# Evaluating the Impact of Spatial Locality in Few-Shot Segmentation for Seismic Facies

1st Gabriel Lima Nunes
*DCC*
*UFMG*
Belo Horizonte, Brazil
nuneslima@dcc.ufmg.br

2st George L. M. Teodoro
*DCC*
*UFMG*
Belo Horizonte, Brazil
george@dcc.ufmg.br

***Abstract*—Seismic facies segmentation is a fundamental task in subsurface interpretation, yet the high cost and subjectivity of manual labeling severely limit the availability of annotated data for training supervised models. Few-Shot Semantic Segmentation (FSSS) provides a promising alternative by enabling segmentation from a small number of annotated examples, but its application to seismic data remains underexplored and sensitive to the spatial properties of seismic volumes. In this work, we investigate the combined impact of spatial data partitioning and support sample locality on FSSS for seismic facies interpretation. We evaluate four representative few-shot segmentation models, PANet, PFENet, HSNet, and ASNet, under a controlled 1-way 1-shot setting using two 3D seismic datasets, Parihaka and Penobscot. Two partitioning strategies are considered: dense sub-volumes (Large Rectangular Prisms) and sparse global sampling (Equally Distant Slices), together with two support selection modes: unrestricted sampling and locality-constrained sampling from the same seismic slice as the query. Experimental results demonstrate that enforcing support locality consistently improves segmentation performance across models and datasets. Moreover, combining sparse global partitioning with locality-aware support selection achieves competitive or superior performance while substantially reducing annotation effort. Our analysis further reveals notable differences in model robustness under severe class imbalance, highlighting the importance of domain-aware episodic design for reliable few-shot seismic facies segmentation.**

***Index Terms*—Seismic facies segmentation, Few-shot semantic segmentation, Seismic interpretation, Support selection, Spatial locality, Data partitioning, Low-label learning**

## I. Introduction

Seismic data interpretation is a fundamental component of subsurface characterization in the oil and gas industry, supporting critical decisions such as reservoir delineation and optimal borehole placement [1]. A central task in this process is seismic facies analysis, which aims to identify and characterize geological units based on reflection patterns observed in seismic data [2], [3]. Beyond hydrocarbon exploration, seismic facies interpretation has also gained relevance in emerging applications such as carbon capture and storage monitoring [4].

Traditionally, seismic facies analysis has relied on manual interpretation, where experts delineate facies boundaries by visually inspecting seismic sections. Although effective, this process is time-consuming, labor-intensive, and inherently subjective, with interpretation quality strongly dependent on expert experience, particularly in complex geological settings [5]–[7]. These limitations have motivated the adoption of automated approaches based on deep learning (DL), especially convolutional neural networks (CNNs), which have shown promising performance in seismic facies classification and segmentation tasks [3].

Despite recent advances, supervised DL methods remain heavily constrained by their dependence on large volumes of pixel-wise annotated data. In geosciences, such datasets are scarce due to the high cost and specialized expertise required for manual labeling [3], [8]. This limitation significantly hinders the scalability and practical deployment of conventional supervised semantic segmentation methods in seismic interpretation workflows.

Few-Shot Semantic Segmentation (FSSS) has emerged as a viable alternative for addressing data scarcity by enabling models to segment target classes using only a small number of annotated examples [4], [8]. In this paradigm, models are trained to generalize from limited support information, making FSSS particularly appealing for seismic interpretation scenarios where annotations are scarce. However, FSSS remains relatively underexplored in the seismic domain, where data exhibits unique spatial characteristics [3].

Seismic volumes are characterized by strong lateral continuity within individual inlines or crosslines, alongside significant structural variability across distant slices. As a result, the spatial relationship between training data, support samples, and query images plays a critical role in few-shot performance, especially under extreme label scarcity [3], [9]. While prior studies have independently investigated spatial partitioning strategies and few-shot learning in seismic settings, the combined impact of data partitioning and support locality during few-shot episode construction remains largely unexplored.

In this work, we address this gap by systematically analyzing how spatial data partitioning and support locality affect Few-Shot Semantic Segmentation of seismic facies. We evaluate two contrasting partitioning strategies, Large Rectangular Prisms (LRP) and Equally Distant Slices (EDS), together with two support selection modes, Anywhere and Same Slice, across four representative FSSS models: PANet [10], PFENet [11], HSNet [12], and ASNet [13]. Experiments are conducted on the Parihaka and Penobscot 3D seismic datasets under a controlled 1-way 1-shot regime.

The main contributions of this work are threefold: (i) we quantify the impact of support locality, showing that Same Slice support consistently outperforms unconstrained support selection; (ii) we demonstrate that combining spatially sparse partitioning with locality-aware support yields strong performance with reduced annotation effort; and (iii) we provide a comparative analysis of model robustness under severe class imbalance, highlighting stability differences among commonly used few-shot segmentation architectures.

## II. Related Work

### A. Supervised Seismic Facies Segmentation Under Data Scarcity

Deep learning has been widely adopted for seismic facies segmentation by adapting semantic segmentation architectures originally developed for natural images [3]. Encoder-decoder models, particularly U-Net and its variants [14]–[16], are commonly employed due to their ability to combine multi-scale contextual information while preserving spatial detail. Other approaches, such as DeepLab-style architectures, leverage dilated convolutions to expand receptive fields and capture broader geological structures [17].

Despite architectural refinements, these methods rely on dense, pixel-wise annotations, which are costly and difficult to obtain in geophysical applications. As a result, supervised seismic segmentation methods struggle to scale to large volumes or new datasets, motivating the adoption of learning paradigms explicitly designed for low-label regimes.

### B. Few-Shot Semantic Segmentation

Few-Shot Semantic Segmentation (FSSS) addresses the data scarcity problem by learning to segment novel classes from a limited number of annotated support examples, typically using episodic meta-learning [8]. Most state-of-the-art FSSS methods were originally developed for natural image benchmarks and can be broadly categorized based on how they model the interaction between query and support samples.

Prototype-based methods, such as PANet [10], represent each class using a feature prototype and perform segmentation via similarity-based matching. PFENet [11] extends this paradigm by incorporating prior-guided feature enrichment to alleviate spatial inconsistency between query and support features. These approaches are computationally efficient but often compress spatial information during prototype construction.

Correlation-based models aim to preserve dense spatial relationships between query and support images. HSNet [12] constructs multi-level hypercorrelation tensors to capture fine-grained correspondence, allowing improved localization at the cost of increased sensitivity to support-query mismatch. More recently, the Integrative Few-Shot Learning (iFSL) framework introduced ASNet [13], which utilizes attention mechanisms over semantic correlation maps and extends few-shot segmentation to joint classification and segmentation tasks.

It is important to note that Few-Shot Semantic Segmentation models do not learn explicit class detectors during training. Instead, they learn how to perform class-conditional segmentation based on the support example provided in each episode. Consequently, even if all facies classes appear somewhere in the training partition, the model does not memorize or internalize them. The evaluation on unseen query patches therefore remains challenging regardless of the number or diversity of classes present in the training set, as the model must rely solely on the support sample to infer the target class in each episode.

While these methods demonstrate strong performance on natural image benchmarks, they make implicit assumptions, such as random support selection and spatial homogeneity, that may not hold in seismic settings.

### C. Low-Label Learning for Seismic Segmentation

Several approaches have been proposed to mitigate annotation scarcity specifically in seismic interpretation. AdaSemSeg [4] adapts few-shot learning principles to seismic facies segmentation by combining contrastive pretraining with Gaussian process-based adaptation, enabling cross-dataset generalization without fine-tuning. Other studies employ self-supervised learning strategies to pretrain seismic encoders using pretext tasks such as rotation prediction and slice ordering [18].

Although these methods reduce dependence on labeled data, they largely focus on representation learning and model adaptation. They do not explicitly investigate how the spatial organization of training data or the locality of support samples affects few-shot segmentation performance.

### D. Locality and Support Selection in Seismic Few-Shot Segmentation

Seismic volumes exhibit strong lateral continuity within individual slices and increasing variability across spatially distant sections [18]. This characteristic implies that the spatial relationship between support and query samples is likely to influence few-shot segmentation performance. Prior work has explored spatial partitioning strategies, showing that sparse sampling approaches such as Equally Distant Slices (EDS) can achieve competitive performance relative to dense subvolume training while significantly reducing annotation effort [9].

In parallel, studies in the general few-shot literature have shown that non-random, similarity-based support selection can substantially improve segmentation accuracy compared to random sampling [19]. However, these approaches typically rely on learned similarity measures and do not incorporate explicit physical constraints.

In the context of seismic interpretation, locality-aware support selection represents a physically meaningful and computationally simple alternative. The present work advances this line of research by systematically analyzing the combined effects of spatial partitioning (LRP versus EDS) and support locality (Anywhere versus Same Slice) across multiple benchmark FSSS models, bridging a gap between few-shot learning theory and the structural properties of seismi

## III. METHODOLOGY

This work analyzes the impact of spatial data partitioning and support locality on Few-Shot Semantic Segmentation (FSSS) for seismic facies interpretation. We adopt established few-shot segmentation frameworks and adapt their episodic construction to the seismic domain by explicitly controlling how training data are partitioned and how support samples are selected relative to the query image.

### A. Few-Shot Formulation

Seismic facies segmentation is formulated as a Few-Shot Semantic Segmentation problem following the standard episodic learning paradigm [13]. All experiments are conducted under a 1-way 1-shot setting, where each episode targets a single facies class using one annotated support example.

In each episode, the query image $I_q$ is a $400 \times 400$ 2D seismic patch. The support image $S$ is also a $400 \times 400$ patch that contains the target facies, guaranteeing class presence without requiring class purity. This formulation is consistent with standard FSSS methods, where support images may include background and other classes [10]–[13]. Aside from support construction, the underlying few-shot frameworks are preserved.

### B. Datasets

Experiments are conducted on two public 3D seismic facies datasets [3], [9]. The Parihaka dataset contains six facies classes and exhibits strong class imbalance, with the dominant class representing approximately 42% of the voxels. The Penobscot dataset comprises eight facies intervals and presents an even more challenging imbalance scenario, with the majority class accounting for approximately 57% of the volume.

These datasets provide representative and realistic testbeds for evaluating few-shot segmentation under severe label scarcity and geological variability.

### C. Spatial Data Partitioning

To control spatial diversity during training, we evaluate two volume partitioning strategies introduced in prior seismic studies [9].

The *Large Rectangular Prisms (LRP)* strategy selects a contiguous subvolume for training, providing dense spatial coverage. In our experiments, this corresponds to using 57.96% of the Parihaka volume and 56.08% of the Penobscot volume for training.

The *Equally Distant Slices (EDS)* strategy instead samples thin, uniformly spaced slices across the volume. We use five inlines and five crosslines, each one pixel thick, resulting in substantially reduced labeled coverage: 1.49% of the Parihaka volume and 2.07% of the Penobscot volume.

Despite the drastic reduction in labeled samples, EDS exposes the model to a broader range of geological variability, making it a strong candidate for low-label seismic learning scenarios.

### D. Support Selection Strategies

Support locality is explicitly controlled during episode construction using two strategies. In the *Anywhere* mode, the support patch is sampled from any spatial location within the training region defined by the selected partitioning strategy. In the *Same Slice* mode, the support patch is sampled from the same inline or crossline as the query patch, enforcing a strict locality constraint.

All other episode parameters remain unchanged across support selection modes, ensuring that performance differences can be attributed solely to support locality.

### E. Training Protocol and Models

All models are trained for 200 epochs using episodic learning under identical experimental conditions. No architectural modifications are introduced, and the same backbone network (ResNet-101) is used in all experiments.

We evaluate four representative few-shot segmentation models originally proposed for natural images: PANet [10], PFENet [11], HSNet [12], and ASNet [13]. Using a fixed backbone and training protocol isolates the effects of data partitioning and support locality.

### F. Evaluation Metrics

Segmentation performance is evaluated using two complementary metrics: mean Intersection over Union (mIoU) and Macro F1. The mIoU measures per-class overlap between predictions and ground truth, making it suitable for highly imbalanced seismic datasets. Macro F1 complements this by computing the unweighted average F1-score across classes, penalizing cases where the model fails to identify minority facies.

Together, these metrics provide a more complete assessment of few-shot performance, capturing both segmentation accuracy (mIoU) and class-level balance (Macro F1). Near-zero values in either metric are used to identify collapse cases, where predictions become dominated by the majority class.

## IV. RESULTS

This section presents the performance of the four evaluated Few-Shot Semantic Segmentation models under different combinations of spatial data partitioning and support locality. Results are reported for the Parihaka and Penobscot datasets using two complementary metrics: Macro F1 and mIoU. Tables I and II summarize the outcomes for each dataset under the 1-way 1-shot regime.

TABLE I
MACRO F1 | MIOU RESULTS ON THE PARIHAKA DATASET.

| Model | LRP-A | LRP-SS | EDS-A | EDS-SS |
|---|---|---|---|---|
| PANet | 63.04 \| 43.98 | 63.04 \| 43.98 | 73.43 \| 48.28 | 73.43 \| 48.28 |
| PFENet | 67.46 \| 48.60 | 67.46 \| 48.60 | 76.84 \| 50.89 | 76.84 \| 50.89 |
| HSNet | 40.62 \| 22.47 | 40.62 \| 22.47 | 55.54 \| 30.51 | 55.54 \| 30.51 |
| ASNet | 26.72 \| 10.73 | 26.72 \| 10.73 | 15.35 \| 0.00 | 15.35 \| 0.00 |

TABLE II
MACRO F1 | MIOU RESULTS ON THE PENOBSCOT DATASET.

| Model | LRP-A | LRP-SS | EDS-A | EDS-SS |
|---|---|---|---|---|
| PANet | 46.89 \| 27.59 | 51.14 \| 26.71 | 49.39 \| 24.57 | 49.04 \| 24.91 |
| PFENet | 51.76 \| 28.67 | 47.57 \| 26.31 | 43.41 \| 22.55 | 52.09 \| 28.78 |
| HSNet | 24.90 \| 8.20 | 21.72 \| 6.36 | 21.22 \| 6.35 | 21.48 \| 6.76 |
| ASNet | 21.53 \| 6.09 | 11.75 \| 0.00 | 11.67 \| 0.00 | 22.76 \| 7.75 |

### A. Overall Performance

Across both datasets, PFENet and PANet achieve the strongest results in terms of both Macro F1 and mIoU, indicating better class-level balance and more accurate segmentation. HSNet exhibits moderate performance on Parihaka but performs poorly on Penobscot, particularly in Macro F1, highlighting difficulty in detecting minority facies. ASNet presents the highest instability, with several configurations collapsing to majority-class predictionS, visible as near-zero mIoU and very low Macro F1.

Performance on Penobscot is consistently lower than on Parihaka across all models and metrics, reflecting the dataset's stronger class imbalance and more heterogeneous facies distribution.

### B. Impact of Support Locality

Support locality has a consistent effect across most models and partitioning strategies. In the majority of cases, configurations using the *Same Slice* strategy outperform their corresponding *Anywhere*-based counterparts, regardless of whether LRP or EDS partitioning is employed. This effect is particularly evident for PFENet and ASNet on the Parihaka dataset, where enforcing support locality leads to large performance gains.

Although a small number of settings show marginal differences or slight performance degradation, no systematic disadvantage is observed when restricting the support to the same seismic slice as the query image. These results indicate that locality-aware support selection generally provides more informative guidance in few-shot seismic segmentation scenarios.

### C. Effect of Data Partitioning Strategy

Comparing spatial partitioning strategies, EDS-based configurations frequently match or outperform LRP-based setups, despite relying on a more spatially sparse training set. On the Parihaka dataset, the highest mIoU values for all evaluated models occur under the EDS-SS configuration.

On Penobscot, the effect of partitioning is more heterogeneous. While PFENet and ASNet benefit from the EDS-SS setup, PANet achieves its best results under LRP-based partitioning. Overall, replacing dense LRP partitioning with the sparse EDS strategy does not lead to systematic performance degradation.

### D. Model Stability

Model collapse, characterized by near-zero mIoU values, occurs predominantly for ASNet and, to a lesser extent, for HSNet, particularly on the Penobscot dataset. In these cases, the models produce predictions dominated by a single class, resulting in negligible overlap with the target facies.

PANet and PFENet exhibit substantially fewer collapse cases and maintain more consistent performance across all evaluated configurations, indicating greater robustness to variations in both support locality and spatial data partitioning.

## V. DISCUSSION

The results reported in the previous section demonstrate that both support locality and spatial data partitioning play a decisive role in Few-Shot Semantic Segmentation of seismic facies. Although the evaluated models were originally designed for natural images, their performance in the seismic domain is highly sensitive to how few-shot episodes are constructed, indicating that domain-aware experimental design is essential.

### A. Effect of Support Locality

Across nearly all evaluated configurations, restricting the support sample to the same seismic slice as the query leads to consistent improvements in segmentation performance. This behavior can be attributed to the strong lateral continuity that characterizes seismic facies within individual inlines or crosslines. When support and query patches originate from the same slice, they tend to share similar reflectivity patterns, structural context, and noise characteristics, facilitating more reliable feature matching during few-shot inference.

In contrast, the *Anywhere* strategy may select support samples that are spatially distant from the query and, consequently, exhibit markedly different geological structures. In few-shot settings, where the model must extrapolate from a single support example, such variability reduces the relevance of the support signal and degrades segmentation performance. The observed gains obtained with Same Slice support suggest that spatial proximity serves as an effective proxy for similarity in seismic facies appearance.

### B. Interaction Between Partitioning and Locality

The combination of Equally Distant Slices (EDS) with Same Slice support consistently yields strong performance across models, particularly on the Parihaka dataset. This result highlights a complementary effect between global spatial diversity and local consistency. While EDS exposes the model to a wide range of geological variability during training, the Same Slice restriction ensures that few-shot inference relies on locally coherent information.

These findings indicate that dense spatial coverage during training is not a strict requirement for effective few-shot segmentation in seismic data. Instead, selectively sampling representative slices across the volume, combined with locality-aware episodic design, can achieve competitive or superior results while substantially reducing annotation effort.

### C. Model Robustness and Collapse Behavior

Model stability varies significantly across the evaluated methods. PANet and PFENet demonstrate the most consistent performance across both datasets and experimental configurations. Their reliance on prototype-based representations and explicit feature aggregation appears to confer greater robustness under severe class imbalance and limited supervision. In particular, PFENet exhibits strong resilience to both sparse partitioning and locality constraints.

Conversely, ASNet shows pronounced susceptibility to model collapse, especially under EDS-A configurations. In several cases, the model converges to predicting predominantly the dominant class, resulting in near-zero mIoU. This behavior suggests that attention-based correlation mechanisms may be more sensitive to support-query mismatch and class imbalance when applied to seismic data. HSNet exhibits intermediate behavior, achieving competitive results when locality constraints are enforced but degrading substantially when support selection is unconstrained.

### D. Dataset-Specific Difficulty

Performance differences between the Parihaka and Penobscot datasets are consistent across all models and configurations. Penobscot presents a notably more challenging scenario, with lower overall mIoU and Macro F1 and increased instability. This difficulty can be attributed to its larger number of classes, stronger class imbalance, and more heterogeneous seismic patterns. The persistence of locality gains on both datasets, however, indicates that the benefits of Same Slice support are not dataset-specific but rather reflect a fundamental property of seismic data.

### E. Implications for Few-Shot Seismic Segmentation

Taken together, these results demonstrate that naive application of existing few-shot segmentation frameworks to seismic data is suboptimal. Instead, incorporating domain-specific constraints, particularly spatial locality, into episodic construction is critical for achieving stable and accurate segmentation. The findings suggest that future advances in few-shot seismic interpretation may rely less on increasingly complex architectures and more on principled, physically informed strategies for data partitioning and support selection.

## VI. Conclusion and Future Work

This work investigated the role of spatial data partitioning and support locality in Few-Shot Semantic Segmentation of seismic facies. By systematically evaluating four representative few-shot segmentation models under different partitioning strategies and support selection modes, we demonstrated that episodic design choices have a decisive impact on segmentation performance and model stability. In particular, constraining the support sample to the same seismic slice as the query consistently improves performance across datasets and models. Moreover, the combination of sparse global partitioning through Equally Distant Slices and locality-aware support selection achieves competitive or superior results compared to dense training schemes, while substantially reducing annotation requirements.

The analysis further revealed notable differences in model robustness. Prototype-based approaches, especially PANet and PFENet, show greater stability under severe class imbalance and limited supervision, whereas correlation- and attention-based models are more susceptible to collapse when support-query mismatch occurs. These findings highlight that effective few-shot seismic segmentation depends not only on model architecture but also on physically informed construction of few-shot episodes that respect the inherent structure of seismic data.

Several directions are identified for future work. First, extending the experimental analysis to include the Netherlands F3 dataset would enable a broader evaluation across seismic volumes with different geological characteristics. Second, increasing the number of shots and considering multi-way settings would provide insight into how the observed locality effects scale with larger support sets and more complex episodic configurations. Third, more advanced support selection strategies, beyond strict spatial constraints, merit investigation, particularly approaches that balance structural similarity and spatial proximity [19]. Another important direction is evaluating cross-dataset generalization, assessing whether models trained on one seismic dataset can effectively adapt to unseen datasets without retraining. Finally, incorporating and benchmarking few-shot models designed specifically for seismic data, such as AdaSemSeg [4], would allow a direct comparison between generic computer vision approaches and domain-specific solutions.

Together, these directions aim to further advance few-shot seismic facies segmentation toward practical, scalable interpretation workflows with minimal annotation effort.

## References


[1] B. Abid, B. M. Khan, and R. A. Memon, "Seismic facies segmentation using ensemble of convolutional neural networks," *Wireless Communications and Mobile Computing*, vol. 2022, no. 1, p. 7762543, 2022.

[2] L. Baroni, R. M. Silva, R. S. Ferreira, D. Civitarese, D. Szwarcman, and E. V. Brazil, "Penobscot dataset: Fostering machine learning development for seismic interpretation," 2019. [Online]. Available: https://arxiv.org/abs/1903.12060

[3] B. A. Monteiro and et al., "Literature review on deep learning for the segmentation of seismic images," *Earth-Science Reviews*, vol. 258, p. 104955, 2024.

[4] S. Saha and R. Whitaker, "Adasemseg: An adaptive few-shot semantic segmentation of seismic facies," 2025. [Online]. Available: https://arxiv.org/abs/2501.16760

[5] T. Ore and D. Gao, "Seissegdiff: A label-efficient few-shot texture segmentation diffusion model for seismic facies classification," *Computers Geosciences*, vol. 196, p. 105823, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0098300424003066

[6] H. Kaur, N. Pham, S. Fomel, Z. Geng, L. Decker, B. Gremillion, M. Jervis, R. Abma, and S. Gao, "A deep learning framework for seismic facies classification," *Interpretation*, vol. 11, no. 1, pp. T107–T116, 2023.

[7] A. U. Waldeland, A. C. Jensen, L.-J. Gelius, and A. H. S. Solberg, "Convolutional neural networks for automated seismic interpretation," *The Leading Edge*, vol. 37, no. 7, pp. 529–537, 2018.

[8] N. Catalano and M. Matteucci, "Few shot semantic segmentation: a review of methodologies, benchmarks, and open challenges," 2024. [Online]. Available: https://arxiv.org/abs/2304.05832

[9] G. Lima, G. Amarante, W. Barreiros Jr. *et al.*, "A systematic evaluation methodology of deep learning on seismic facies classification," *Earth Science Informatics*, vol. 18, p. 530, 2025. [Online]. Available: https://doi.org/10.1007/s12145-025-02024-3

[10] K. Wang, J. H. Liew, Y. Zou, D. Zhou, and J. Feng, "Panet: Few-shot image semantic segmentation with prototype alignment," 2020. [Online]. Available: https://arxiv.org/abs/1908.06391

[11] Z. Tian, H. Zhao, M. Shu, Z. Yang, R. Li, and J. Jia, "Prior guided feature enrichment network for few-shot segmentation," 2020. [Online]. Available: https://arxiv.org/abs/2008.01449

[12] J. Min, D. Kang, and M. Cho, "Hypercorrelation squeeze for few-shot segmentation," 2021. [Online]. Available: https://arxiv.org/abs/2104.01538

[13] D. Kang and M. Cho, "Integrative few-shot learning for classification and segmentation," 2022. [Online]. Available: https://arxiv.org/abs/2203.15712

[14] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," 2015. [Online]. Available: https://arxiv.org/abs/1505.04597

[15] X. Chai, W. Nie, K. Lin, G. Tang, T. Yang, J. Yu, and W. Cao, "An open-source package for deep-learning-based seismic facies classification: Benchmarking experiments on the seg 2020 open data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–19, 2022.

[16] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, B. Glocker, and D. Rueckert, "Attention u-net: Learning where to look for the pancreas," in *Medical Imaging with Deep Learning*, 2018. [Online]. Available: https://openreview.net/forum?id=Skft7cijM

[17] Z. Tian, X. Lai, L. Jiang, S. Liu, M. Shu, H. Zhao, and J. Jia, "Generalized few-shot semantic segmentation," 2022. [Online]. Available: https://arxiv.org/abs/2010.05210

[18] B. A. A. Monteiro, H. Oliveira, and J. A. d. Santos, "Self-supervised learning for seismic image segmentation from few-labeled samples," *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022.

[19] W. Shao, H. Qi, and X. Dong, "A learnable support selection scheme for boosting few-shot segmentation," *Pattern Recognition*, vol. 148, p. 110202, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0031320323008993