Improving GPR Detection Accuracy under Limited Data via Structure-Constrained Transfer Learning

Ming Li[1], Christopher Gomez[1], Shusuke Miyata[2], Norifumi Hotta[3]
[1] Sabo Laboratory, Faculty of Maritime Sciences, Kobe University, Kobe, Japan
[2] Disaster Prevention Research Institute, Kyoto University, Kyoto, Japan
[3] Department of Forest Science, Graduate School of Agricultural and Life Sciences, The University of Tokyo, Tokyo, Japan

I. Introduction

Ground Penetrating Radar (GPR) has been widely utilized in erosion control (sabo) engineering owing to its non-destructive nature, rapid data acquisition capability, and high spatial resolution. It has been applied in various tasks such as subsurface structural characterization, inspection of erosion control facilities, and detection of floating wood [1]. With the increasing availability of GPR data, deep learning techniques have emerged as the dominant paradigm for GPR image interpretation in recent years. However, its generalization remains limited due to differences in subsurface targets, survey conditions, and the scarcity of labeled data. Unlike general image recognition tasks, GPR images exhibit similar reflection patterns across different targets and environments due to their underlying physical imaging mechanisms, which makes them well suited for transfer learning. Motivated by these observations, we propose a symmetry-constrained transfer learning approach for GPR image analysis, which explicitly incorporates domain-specific structural characteristics into the transfer process and leads to improved adaptation performance under limited labeled data conditions.

II. Methodology

In this work, we consider a transfer learning setting for GPR-based subsurface structure identification, where a model trained for cavity detection serves as the source model and is adapted to utility (hyperbolic target) detection in the target domain. When the cavity-trained model is directly applied to utility detection, a substantial performance degradation is observed, indicating limited cross-target generalization.

In standard GPR target detection, the model is trained using a detection loss composed of a classification term and a localization regression term. Given a GPR B-scan image x and its corresponding annotations, the baseline training objective is defined as

$$\mathcal{L}_{\mathrm{base}} = \mathcal{L}_{\mathrm{cls}} + \mathcal{L}_{\mathrm{reg}},$$

Fine-tuning with limited target-domain samples often leads to suboptimal performance, motivating a symmetry-based consistency constraint guided by domain-specific structural priors. As illustrated in Fig. 1, GPR images of subsurface utilities typically exhibit an approximate hyperbolic structure with inherent symmetry characteristics. To incorporate this prior knowledge, we introduce a symmetry-based consistency loss.

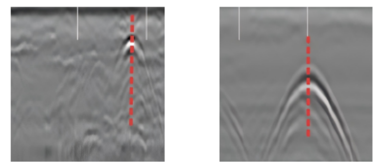

Fig. 1. Approximate symmetry in GPR B-scan images of subsurface utilities.

Let $T(\cdot)$ denote a horizontal flipping operator applied to a GPR B-scan image. The symmetry constraint enforces consistency between the prediction $f(x)$, and the inverse-transformed prediction $T^{-1}(f(T(x)))$. Accordingly, the symmetry loss is defined as

$$\mathcal{L}_{\mathrm{sym}} = \left\| f(x) - T^{-1}\big(f(T(x))\big) \right\|_2,$$

In practice, the symmetry loss is implemented at the prediction level as the sum of classification and bounding box regression consistency losses, computed between original predictions and inverse-transformed predictions of flipped B-scan images. The final training objective is formulated as

$$\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{base}} + \lambda\,\mathcal{L}_{\mathrm{sym}},$$

The value of λ is selected empirically and fixed for all target-domain data ratios to ensure a fair comparison. By enforcing symmetric prediction behavior, the proposed loss encourages structure-consistent outputs that reflect the geometric symmetry of GPR hyperbolic reflections, leading to improved robustness under limited target-domain samples.

III. Results

The experiments are conducted on a publicly available GPR dataset for intelligent recognition of subsurface utilities and cavities [2].

First, a cavity detection model is trained using the cavity dataset, with 70% of the data for training, 15% for validation, and the remaining 15% for testing. As shown in Table I, the trained model achieves an mAP@0.5 of 89.6%, indicating strong detection performance on the cavity domain. However, when the cavity-trained model is directly applied to utility detection without adaptation, the performance drops significantly to an mAP@0.5 of 33.8%. This performance degradation can be attributed to the substantial differences in structural characteristics between cavities and utilities in

TABLE I
CROSS-DOMAIN GPR DETECTION PERFORMANCE.

Train data	Test data	mAP@0.5 (%)
Cavities	Cavities	89.6
Cavities	Utilities	33.8

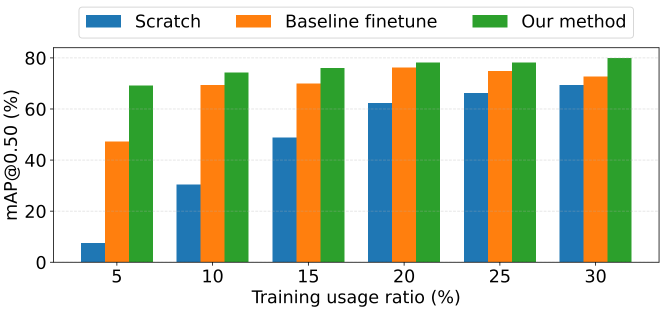

Fig. 2. Detection performance (mAP@0.5) of different training strategies under varying proportions of labeled utility data.

GPR B-scan images, which limits the cross-domain generalization capability of the learned model. The performance of different training strategies under varying labeled data ratios is shown in Fig. 2 . When trained from scratch, the model exhibits poor performance in low-data regimes. In contrast, the proposed method consistently achieves the highest mAP@0.5 across all training ratios. In the extreme low-data setting with 5% labeled samples, our approach improves mAP@0.5 from 47% obtained by baseline fine-tuning to 69%. Comparable performance requires 10% labeled data with fine-tuning or approximately 30% with training from scratch. These quantitative results demonstrate that the proposed method significantly improves detection performance in low-data regimes. Qualitative comparison results of different training strategies are shown in Fig. 3 . When trained from scratch, the model fails to correctly identify the hyperbolic target. With standard fine-tuning, although the true hyperbola is partially detected, large regions of the background are incorrectly recognized as hyperbolic structures, resulting in a considerable number of false positives. The proposed method effectively suppresses background-induced false detections and accurately localizes the target hyperbola by enforcing symmetry-based structural consistency during fine-tuning. Sensitivity analysis of the symmetry weight is shown in Fig. 4. Under training usage ratios of 0.05, 0.10, and 0.15, the detection performance varies with different values of λ. Overall, a weighting value of $\lambda = 2$ yields the best performance. In particular, under the most challenging setting with only 5% labeled training data, this value leads to a clear improvement in mAP@0.5, which is consistent with the intended role of the symmetry-based constraint in enhancing small-sample adaptation.

IV. DISCUSSION

The symmetry constraint can also be applied to training from scratch and yields consistent performance improvements across different data ratios, although the gains remain modest compared to fine-tuning. This indicates that while the

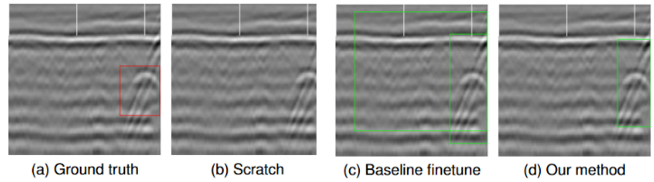

Fig. 3. Detection results under different training strategies.

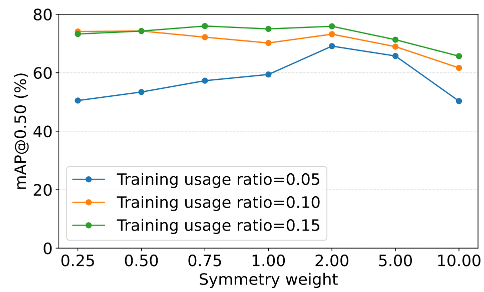

Fig. 4. Sensitivity analysis of the weighting parameter.

symmetry constraint provides useful structural regularization, its effectiveness is enhanced when combined with knowledge transfer from the source domain, suggesting that the two are complementary. However, the effectiveness of the proposed method depends on the presence of stable geometric structures in GPR reflections. In scenarios with highly irregular targets or complex environments (e.g., strong noise, multipath effects, or unstable media), the structural prior becomes less reliable, and large-scale data-driven learning is still required.

V. CONCLUSION

This paper proposes a structure-aware transfer learning approach for GPR-based subsurface object detection. By introducing a symmetry-based structural constraint during fine-tuning, the proposed method improves detection performance under limited training data conditions. Experimental results on real-world GPR datasets demonstrate consistent performance gains over conventional fine-tuning strategies, with more pronounced improvements observed in extremely low-data scenarios. Future work will explore extending the proposed structure-aware transfer learning framework to a broader range of subsurface targets and GPR acquisition settings.

ACKNOWLEDGMENTS

This research was funded under the Japanese Research Fund Kakenhi-C 23K04329.

REFERENCES

[1] C. Gomez, N. Hotta, S. Miyata, B. Bradak, M. Kataoka, K. Ashikaga, and F. C. Persendt, "Detection and velocimetry of floating wood for flood disaster risk management using electromagnetic imaging," *Proceedings*, vol. 87, no. 1, 2023.
[2] A. Mojahid, D. E. Ouai, K. E. Amraoui, K. El-Hami, and H. Aitbenamer, "Intelligent recognition of subsurface utilities and voids: A ground penetrating radar dataset for deep learning applications," *Data in Brief*, vol. 59, p. 111338, 2025.