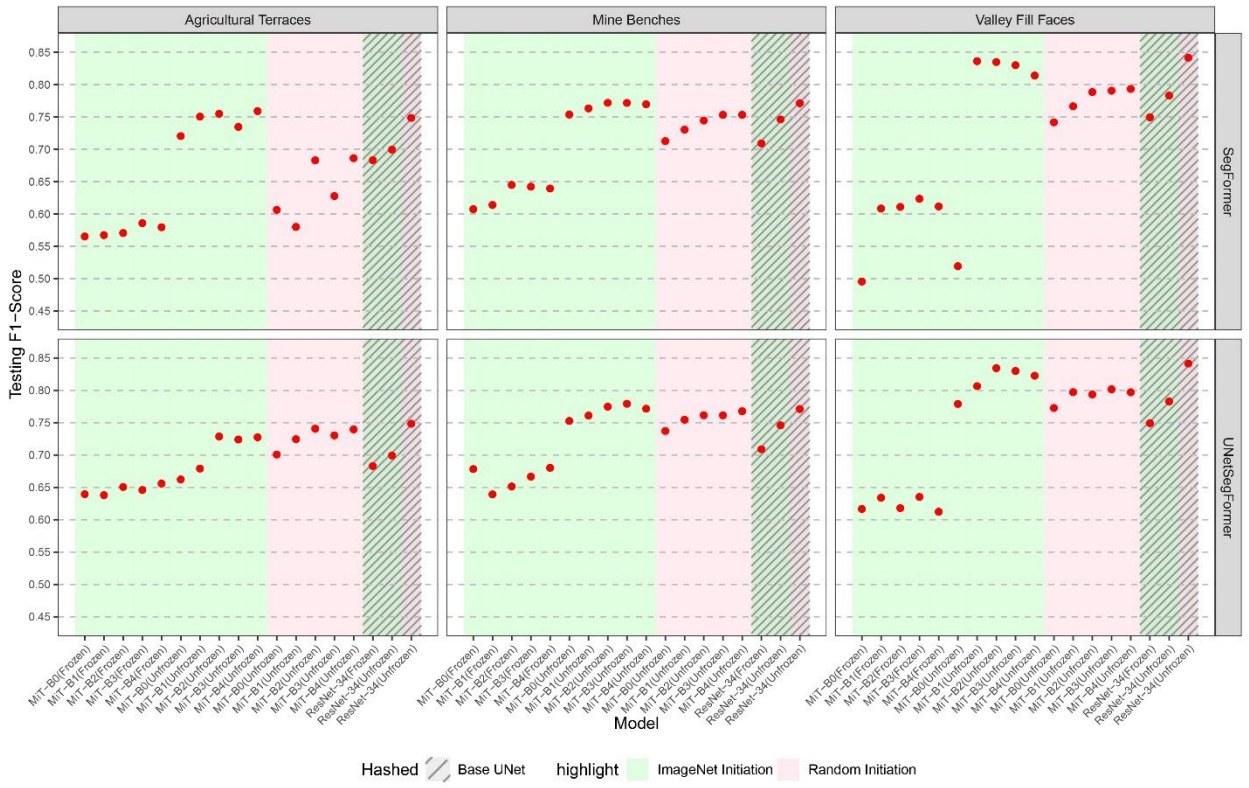

Figure 6. F1-score calculated from the withheld test data using different models and architectural configurations.

Table 4. Testing set assessment metrics for prediction of agricultural terraces. Best performance for each architecture relative to each metric is bolded.

Model	Encoder	Initiation	Frozen/ Unfrozen	OA	F1-Score	Recall	Precision
SegFormer	MiT-B0	ImageNet	Frozen	0.9884	0.5651	0.5639	0.5663
			Unfrozen	0.9937	0.7203	0.682	0.7797
		Random	Unfrozen	0.9791	0.6063	0.7107	0.5728
	MiT-B1	ImageNet	Frozen	0.9886	0.5672	0.5651	0.5694
			Unfrozen	0.9937	0.7504	0.7328	0.7708
		Random	Unfrozen	0.9656	0.5799	0.7559	0.5513
	MiT-B2	ImageNet	Frozen	0.9866	0.5705	0.5810	0.5625
			Unfrozen	0.9935	0.7548	0.7479	0.7621
		Random	Unfrozen	0.9876	0.6828	0.7634	0.6408
	MiT-B3	ImageNet	Frozen	0.9906	0.5856	0.5715	0.6073
			Unfrozen	**0.9939**	0.7345	0.6983	**0.7873**
		Random	Unfrozen	0.9786	0.6276	**0.7740**	0.5854
	MiT-B4	ImageNet	Frozen	0.9872	0.5794	0.5890	0.5717
			Unfrozen	0.9936	**0.7589**	0.7535	0.7646
		Random	Unfrozen	0.9904	0.6862	0.7098	0.6675
SegUNet	MiT-B0	ImageNet	Frozen	0.9892	0.6395	0.6548	0.6271
			Unfrozen	0.9916	0.6623	0.6497	0.6774
		Random	Unfrozen	0.9928	0.7008	0.6788	0.7292
	MiT-B1	ImageNet	Frozen	0.9908	0.6380	0.6299	0.6473
			Unfrozen	0.9924	0.6791	0.6560	0.7105
		Random	Unfrozen	0.9933	0.7246	0.7021	0.7527
	MiT-B2	ImageNet	Frozen	0.9911	0.6506	0.6420	0.6604
			Unfrozen	0.9935	0.7287	0.7035	0.7613
		Random	Unfrozen	0.9930	**0.7408**	**0.7396**	0.7421
	MiT-B3	ImageNet	Frozen	0.9910	0.6460	0.6373	0.6560
			Unfrozen	0.9933	0.7240	0.7004	0.7541
		Random	Unfrozen	0.9937	0.7304	0.6977	0.7764
	MiT-B4	ImageNet	Frozen	0.9909	0.6560	0.6515	0.6607
			Unfrozen	**0.9938**	0.7275	0.6898	**0.7844**
		Random	Unfrozen	0.9935	0.7397	0.7213	0.7614
UNet	ResNet-34	ImageNet	Frozen	0.9926	0.6830	0.6568	0.7202
			Unfrozen	0.9929	0.6993	0.6723	0.7365
		Random	Unfrozen	**0.9936**	**0.7484**	**0.7347**	**0.7640**

Table 5. Testing set assessment metrics for prediction of mine benches. Best performance for each architecture relative to each metric is bolded.

Model	Encoder	Initiation	Frozen/ Unfrozen	OA	F1-Score	Recall	Precision
SegFormer	MiT-B0	ImageNet	Frozen	0.9558	0.6073	0.6639	0.5821
			Unfrozen	0.9826	0.7537	0.7389	0.7706
		Random	Unfrozen	0.9747	0.7126	0.7507	0.6851
	MiT-B1	ImageNet	Frozen	0.9520	0.6138	0.6934	0.5840
			Unfrozen	0.9830	0.7631	0.7516	0.7757
		Random	Unfrozen	0.9758	0.7302	**0.7767**	0.6979
	MiT-B2	ImageNet	Frozen	0.9680	0.6448	0.6747	0.6242
			Unfrozen	**0.9836**	**0.7717**	0.7600	0.7846
		Random	Unfrozen	0.9792	0.7443	0.7652	0.7266
	MiT-B3	ImageNet	Frozen	0.9662	0.6421	0.6801	0.6184
			Unfrozen	0.9832	0.7716	0.7673	0.7762
		Random	Unfrozen	0.9805	0.7532	0.7677	0.7403
	MiT-B4	ImageNet	Frozen	0.9591	0.6393	0.7139	0.6058
			Unfrozen	0.9836	0.7695	0.7558	**0.7849**
		Random	Unfrozen	0.9806	0.7532	0.7656	0.7420
SegUNet	MiT-B0	ImageNet	Frozen	0.9726	0.6784	0.7033	0.6593
			Unfrozen	0.9822	0.7527	0.7436	0.7626
		Random	Unfrozen	0.9808	0.7373	0.7315	0.7433
	MiT-B1	ImageNet	Frozen	0.9649	0.6393	0.6819	0.614
			Unfrozen	0.9832	0.7611	0.745	0.7796
		Random	Unfrozen	0.9822	0.7547	0.7475	0.7623
	MiT-B2	ImageNet	Frozen	0.9633	0.6515	0.7145	0.6189
			Unfrozen	0.9838	0.7748	0.7642	0.7863
		Random	Unfrozen	0.9833	0.7615	0.7445	0.7813
	MiT-B3	ImageNet	Frozen	0.9684	0.6666	0.7116	0.6385
			Unfrozen	0.9841	**0.7793**	**0.7692**	0.7901
		Random	Unfrozen	0.9816	0.7614	0.7689	0.7544
	MiT-B4	ImageNet	Frozen	0.9732	0.6803	0.7024	0.6628
			Unfrozen	**0.9844**	0.7716	0.7472	**0.8017**
		Random	Unfrozen	0.9842	0.7678	0.7438	0.7973
UNet	ResNet-34	ImageNet	Frozen	0.9772	0.7090	0.7190	0.7000
			Unfrozen	0.9825	0.7460	0.7263	0.7697
		Random	Unfrozen	**0.9848**	**0.7711**	**0.7413**	**0.8100**

Table 6. Testing set assessment metrics for prediction of valley fill faces. Best performance for each architecture relative to each metric is bolded.

Model	Encoder	Initiation	Frozen/ Unfrozen	OA	F1-Score	Recall	Precision
SegFormer	MiT-B0	ImageNet	Frozen	0.9184	0.4954	0.5077	0.5021
			Unfrozen	0.9590	0.5192	0.5250	0.5163
		Random	Unfrozen	0.9811	0.7416	0.7642	0.7227
	MiT-B1	ImageNet	Frozen	0.9803	0.6083	0.5819	0.6636
			Unfrozen	0.9905	**0.8361**	**0.7913**	0.8982
		Random	Unfrozen	0.9856	0.7664	0.7444	0.7930
	MiT-B2	ImageNet	Frozen	0.9792	0.6108	0.5886	0.6500
			Unfrozen	**0.9907**	0.8348	0.7815	0.9144
		Random	Unfrozen	0.9874	0.7884	0.7558	0.8310
	MiT-B3	ImageNet	Frozen	0.9805	0.6233	0.5957	0.6763
			Unfrozen	0.9904	0.8299	0.7777	0.9077
		Random	Unfrozen	0.9863	0.7905	0.7840	0.7973
	MiT-B4	ImageNet	Frozen	0.9787	0.6115	0.5913	0.6446
			Unfrozen	0.9901	0.8138	0.7495	**0.9258**
		Random	Unfrozen	0.9877	0.7932	0.7597	0.8372
SegUNet	MiT-B0	ImageNet	Frozen	0.9780	0.6167	0.5999	0.6410
			Unfrozen	0.9880	0.7790	0.7258	0.8673
		Random	Unfrozen	0.9863	0.7728	0.7457	0.8069
	MiT-B1	ImageNet	Frozen	0.9816	0.6340	0.6010	0.7032
			Unfrozen	0.9891	0.8065	0.7569	0.8814
		Random	Unfrozen	0.9877	0.7973	0.7685	0.8334
	MiT-B2	ImageNet	Frozen	0.9815	0.6180	0.5859	0.6946
			Unfrozen	0.9901	**0.8343**	**0.7984**	0.8806
		Random	Unfrozen	0.9886	0.7938	0.7418	0.8760
	MiT-B3	ImageNet	Frozen	0.9796	0.6353	0.6134	0.6687
			Unfrozen	**0.9902**	0.8300	0.7833	0.8962
		Random	Unfrozen	0.9884	0.8017	0.7641	0.8523
	MiT-B4	ImageNet	Frozen	0.9825	0.6124	0.5768	0.7220
			Unfrozen	0.9901	0.8227	0.7686	**0.9056**
		Random	Unfrozen	0.9890	0.7970	0.7403	0.8911
UNet	ResNet-34	ImageNet	Frozen	0.9870	0.7493	0.6937	0.8529
			Unfrozen	0.9884	0.7829	0.7260	0.8807
		Random	Unfrozen	**0.9908**	**0.8415**	**0.7958**	**0.9046**

4.3. Qualitative spatial predictions

To visualize model performance, we applied the best models to predict spatial subsets from the respective test areas for the terraceDL, mineBenchDL, and vfillDL problems. Figures 7, 8, and 9 display the resulting predictions for the terraceDL, mineBenchDL, and vfillDL testing areas, respectively.

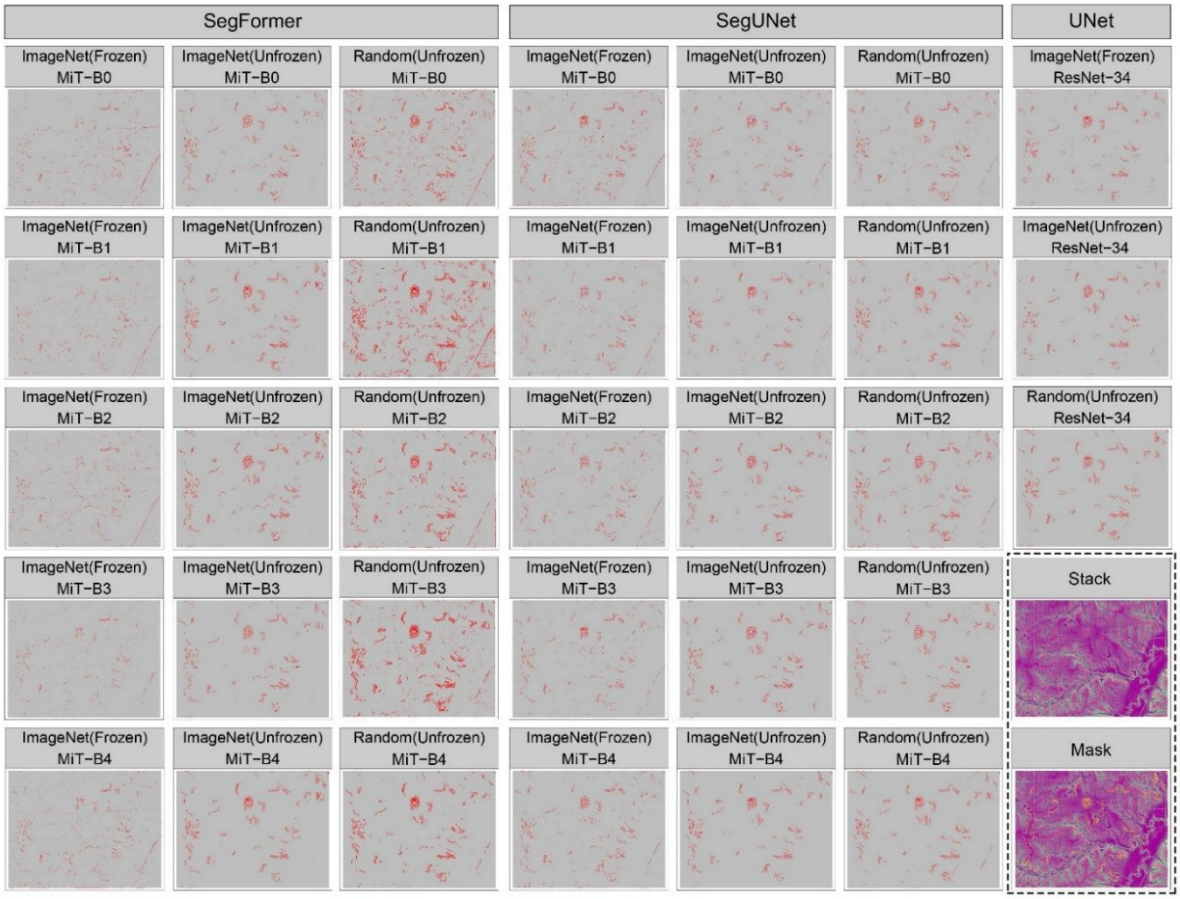

Figure 7. Comparison of agricultural terrace detection results using different models and architectural configurations.

Agricultural terraces are the most challenging: frozen ImageNet SegFormer maps are sparse and fragmented, omitting narrow segments. Unfreezing fills gaps and restores continuity, consistent with the +0.15–0.20 F1-score gains. SegUNet and UNet yield the cleanest delineations by contiguous, thicker ribbons with smoother boundaries and less speckle. Random initialization for these UNet-style models matches or exceeds ImageNet-unfrozen; SegFormer (Random) improves recall but is more commission error-prone along linear artifacts (e.g., roads/field edges). Depth helps modestly (B0→B4), without changing the ranking.

For mine benches, unfreezing again improves completeness, yielding more continuous bench polygons that capture switchbacks and curvature. SegUNet and UNet produce the most faithful shapes (preserved widths; boundaries aligned with breaks-in-slope), mirroring their top test F1-scores. Random-initiated SegUNet/UNet are at least as good as ImageNet-unfrozen configurations. SegFormer (random) boosts recall over frozen ImageNet but introduces extra commission errors along roads and cut edges. Deeper MiT backbones reduce speckle slightly; improvements from B3→B4 are minor.

Of the three problems, valley fill faces are visually easiest: the best models form compact, contiguous polygons with minimal holes. Unfreezing reduces fragmentation in SegFormer outputs; UNet (ResNet-34, Random) delivers the crispest interiors and edges, with SegUNet (MiT-B2, ImageNet, unfrozen) and SegFormer (MiT-B2, ImageNet, unfrozen) close behind, tending toward slightly higher recall (SegUNet) and precision (SegFormer). Common error modes include (i) false negatives on faint/eroded terraces, subtle benches, and thin valley-fill lobes where relief is weak, and (ii) false positives along linear anthropogenic features or steep gullies that mimic target geometry. These errors are most evident in frozen ImageNet SegFormer predictions and least frequent in UNet-style outputs.

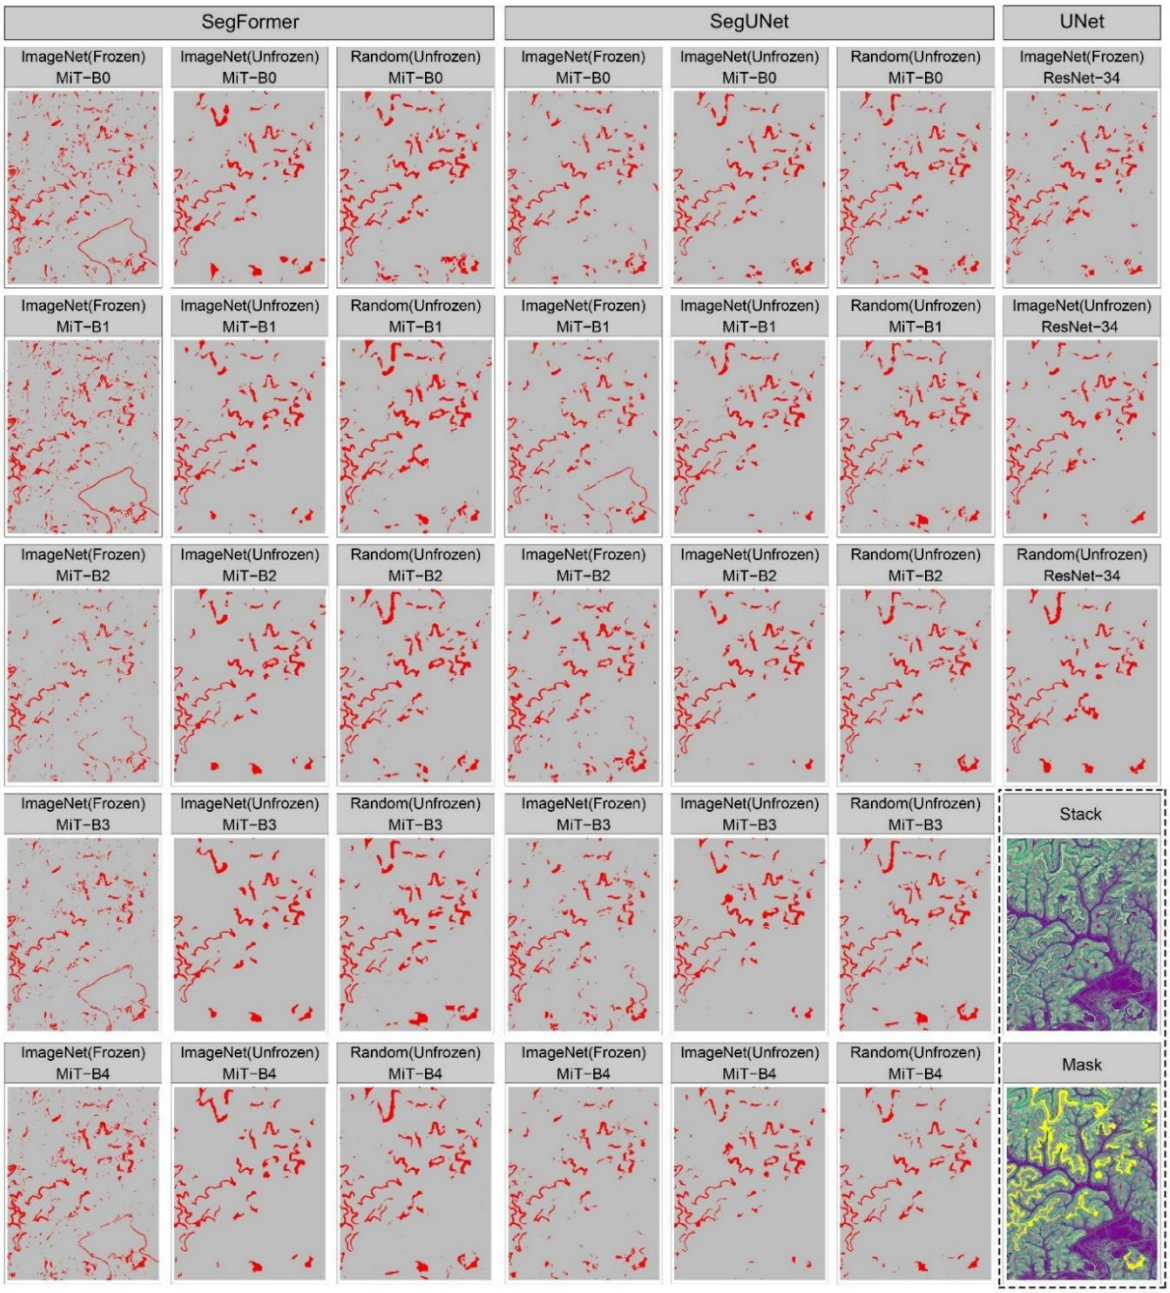

Figure 8. Comparison of historic mine bench detection results using different models and architectural configurations.

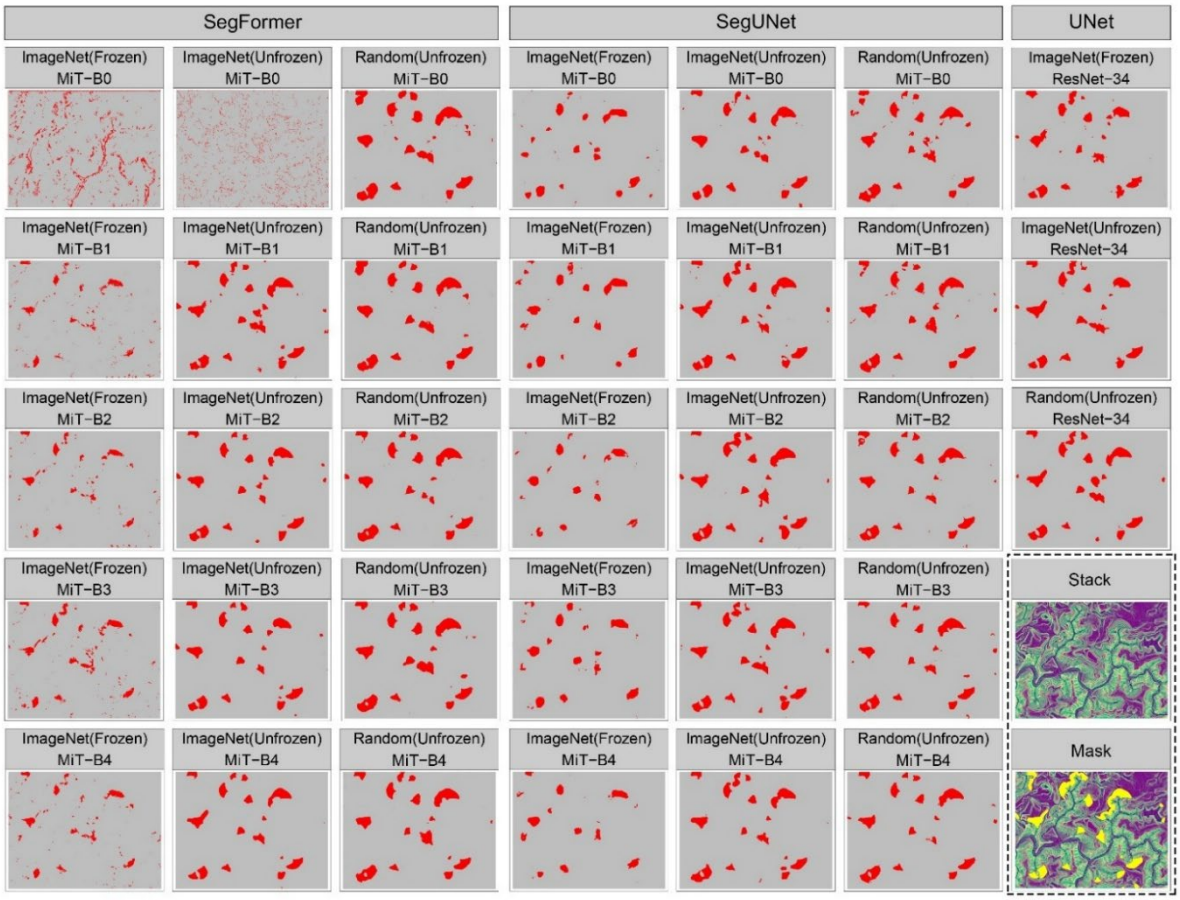

Figure 9. Comparison of valley fill face detection results using different models and architectural configurations.

4.4. Analysis of model performance variability

To further assess model robustness, we used the final models to re-calculate assessment metrics based on 10 subsamples, each consisting of a random selection of 70% from the available test chips. Figure 10 summarizes the overall accuracy (OA), F1-score, precision, and recall for every configuration across the terraceDL, mineBenchDL, and vfillDL problems; green/pink bands indicate ImageNet vs. random initialization and the hatched block highlights the base UNet (ResNet-34). Error bars show the range across subsamples, providing a direct view of metric variability.

Across all tasks and metrics, error bars are generally short, indicating that predictions are stable to which subset of test chips is sampled. This stability is most pronounced for OA, which is consistently high and least variable, reflecting class imbalance (large background relative to target pixels). As expected, F1-scores, precision, and recall display more spread and thus better differentiate configurations.

The bootstrap trends mirror Sections 4.1–4.3. Configurations with unfrozen encoders maintain higher mean F1-scores and smaller variability than their frozen counterparts, with the contrast most evident for SegFormer. In comparison, UNet-style decoders (SegUNet and the base UNet) show improved F1-scores with tight ranges, while SegFormer (ImageNet, frozen) occupies the lower end with wider spreads. Random initialization remains competitive, often matching or exceeding

ImageNet-unfrozen, without increasing variability. Backbone depth (MiT-B0→B4) yields only incremental gains and does not materially change the ordering or variance patterns.

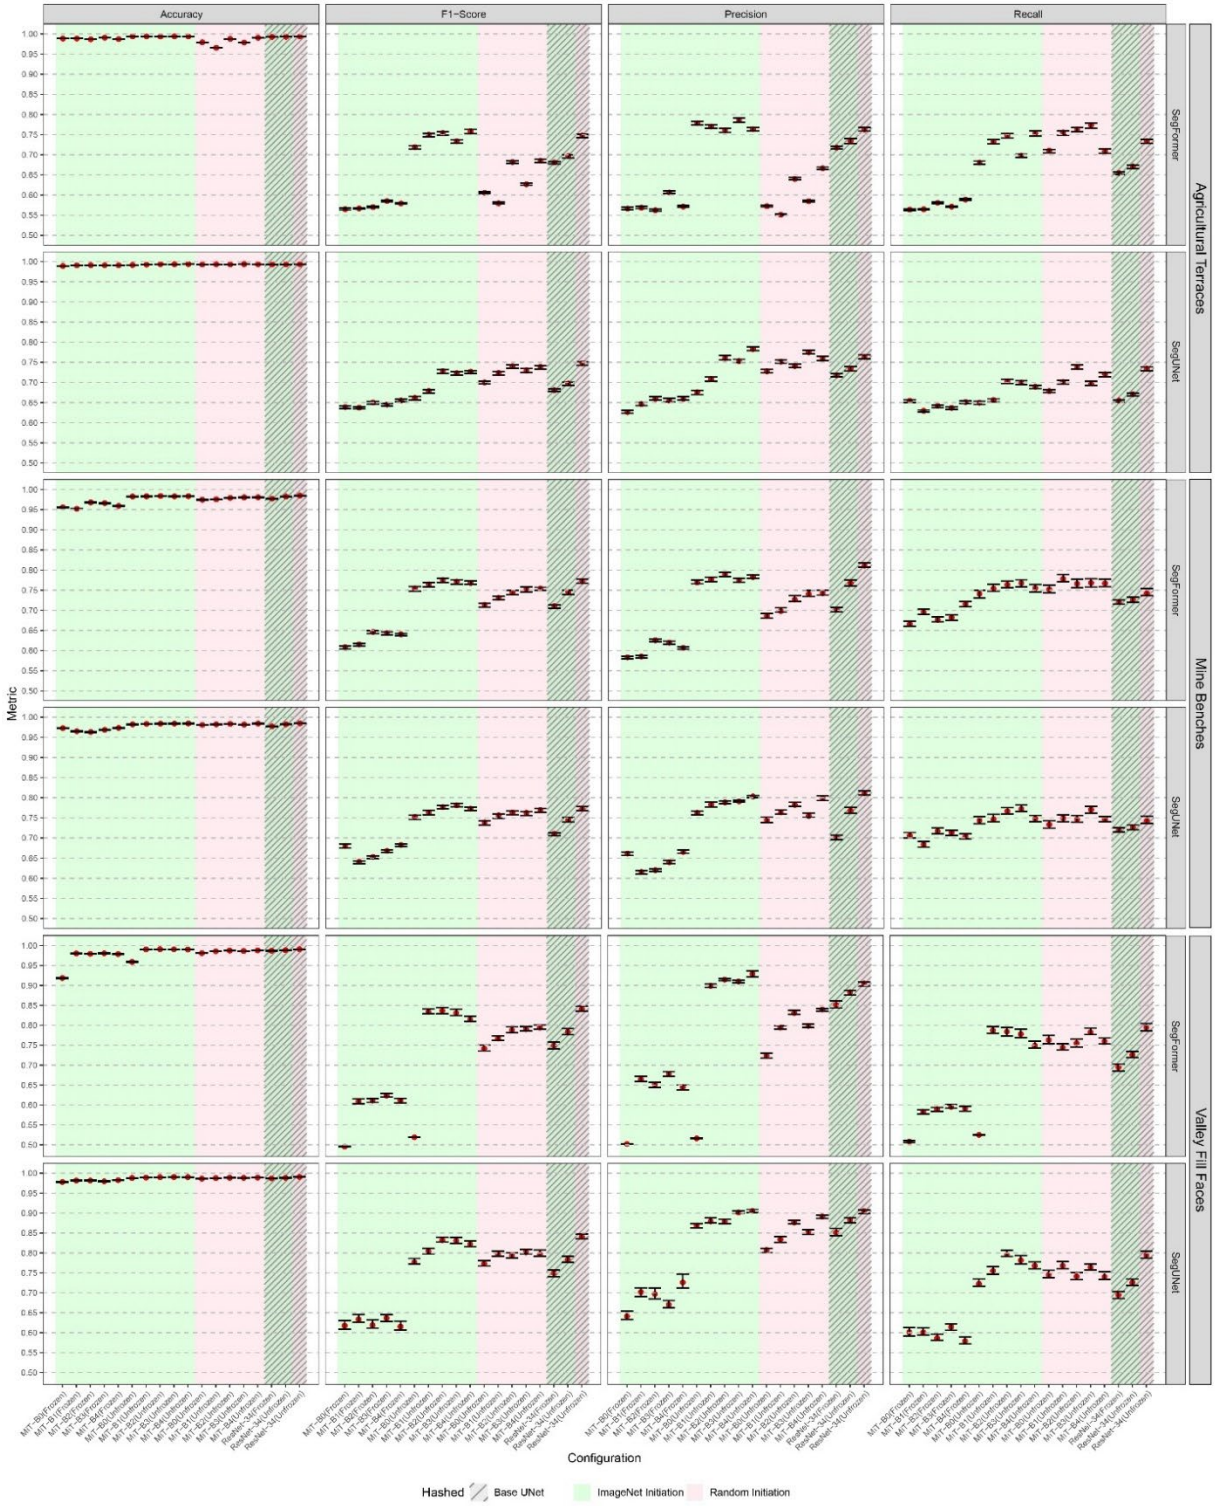

Figure 10. Variability evaluation of overall accuracy, F1-score, precision, and recall for the mineBenchDL, terraceDL, and vfillDL datasets. Error bars indicate variability across 10 test set subsamples, illustrating differences in performance stability among model configurations.

5. Discussion

This study evaluated three segmentation architectures: SegFormer, a hybrid SegUNet (SegFormer encoder + UNet-style decoder), and a fully CNN-based UNet with a ResNet-34 encoder, on three lidar-derived LSP datasets (terraceDL, mineBenchDL, and vfillDL) for anthropogenic geomorphic feature extraction tasks. To situate these models in a broader context, it is worth noting that SegFormer has demonstrated strong performance across diverse applications, including building detection (Li, Rui, Yang, Liu, et al. 2023; Li, Rui, Yang, Zuo, et al. 2023; Li, Wang, and Yan 2024), water body and wetland mapping (Zhang et al. 2023; Yang et al. 2023; Lin et al. 2023), and even medical image segmentation (Sourget et al. 2023), often outperforming CNN-based models.

Several critical findings emerged from this work, offering valuable insights into the performance and optimization of these architectures and configurations for anthropogenic geomorphic feature extraction:

End-to-end training (unfreezing the encoder) was consistently beneficial. Across MiT-B0→B4 backbones and all three problems, unfreezing ImageNet encoders lowered training loss and raised validation/test F1-scores relative to frozen counterparts; the contrast was most pronounced for SegFormer. Qualitative maps echoed this: unfrozen models closed gaps and reduced feature fragmentation.

UNet-style decoders generalized more robustly than pure SegFormer. SegUNet variants and the ResNet-34 UNet clustered at the top of validation/test F1-scores with smoother optimization; SegFormer (ImageNet) showed mid-training volatility and occasional large F1-score dips, especially for the agricultural terraces problem. On test data, UNet-style decoders provided the most reliable precision–recall trade-offs and competitive F1-score values.

Random initialization often matched or exceeded ImageNet pretraining for LSP inputs. Across tasks, random initialized SegUNet and UNet frequently equalled or surpassed ImageNet initialized models without added variability, suggesting a domain gap between ImageNet features and LSP-based geomorphic characterization. Interestingly, this was in contrast to our prior study (Maxwell, Farhadpour, and Ali 2024a).

Task difficulty tracked with landform geometry and scale. Agricultural terraces (narrow, elongated) were generally most difficult to extract while mine benches were intermediate and valley fill faces (compact) were easiest. We attribute this at least partially to higher boundary-pixel fractions for narrow features and with differing sample sizes across datasets. In practice, easier valley fill face shapes yielded several >0.80 test F1-score configurations, while agricultural terraces topped out lower. Thus, it is important to consider the shape of features and the relative proportion of edge pixels when interpreting assessment metrics.

Backbone depth for SegFormer yields modest gains; a well-tuned CNN baseline remains strong. Moving from MiT-B0to B4 slightly improved convergence and the F1-score but did not alter the architecture ranking; the ResNet-34 UNet converged quickly and was often competitive with the best MiT-SegUNet settings.

This study generally highlights the value of incorporating CNN components, as the SegUNet models generally outperformed SegFormer without the CNN-based UNet encoder. The hybrid SegUNet combines the global context awareness of SegFormer's hierarchical transformer encoder with the localized spatial feature reconstruction of UNet's decoder that may enhance semantic segmentation tasks.

As shown in Figure 11, in this hybrid approach, the encoder path comprises the hierarchical transformer blocks of SegFormer, generating multi-scale contextual feature representations where

the size of the data representations in the spatial dimensions decreases between encoder blocks using overlap patch merging. The decoder retains UNet's design, progressively upsamples feature maps through transpose convolutions, and uses skip connections from each transformer block to their corresponding decoder blocks. This architecture combines the global context awareness of transformers with the localized spatial context extraction of a CNN-based decoder, potentially achieving superior segmentation performance.

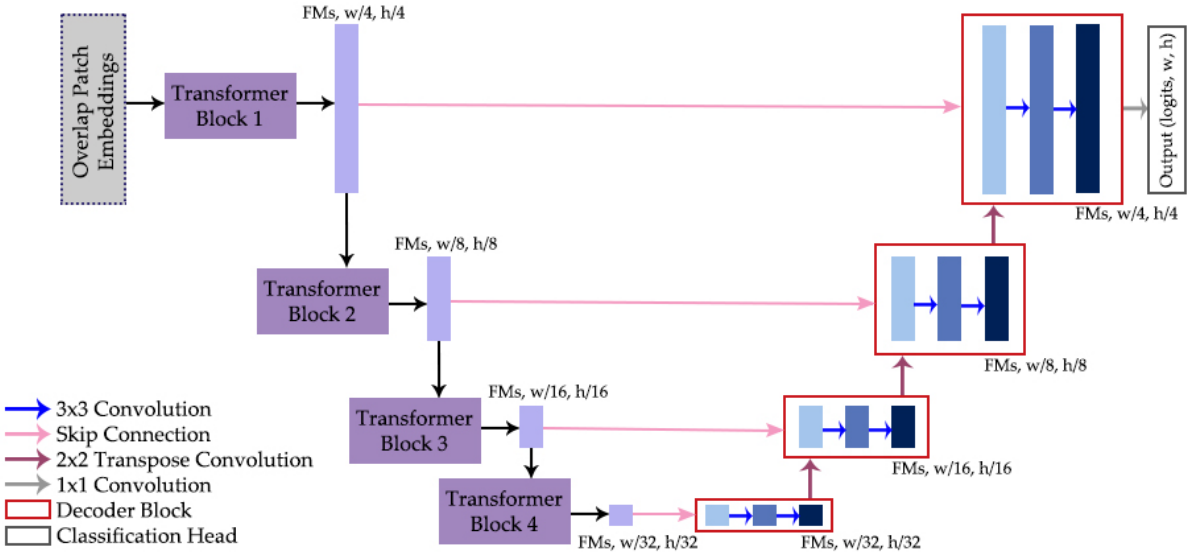

Figure 11. Conceptual design of SegUNet, a UNet model with a SegFormer backbone, showing the hierarchical transformer encoder linked via skip connections to a UNet-style decoder.

While transformer encoders are highly effective at modelling long-range dependencies and capturing global contextual relationships, they often lose fine-grained spatial detail. CNN decoders, in contrast, are adept at reconstructing local textures and refining object boundaries. One possible reason for this advantage is that CNN layers embed a strong inductive bias toward locality and translation equivariance, which helps preserve structural continuity in boundary-heavy or elongated features. Compared to a simple MLP head, a CNN-based decoder can also produce richer spatial feature representations by explicitly filtering and aggregating neighbourhood information, rather than treating all tokens independently. This may explain why CNN decoders are more effective in restoring spatial precision when paired with transformer encoders. For example, in the CTCANet framework, a transformer encoder was paired with a cascaded CNN decoder, where skip connections restored high spatial resolution details and reduced information loss across layers, resulting in more accurate change detection maps (Yin, Chen, and Zhang 2023). Similarly, the FCTN (Full-Scale Connected CNN–Transformer Network) model introduced a full-scale skip-connected CNN decoder to complement its transformer encoder, enabling both global dependency modelling and spatially precise predictions (Chen, Zhang, et al. 2023).

Other studies further emphasize this hybrid advantage. The SSNet (semantic segmentation network) architecture, for instance, combined a transformer encoder with a CNN-based decoder to ensure that global dependencies were captured without sacrificing edge sharpness and local detail, highlighting that CNN decoders are essential for refining structural features (Yao et al. 2024). Beyond individual models, recent surveys of transformers in remote sensing also confirm this trend, noting that hybrid encoder–decoder designs consistently outperform purely transformer-based or CNN-only approaches, since the transformer provides global context while the CNN

decoder restores spatial resolution and object boundaries (Aleissaee et al. 2023). Together, these studies underscore the value of incorporating CNN-based decoders alongside transformer encoders, providing both the contextual awareness and the fine spatial fidelity required for complex remote sensing tasks.

Beyond the encoder–decoder design, the mixed results for pre-training highlight a domain gap between natural image features and LSP representations. Random initialization rarely hurt and frequently helped for UNet-style models; it also improved SegFormer's performance relative to its frozen ImageNet counterpart, albeit with a tendency toward extra commission error in highly textured terrain. These outcomes suggest that when labels are sufficient and inputs diverge from RGB semantics, learning representations from scratch can be preferable to reusing ImageNet features.

Our results also highlight the value of updating the encoder parameters during training, as this consistently reduced training loss and raised validation/test F1-scores across MiT backbones and datasets, with the largest gains for SegFormer. The qualitative maps corroborate this pattern. Frozen ImageNet models, especially SegFormer, were prone to sparse, fragmented predictions and missed interiors, while unfreezing closed gaps, restored continuity of elongated features, and reduced salt-and-pepper noise. These gains persisted at test time and were evident across all three targets.

Performance differences among tasks aligned with object morphology. Valley fill faces, being compact with high interior-to-boundary ratios, achieved the highest F1-scores across multiple configurations. Mine benches were more challenging than valley fill faces but less so than terraces; their curvilinear polygons were captured well by UNet-style decoders that preserved width and continuity. Terraces, being narrow, elongated, and often low-contrast, posed the greatest difficulty; even the best models under-segmented faint segments and confused terrace-like linear artifacts (e.g., roads and field edges), which explains the precision and recall trade-offs observed.

The performance variability analysis adds a robustness perspective. Metric ranges across ten resampled test subsets were generally fairly consistent, indicating that conclusions are not driven by a particular test split. Accuracy stayed high and uninformative because of class imbalance; F1-score, precision, and recall provided more meaningful comparisons. Variability patterns echoed the architecture findings: unfrozen encoders reduced spread, UNet-style decoders paired strong central scores with narrow ranges, and differences between top configurations were modest. Notably, mine benches often exhibited the tightest error bars, agricultural terraces were moderate, and valley fill faces showed wider spreads for some SegFormer variants despite high means, underscoring that "easier" classes can still vary across models.

Practically, these insights support several recommendations for geomorphic mapping with lidar-derived LSPs. First, we recommend unfreezing encoders and using a learning-rate strategy that stably trains them end-to-end. Second, there is value in incorporating a CNN-based decoder (i.e., SegUNet or a well-tuned UNet); SegFormer can be competitive when fully fine-tuned but is more sensitive to settings. Third, do not assume ImageNet pretraining will help; random initialization can be a credible starting point. Finally, report F1-score with precision and recall, and consider robustness summaries (e.g., metric variability estimates) to quantify stability across plausible test subsets.

6. Conclusion

We compared SegFormer, a hybrid SegUNet, and a UNet with a CNN-based ResNet-34 backbone for extracting three anthropogenic geomorphic features from lidar-derived LSPs. Across backbones and datasets, the most consistent driver of performance was end-to-end training:

unfreezing encoders lowered training loss and improved validation/test F1-score, with the largest gains for SegFormer. UNet-style decoders (SegUNet and UNet) delivered the most reliable generalization and the best precision and recall balance, while SegFormer, though competitive when unfrozen, was more sensitive to training dynamics. Random initialization often surpassed ImageNet pretraining, pointing to a domain gap between natural image features and lidar-based terrain representations. Increasing the SegFormer backbone depth provided only modest improvements and did not change the architecture ranking. Qualitative maps echoed these findings, showing fewer gaps and cleaner boundaries for unfrozen, UNet-style models; a test set subsampling analysis further indicated stable performance.

Practically, we recommend unfrozen encoders and CNN-based, UNet-style decoders as strong defaults. Future work should explore domain-specific pretraining strategies and adaptations for underrepresented classes. Overall, effective geomorphic mapping from lidar-derived LSPs hinges more on how models are trained and decoded than on increasing model size. The findings here offer clear guidance for selecting and optimizing deep learning models for high spatial resolution anthropogenic geomorphic feature extraction.

Data availability

The data supporting the findings of this study are openly available in figshare. The terraceDL dataset is available at “terraceDL: A geomorphology deep learning dataset of agricultural terraces in Iowa, USA.” The mineBenchDL dataset is available at “mineBenchDL: A geomorphology deep learning dataset of historic surface coal mine benches in West Virginia, USA.” The vfillDL dataset is available at “vfillDL: A geomorphology deep learning dataset of valley fill faces resulting from mountaintop removal coal mining in southern West Virginia, eastern Kentucky, and south-western Virginia, USA.”

References

Abraham, N., and N. M. Khan. 2019. A Novel Focal Tversky Loss Function With Improved Attention U-Net for Lesion Segmentation. Paper presented at the 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), 8-11 April 2019.

Agarap, Abien Fred. 2018. "Deep learning using rectified linear units (relu)." *arXiv preprint arXiv:1803.08375*.

Albrecht, C. M., C. Fisher, M. Freitag, H. F. Hamann, S. Pankanti, F. Pezzutti, and F. Rossi. 2019. Learning and Recognizing Archeological Features from LiDAR Data. Paper presented at the 2019 IEEE International Conference on Big Data (Big Data), 9-12 Dec. 2019.

Aleissaee, Abdulaziz A., Amandeep Kumar, Rao M. Anwer, Salman Khan, Hisham Cholakkal, Gui-Song Xia, and Fahad S. Khan. 2023. "Transformers in Remote Sensing: A Survey." *Remote Sensing* 15 (7). https://doi.org/10.3390/rs15071860.

Arundel, S. T., L. A. Phillips, A. J. Lowe, J. Bobinmyer, K. S. Mantey, C. A. Dunn, E. W. Constance, and E. L. Usery. 2015. "Preparing The National Map for the 3D Elevation Program – products, process and research." *Cartography and Geographic Information Science* 42 (sup1):40-53. https://doi.org/10.1080/15230406.2015.1057229.

Azarafza, Mohammad, Mehdi Azarafza, Haluk Akgün, Peter M. Atkinson, and Reza Derakhshani. 2021. "Deep learning-based landslide susceptibility mapping." *Scientific Reports* 11 (1):24112. https://doi.org/10.1038/s41598-021-03585-1.

Bjorck, Nils, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. 2018. "Understanding batch normalization." *Advances in neural information processing systems* 31.

Bossi, G., M. Cavalli, S. Crema, S. Frigerio, B. Quan Luna, M. Mantovani, G. Marcato, L. Schenato, and A. Pasuto. 2015. "Multi-temporal LiDAR-DTMs as a tool for modelling a complex landslide: a

case study in the Rotolon catchment (eastern Italian Alps)." *Nat. Hazards Earth Syst. Sci.* 15 (4):715-722. https://doi.org/10.5194/nhess-15-715-2015.

Cao, Hu, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. 2022. Swin-unet: Unet-like pure transformer for medical image segmentation. Paper presented at the European conference on computer vision.

Cao, Jingzhi, Haiquan Wu, and Yuyou Zou. 2024. "Application of Light Remote Sensing Technology Based on Transfer Learning Algorithm in Sustainable Management of Ecological Economy."

Cao, Kaili, and Xiaoli Zhang. 2020. "An Improved Res-UNet Model for Tree Species Classification Using Airborne High-Resolution Images." *Remote Sensing* 12 (7). https://doi.org/10.3390/rs12071128.

Cazorzi, Federico, Giancarlo Dalla Fontana, Alberto De Luca, Giulia Sofia, and Paolo Tarolli. 2013. "Drainage network detection and assessment of network storage capacity in agrarian landscape." *Hydrological Processes* 27 (4):541-553. https://doi.org/https://doi.org/10.1002/hyp.9224.

Chai, Hua, Long Xia, Lei Zhang, Jiarui Yang, Zhongyue Zhang, Xiangjun Qian, Yuedong Yang, and Weidong Pan. 2021. "An Adaptive Transfer-Learning-Based Deep Cox Neural Network for Hepatocellular Carcinoma Prognosis Prediction." *Frontiers in Oncology* 11. https://doi.org/10.3389/fonc.2021.692774.

Chase, Arlen F., Diane Z. Chase, Christopher T. Fisher, Stephen J. Leisz, and John F. Weishampel. 2012. "Geospatial revolution and remote sensing LiDAR in Mesoamerican archaeology." *Proceedings of the National Academy of Sciences* 109 (32):12916-12921. https://doi.org/10.1073/pnas.1205198109.

Chen, Min, Qiangjiang Zhang, Xuming Ge, Bo Xu, Han Hu, Qing Zhu, and Xin Zhang. 2023. "A Full-Scale Connected CNN–Transformer Network for Remote Sensing Image Change Detection." *Remote Sensing* 15 (22). https://doi.org/10.3390/rs15225383.

Chen, Xin, Dongfen Li, Mingzhe Liu, and Jiaru Jia. 2023. "CNN and Transformer Fusion for Remote Sensing Image Semantic Segmentation." *Remote Sensing* 15 (18). https://doi.org/10.3390/rs15184455.

Chen, Zhen, and Yiyang Zheng. 2023. "Research on Surface Information Extraction Based on Deep Learning and Transfer Learning." *Journal of Geoscience and Environment Protection* 11 (10):67-78.

Chen, Zhengchao, Bing Zhang, Yongshun Han, Zhengli Zuo, and Xiaoyong Zhang. 2014. "Modeling Accumulated Volume of Landslides Using Remote Sensing and DTM Data." *Remote Sensing* 6 (2):1514-1537. https://doi.org/10.3390/rs6021514.

Chen, Ziyue, Bingbo Gao, and Bernard Devereux. 2017. "State-of-the-Art: DTM Generation Using Airborne LIDAR Data." *Sensors* 17 (1). https://doi.org/10.3390/s17010150.

Cheng, Ruoxi. 2023. "Expansion of the CT-scans image set based on the pretrained DCGAN for improving the performance of the CNN." *Journal of Physics: Conference Series* 2646 (1):012015. https://doi.org/10.1088/1742-6596/2646/1/012015.

"CUDA Deep Neural Network (cuDNN) | NVIDIA Developer [WWW Document], n.d. ." *https://developer.nvidia.com/cudnn (accessed 10.30.23).*

"CUDA Toolkit - Free Tools and Training [WWW Document], n.d. ." *https://developer.nvidia.com/cuda-toolkit (accessed 10.30.23).*

Deng, J., W. Dong, R. Socher, L. J. Li, Li Kai, and Fei-Fei Li. 2009. ImageNet: A large-scale hierarchical image database. Paper presented at the 2009 IEEE Conference on Computer Vision and Pattern Recognition, 20-25 June 2009.

Diakogiannis, Foivos I., François Waldner, Peter Caccetta, and Chen Wu. 2020. "ResUNet-a: A deep learning framework for semantic segmentation of remotely sensed data." *ISPRS Journal of Photogrammetry and Remote Sensing* 162:94-114. https://doi.org/https://doi.org/10.1016/j.isprsjprs.2020.01.013.

Dosovitskiy, Alexey, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, and Sylvain Gelly. 2020. "An

image is worth 16x16 words: Transformers for image recognition at scale." *arXiv preprint arXiv:2010.11929*.

Farhadpour, Sarah, and Aaron E. Maxwell. 2025. "Comparing UNet configurations for anthropogenic geomorphic feature extraction from land surface parameters." *PLOS ONE* 20 (6):e0325904. https://doi.org/10.1371/journal.pone.0325904.

Farhadpour, Sarah, Timothy A. Warner, and Aaron E. Maxwell. 2024. "Selecting and Interpreting Multiclass Loss and Accuracy Assessment Metrics for Classifications with Class Imbalance: Guidance and Best Practices." *Remote Sensing* 16 (3). https://doi.org/10.3390/rs16030533.

Fernandez-Diaz, J. C., W. E. Carter, R. L. Shrestha, S. J. Leisz, C. T. Fisher, A. M. González, D. Thompson, and S. Elkins. 2014. Archaeological prospection of north Eastern Honduras with airborne mapping LiDAR. Paper presented at the 2014 IEEE Geoscience and Remote Sensing Symposium, 13-18 July 2014.

Franklin, Steven E. 2020. "Interpretation and use of geomorphometry in remote sensing: a guide and review of integrated applications." *International Journal of Remote Sensing* 41 (19):7700-7733. https://doi.org/10.1080/01431161.2020.1792577.

Graf, Lukas, Mariano Moreno-de-las-Heras, Maurici Ruiz, Aleix Calsamiglia, Julián García-Comendador, Josep Fortesa, José A. López-Tarazón, and Joan Estrany. 2018. "Accuracy Assessment of Digital Terrain Model Dataset Sources for Hydrogeomorphological Modelling in Small Mediterranean Catchments." *Remote Sensing* 10 (12). https://doi.org/10.3390/rs10122014.

Guyot, Alexandre, Marc Lennon, Thierry Lorho, and Laurence Hubert-Moy. 2021. "Combined Detection and Segmentation of Archeological Structures from LiDAR Data Using a Deep Learning Approach." *Journal of Computer Applications in Archaeology* 4 (1):1-19. https://doi.org/https://doi.org/10.5334/jcaa.64.

He, K., X. Zhang, S. Ren, and J. Sun. 2016. Deep Residual Learning for Image Recognition. Paper presented at the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 27-30 June 2016.

Hesse, Ralf. 2010. "LiDAR-derived Local Relief Models – a new tool for archaeological prospection." *Archaeological Prospection* 17 (2):67-72. https://doi.org/https://doi.org/10.1002/arp.374.

Hijmans, Robert J. "terra: Spatial Data Analysis." Accessed 2024. https://CRAN.R-project.org/package=terra.

Huang, Gao, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. 2017. Densely connected convolutional networks. Paper presented at the Proceedings of the IEEE conference on computer vision and pattern recognition.

Ilich, A., V. Lecours, B. Misiuk, and S. Murawski. 2024. "MultiscaleDTM: Multi-Scale Geomorphometric Terrain Attributes (v0.8.3). Zenodo." In.

Ioffe, Sergey. 2015. "Batch normalization: Accelerating deep network training by reducing internal covariate shift." *arXiv preprint arXiv:1502.03167*.

Jaboyedoff, Michel, Thierry Oppikofer, Antonio Abellán, Marc-Henri Derron, Alex Loye, Richard Metzger, and Andrea Pedrazzini. 2012. "Use of LIDAR in landslide investigations: a review." *Natural Hazards* 61 (1):5-28. https://doi.org/10.1007/s11069-010-9634-2.

Jin, Youngsaeng, David Han, and Hanseok Ko. 2021. "TrSeg: Transformer for semantic segmentation." *Pattern Recognition Letters* 148:29-35. https://doi.org/https://doi.org/10.1016/j.patrec.2021.04.024.

Khan, Rayyan Azam, Yigang Luo, and Fang-Xiang Wu. 2022. "RMS-UNet: Residual multi-scale UNet for liver and lesion segmentation." *Artificial Intelligence in Medicine* 124:102231. https://doi.org/https://doi.org/10.1016/j.artmed.2021.102231.

Krithika alias AnbuDevi, M., and K. Suganthi. 2022. "Review of Semantic Segmentation of Medical Images Using Modified Architectures of UNET." *Diagnostics* 12 (12). https://doi.org/10.3390/diagnostics12123064.

Kugelman, Jason, Joseph Allman, Scott A. Read, Stephen J. Vincent, Janelle Tong, Michael Kalloniatis, Fred K. Chen, Michael J. Collins, and David Alonso-Caneiro. 2022. "A comparison of deep

learning U-Net architectures for posterior segment OCT retinal layer segmentation." *Scientific Reports* 12 (1):14888. https://doi.org/10.1038/s41598-022-18646-2.

Li, Deliang, Haokun Wang, and Long Yan. 2024. "Building Extraction on SegFormer Model and Post Processing from Remote Sensing Images." *PREPRINT (Version 1) available at Research Square*. https://doi.org/https://doi.org/10.21203/rs.3.rs-4201466/v1.

Li, Meilin, Jie Rui, Songkun Yang, Zhi Liu, Liqiu Ren, Li Ma, Qing Li, Xu Su, and Xibing Zuo. 2023. "Method of Building Detection in Optical Remote Sensing Images Based on SegFormer." *Sensors* 23 (3). https://doi.org/10.3390/s23031258.

Li, Meilin, Jie Rui, Songkun Yang, Xibing Zuo, Yan Pang, and Yongkai Shi. 2023. The method of building detection in remote sensing images based on improved Segformer. Paper presented at the Proc.SPIE.

Li, Sijin, Liyang Xiong, Guoan Tang, and Josef Strobl. 2020. "Deep learning-based approach for landform classification from integrated data sources of digital elevation model and imagery." *Geomorphology* 354:107045. https://doi.org/https://doi.org/10.1016/j.geomorph.2020.107045.

Lillesand, Thomas, Ralph W Kiefer, and Jonathan Chipman. 2015. *Remote sensing and image interpretation*: John Wiley & Sons.

Lin, Xufeng, Youwei Cheng, Gong Chen, Wenjing Chen, Rong Chen, Demin Gao, Yinlong Zhang, and Yongbo Wu. 2023. "Semantic Segmentation of China's Coastal Wetlands Based on Sentinel-2 and Segformer." *Remote Sensing* 15 (15). https://doi.org/10.3390/rs15153714.

Liu, Naihao, Tao He, Yajun Tian, Bangyu Wu, Jinghuai Gao, and Zongben Xu. 2020. "Common-azimuth seismic data fault analysis using residual UNet." *Interpretation* 8 (3):SM25-SM37. https://doi.org/10.1190/INT-2019-0173.1.

Long, Jonathan, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. Paper presented at the Proceedings of the IEEE conference on computer vision and pattern recognition.

Loshchilov, Ilya, and Frank Hutter. 2017. "Decoupled weight decay regularization." *arXiv preprint arXiv:1711.05101*.

Luo, Ping, Xinjiang Wang, Wenqi Shao, and Zhanglin Peng. 2018. "Towards understanding regularization in batch normalization." *arXiv preprint arXiv:1809.00846*.

Ma, Lei, Yu Liu, Xueliang Zhang, Yuanxin Ye, Gaofei Yin, and Brian Alan Johnson. 2019. "Deep learning in remote sensing applications: A meta-analysis and review." *ISPRS Journal of Photogrammetry and Remote Sensing* 152:166-177. https://doi.org/https://doi.org/10.1016/j.isprsjprs.2019.04.015.

Ma, Yuchi, Shuo Chen, Stefano Ermon, and David B. Lobell. 2024. "Transfer learning in environmental remote sensing." *Remote Sensing of Environment* 301:113924. https://doi.org/https://doi.org/10.1016/j.rse.2023.113924.

Maxwell, Aaron. 2023a. "terraceDL: A geomorphology deep learning dataset of agricultural terraces in Iowa, USA. figshare. Dataset." In.

Maxwell, Aaron 2023b. "vfillDL: A geomorphology deep learning dataset of valley fill faces resulting from mountaintop removal coal mining (southern West Virginia, eastern Kentucky, and southwestern Virginia, USA). figshare. Dataset." In.

Maxwell, Aaron E., Sarah Farhadpour, and Muhammad Ali. 2024a. "Exploring Transfer Learning for Anthropogenic Geomorphic Feature Extraction from Land Surface Parameters Using UNet." *Remote Sensing* 16 (24). https://doi.org/10.3390/rs16244670.

Maxwell, Aaron E., William E. Odom, Charles M. Shobe, Daniel H. Doctor, Michelle S. Bester, and Tobi Ore. 2023. "Exploring the Influence of Input Feature Space on CNN-Based Geomorphic Feature Extraction From Digital Terrain Data." *Earth and Space Science* 10 (5):e2023EA002845. https://doi.org/https://doi.org/10.1029/2023EA002845.

Maxwell, Aaron E., and Charles M. Shobe. 2022. "Land-surface parameters for spatial predictive mapping and modeling." *Earth-Science Reviews* 226:103944. https://doi.org/https://doi.org/10.1016/j.earscirev.2022.103944.

Maxwell, Aaron E., Timothy A. Warner, and Luis Andrés Guillén. 2021a. "Accuracy Assessment in Convolutional Neural Network-Based Deep Learning Remote Sensing Studies—Part 1: Literature Review." *Remote Sensing* 13 (13):2450.

———. 2021b. "Accuracy Assessment in Convolutional Neural Network-Based Deep Learning Remote Sensing Studies—Part 2: Recommendations and Best Practices." *Remote Sensing* 13 (13):2591.

Maxwell, Aaron, Sarah Farhadpour, and Muhammad Ali. 2024b. "mineBenchDL: A geomorphology deep learning dataset of historic surface coal mine benches in West Virginia, USA." In.: figshare.

McNeely, ROBIN, AA Logan, JOSHUA Obrecht, J Giglierano, and CALVIN Wolter. 2017. "Iowa best management practices (BMP) mapping project handbook." *Iowa State University (ISU)[Ames, IA]*.

Miller, Andrew J., and Nicolas P. Zégre. 2014. "Mountaintop Removal Mining and Catchment Hydrology." *Water* 6 (3):472-499. https://doi.org/10.3390/w6030472.

Novakovsky, Gherman, Manu Saraswat, Oriol Fornes, Sara Mostafavi, and Wyeth W. Wasserman. 2021. "Biologically relevant transfer learning improves transcription factor binding prediction." *Genome Biology* 22 (1):280. https://doi.org/10.1186/s13059-021-02499-5.

Palmer, M. A., E. S. Bernhardt, W. H. Schlesinger, K. N. Eshleman, E. Foufoula-Georgiou, M. S. Hendryx, A. D. Lemly, et al. 2010. "Mountaintop Mining Consequences." *Science* 327 (5962):148-149. https://doi.org/10.1126/science.1180543.

Pires de Lima, Rafael, and Kurt Marfurt. 2020. "Convolutional Neural Network for Remote-Sensing Scene Classification: Transfer Learning Analysis." *Remote Sensing* 12 (1). https://doi.org/10.3390/rs12010086.

"PyTorch [WWW Document], n.d.". *https://www.pytorch.org (accessed 12.31.20).*

R Core Team. 2022. "R: A Language and Environment for Statistical Computing." In.: R Foundation for Statistical Computing, Vienna, Austria.

Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. Paper presented at the Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Cham.

Ross, Matthew R. V., Brian L. McGlynn, and Emily S. Bernhardt. 2016. "Deep Impact: Effects of Mountaintop Mining on Surface Topography, Bedrock Structure, and Downstream Waters." *Environmental Science & Technology* 50 (4):2064-2074. https://doi.org/10.1021/acs.est.5b04532.

Sampson, Christopher C., Andrew M. Smith, Paul D. Bates, Jeffrey C. Neal, and Mark A. Trigg. 2016. "Perspectives on Open Access High Resolution Digital Elevation Models to Produce Global Flood Hazard Layers." *Frontiers in Earth Science* 3. https://doi.org/10.3389/feart.2015.00085.

Sha, Youyang, Yonghong Zhang, Xuquan Ji, and Lei Hu. 2021. "Transformer-unet: Raw image processing with unet." *arXiv preprint arXiv:2109.08417*.

Sharma, Siddharth, Simone Sharma, and Anidhya Athaiya. 2020. "Activation Functions in Neural Networks." *International Journal of Engineering Applied Sciences and Technology*.

Sofia, Giulia, Giancarlo Dalla Fontana, and Paolo Tarolli. 2014. "High-resolution topography and anthropogenic feature extraction: testing geomorphometric parameters in floodplains." *Hydrological Processes* 28 (4):2046-2061. https://doi.org/https://doi.org/10.1002/hyp.9727.

Sourget, Théo, Syed Nouman Hasany, Fabrice Mériaudeau, and Caroline Petitjean. 2023. Can SegFormer be a True Competitor to U-Net for Medical Image Segmentation? Paper presented at the 27th Conference on Medical Image Understanding and Analysis 2023, Aberdeen, United Kingdom, 2023-07-19.

Strudel, Robin, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. 2021. Segmenter: Transformer for semantic segmentation. Paper presented at the Proceedings of the IEEE/CVF international conference on computer vision.

Sugarbaker, Larry J, Eric W Constance, Hans Karl Heidemann, Allyson L Jason, Vicki Lukas, David L Saghy, and Jason M Stoker. 2014. "The 3D Elevation Program initiative: a call for action." In.: US Geological Survey.

Tariku, Girma, Isabella Ghiglieno, Gianni Gilioli, Fulvio Gentilin, Stefano Armiraglio, and Ivan Serina. 2023. "Automated Identification and Classification of Plant Species in Heterogeneous Plant Areas

Using Unmanned Aerial Vehicle-Collected RGB Images and Transfer Learning." *Drones* 7 (10). https://doi.org/10.3390/drones7100599.

Tarolli, Paolo. 2014. "High-resolution topography for understanding Earth surface processes: Opportunities and challenges." *Geomorphology* 216:295-312. https://doi.org/https://doi.org/10.1016/j.geomorph.2014.03.008.

Tharwat, Alaa. 2021. "Classification assessment methods." *Applied Computing and Informatics* 17 (1):168-192. https://doi.org/10.1016/j.aci.2018.08.003.

Thi Ngo, Phuong Thao, Mahdi Panahi, Khabat Khosravi, Omid Ghorbanzadeh, Narges Kariminejad, Artemi Cerda, and Saro Lee. 2021. "Evaluation of deep learning algorithms for national scale landslide susceptibility mapping of Iran." *Geoscience Frontiers* 12 (2):505-519. https://doi.org/https://doi.org/10.1016/j.gsf.2020.06.013.

van der Meij, W. Marijn, Erik W. Meijles, Diego Marcos, Tom T. L. Harkema, Jasper H. J. Candel, and Gilbert J. Maas. 2022. "Comparing geomorphological maps made manually and by deep learning." *Earth Surface Processes and Landforms* 47 (4):1089-1107. https://doi.org/https://doi.org/10.1002/esp.5305.

van Westen, Cees J., Enrique Castellanos, and Sekhar L. Kuriakose. 2008. "Spatial data for landslide susceptibility, hazard, and vulnerability assessment: An overview." *Engineering Geology* 102 (3):112-131. https://doi.org/https://doi.org/10.1016/j.enggeo.2008.03.010.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. "Attention Is All You Need." *Advances in neural information processing systems, 30*.

Wang, Libo, Rui Li, Ce Zhang, Shenghui Fang, Chenxi Duan, Xiaoliang Meng, and Peter M. Atkinson. 2022. "UNetFormer: A UNet-like transformer for efficient semantic segmentation of remote sensing urban scene imagery." *ISPRS Journal of Photogrammetry and Remote Sensing* 190:196-214. https://doi.org/https://doi.org/10.1016/j.isprsjprs.2022.06.008.

"Welcome to Python.org [WWW Document], n.d.". *https://www.python.org/ (accessed 1.5.21).*

Wickham, James, Petra Bohall Wood, Matthew C. Nicholson, William Jenkins, Daniel Druckenbrod, Glenn W. Suter, Michael P. Strager, Christine Mazzarella, Walter Galloway, and John Amos. 2013. "The Overlooked Terrestrial Impacts of Mountaintop Mining." *BioScience* 63 (5):335-348. https://doi.org/10.1525/bio.2013.63.5.7.

Xiang, Jianjian, Jia Liu, Du Chen, Qi Xiong, and Chongjiu Deng. 2023. "CTFuseNet: A Multi-Scale CNN-Transformer Feature Fused Network for Crop Type Segmentation on UAV Remote Sensing Imagery." *Remote Sensing* 15 (4). https://doi.org/10.3390/rs15041151.

Xie, Enze, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. 2021. "SegFormer: Simple and efficient design for semantic segmentation with transformers." *Advances in neural information processing systems* 34:12077-12090.

Xu, Yuexue, Hongchun Zhu, Changyu Hu, Haiying Liu, and Yu Cheng. 2022. "Deep learning of DEM image texture for landform classification in the Shandong area, China." *Frontiers of Earth Science* 16 (2):352-367. https://doi.org/10.1007/s11707-021-0884-y.

Yakubovskiy, Pavel. "Segmentation Models." *GitHub repository. https://github.com/qubvel-org/segmentation_models.pytorch*.

Yang, Xiao, Mingwei Chen, Chengjun Yu, Haozhe Huang, Xiaobin Yue, Bei Zhou, and Ming Ni. 2023. "WaterSegformer: A lightweight model for water body information extraction from remote sensing images." *IET Image Processing* 17 (3):862-871. https://doi.org/https://doi.org/10.1049/ipr2.12678.

Yao, M., Y. Zhang, G. Liu, and D. Pang. 2024. "SSNet: A Novel Transformer and CNN Hybrid Network for Remote Sensing Semantic Segmentation." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 17:3023-3037. https://doi.org/10.1109/JSTARS.2024.3349657.

Yin, Mengmeng, Zhibo Chen, and Chengjian Zhang. 2023. "A CNN-Transformer Network Combining CBAM for Change Detection in High-Resolution Remote Sensing Images." *Remote Sensing* 15 (9). https://doi.org/10.3390/rs15092406.

Yuan, Qiangqiang, Huanfeng Shen, Tongwen Li, Zhiwei Li, Shuwen Li, Yun Jiang, Hongzhang Xu, et al. 2020. "Deep learning in environmental remote sensing: Achievements and challenges." *Remote Sensing of Environment* 241:111716. https://doi.org/https://doi.org/10.1016/j.rse.2020.111716.

Zhang, L., L. Zhang, and B. Du. 2016. "Deep Learning for Remote Sensing Data: A Technical Tutorial on the State of the Art." *IEEE Geoscience and Remote Sensing Magazine* 4 (2):22-40. https://doi.org/10.1109/MGRS.2016.2540798.

Zhang, Tianyi, Chenhao Qin, Weibin Li, Xin Mao, Liyun Zhao, Biao Hou, and Licheng Jiao. 2023. "Water Body Extraction of the Weihe River Basin Based on MF-SegFormer Applied to Landsat8 OLI Data." *Remote Sensing* 15 (19). https://doi.org/10.3390/rs15194697.

Zhao, Yan, and Caiying Feng. 2022. "Remote Sensing Image Water Body Recognition Algorithm Based on Deep Convolution Generating Network and Combined Features." *Wireless Communications and Mobile Computing* 2022 (1):9932251. https://doi.org/https://doi.org/10.1155/2022/9932251.

Zhu, X. X., D. Tuia, L. Mou, G. S. Xia, L. Zhang, F. Xu, and F. Fraundorfer. 2017. "Deep Learning in Remote Sensing: A Comprehensive Review and List of Resources." *IEEE Geoscience and Remote Sensing Magazine* 5 (4):8-36. https://doi.org/10.1109/MGRS.2017.2762307.