remote sensing

MDPI

Article

Development of an Open GPR Dataset for Enhanced Bridge Deck Inspection

Da Hu

Department of Civil and Environmental Engineering, Kennesaw State University, Marietta, GA 30060, USA; dhu3@kennesaw.edu

Abstract

Bridge infrastructure in the United States is aging, necessitating efficient and accurate inspection methods. Ground-penetrating radar (GPR) is a widely used non-destructive testing (NDT) method for detecting subsurface anomalies in bridge decks. However, manual interpretation of GPR scans is labor-intensive, and annotated datasets for deep learning applications are limited. This study investigates YOLO-based deep learning models for automated rebar detection using a combination of real and synthetic GPR data. A dataset comprising 2255 real GPR images from four bridges and 20,000 simulated GPR scans was used to train and evaluate YOLOv8, YOLOv9, YOLOv10, and YOLOv11 under different training strategies. The results show that pretraining on simulated GPR data improves detection accuracy compared to conventional COCO pretraining, demonstrating the effectiveness of domain-specific transfer learning. These findings highlight the potential of simulated GPR data for training deep learning models, reducing reliance on extensive real-world annotations. This study contributes to AI-driven infrastructure monitoring, supporting the development of more scalable and automated GPR-based bridge inspections.

Keywords: Ground-Penetrating Radar (GPR); bridge deck inspection; deep learning; synthetic data; YOLO object detection

check for updates

Academic Editors: José Fernando Borges and Rui Jorge Oliveira

Received: 24 April 2025
Revised: 12 June 2025
Accepted: 23 June 2025
Published: 27 June 2025

Citation: Hu, D. Development of an Open GPR Dataset for Enhanced Bridge Deck Inspection. *Remote Sens.* **2025**, *17*, 2210. https://doi.org/10.3390/rs17132210

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

In the United States, the deterioration of bridge infrastructure has become a critical concern, affecting public safety and economic efficiency. According to the American Road & Transportation Builders Association (ARTBA), over 222,000 of the nation's 617,000 bridges—approximately 36%—require repairs or replacement [1]. Additionally, the Federal Highway Administration (FHWA) reports that 42% of all bridges are at least 50 years old, with 46,154 bridges (7.5%) classified as structurally deficient [2]. These alarming statistics emphasize the urgent need for advanced inspection techniques to ensure structural integrity and extend the service life of aging infrastructure. Timely and accurate bridge deck assessments are essential for prioritizing maintenance efforts, preventing catastrophic failures, and optimizing budget allocation for repairs and replacements.

Ground-penetrating radar (GPR) has emerged as a powerful non-destructive testing (NDT) technique for bridge deck inspections, enabling engineers to detect subsurface anomalies such as rebar, voids, and deterioration without compromising the structure [3]. GPR offers several advantages over traditional inspection methods, including rapid data acquisition, minimal disruption to traffic, and the ability to detect internal defects that may not be visible on the surface. However, despite these benefits, the effectiveness of GPR-based assessments is highly dependent on the accurate interpretation of radargrams, which

requires specialized knowledge and experience. Several traditional signal processing techniques, such as migration and microwave tomography, have been developed to enhance the interpretability of radargrams by improving target resolution and suppressing clutter. These methods aim to reconstruct clearer subsurface images and support manual interpretation. A comprehensive overview of such approaches is provided in Persico's seminal text on GPR data processing [4]. Manual analysis of GPR scans is not only time-consuming and labor-intensive but also prone to human subjectivity and variability, making it difficult to standardize results across large-scale infrastructure networks [5].

To address these challenges, there is growing interest in applying machine learning (ML) and deep learning (DL) techniques to automate and enhance the interpretation of GPR data [6,7]. Deep learning models, particularly object detection frameworks, have shown significant promise in identifying rebar, voids, and other subsurface features from GPR scans with high accuracy and efficiency. Among these models, YOLO (You Only Look Once) architectures have gained popularity due to their ability to perform real-time object detection with strong localization capabilities, making them suitable for processing large-scale GPR datasets efficiently [8]. However, the success of deep learning models in GPR applications is highly dependent on the availability of large, diverse, and well-annotated datasets for training and validation.

Publicly available GPR datasets for bridge inspections remain scarce, posing a significant challenge for the development and benchmarking of robust deep learning-based detection models [9]. Most existing datasets are limited in size, lack diversity in subsurface conditions, and often require extensive manual annotation, which is both costly and time-intensive. To overcome this limitation, this study introduces a comprehensive GPR dataset that combines both real-world and simulated GPR data to improve deep learning model training. The real dataset consists of 2255 GPR images collected from four different bridges, representing various antenna frequencies, bridge deck conditions, and rebar configurations. To supplement this, a large-scale simulated GPR dataset of 20,000 radargrams is generated using gprMax 3.1.7 [10], an open-source GPR simulation software. The simulated dataset is designed to incorporate variations in material properties, rebar spacing, antenna frequencies, and noise conditions, ensuring that the dataset captures a wide range of possible real-world scanning conditions. The integration of both real and synthetic data provides a scalable approach to training deep learning models, enhancing their generalization and robustness when deployed for real-world bridge deck inspections.

This study evaluates the performance of YOLO-based object detection models for automated rebar detection in GPR scans, comparing three training strategies: training from scratch on real GPR data, transfer learning using COCO-pretrained weights, and transfer learning using simulated GPR-pretrained weights. By systematically comparing these approaches, this study aims to assess the impact of domain-specific pretraining and determine whether synthetic GPR data can serve as a viable alternative to real-world datasets for deep learning applications. The results of this work contribute to the advancement of AI-driven infrastructure monitoring, providing insights into the scalability, accuracy, and efficiency of deep learning models for automated bridge deck inspections.

2. Literature Review

Ground-penetrating radar (GPR) technology, a non-invasive method, has become a cornerstone in civil engineering for inspecting subsurface structures. Over the years, ground-penetrating radar (GPR) has advanced into a powerful imaging technology that utilizes electromagnetic waves to detect and analyze subsurface structures [11,12]. GPR's prominence in civil engineering, particularly in bridge deck inspection, is noteworthy. It stands out among other non-destructive testing (NDT) methods due to its ability to

quickly assess large areas, providing detailed insights into the internal condition of bridge decks without causing any damage [13]. The accurate localization of rebar within concrete structures is a crucial aspect of bridge deck inspection. GPR's ability to map rebar layouts aids in evaluating the structural integrity and planning maintenance activities, though challenges remain in interpreting the complex signals associated with dense rebar networks.

2.1. Rebar Localization and Mapping

Evaluating the structural elements within bridge decks, such as rebars, pre-stressed tendons, and ducts, is a critical component of bridge inspection processes. Traditional surface inspection methods, like visual examination, often fall short of accurately detecting these embedded elements. This is where ground-penetrating radar (GPR) comes into play, demonstrating remarkable capability in identifying the condition of such metallic elements, thanks to the sensitivity of electromagnetic (EM) waves to metal [14]. GPR's effectiveness in mapping the intricate network of rebars is noteworthy. Utilizing array GPR antennas or advanced 3D antennas, detailed three-dimensional maps of rebar layouts can be generated, offering invaluable data for maintenance and rehabilitation planning [15]. One of the key aspects of analyzing GPR data for rebar detection involves identifying the hyperbolic patterns indicative of rebar presence [16]. Initially, GPR data interpretation relied heavily on manual analysis by skilled technicians. This involved visually inspecting GPR scan outputs to identify patterns and anomalies indicative of structural defects or features. Despite its effectiveness in certain scenarios, manual interpretation of GPR data is time-consuming and can be subjective. This approach becomes less reliable with the increasing complexity of the structures being examined, highlighting the need for more advanced interpretation methods.

To enhance the precision and efficiency in detecting rebars and assessing their condition, recent innovations have turned towards the integration of machine learning (ML) and deep learning (DL) models. Notably, an ML-based approach employing a Histogram of Oriented Gradients (HOG) feature descriptors combined with a supervised binary classifier, such as AdaBoost, has shown promise in automating the extraction of rebar locations from GPR data [17]. This method, verified through its application on publicly available GPR image datasets like DECKGPRHv1.0, has achieved high accuracy levels in detecting rebars, despite the complexity of field data. The algorithm's effectiveness in identifying rebar corrosion has also been acknowledged, although it is influenced by various training parameters like cell size and block overlap. In addition to ML models, deep learning techniques have made significant strides in this domain. A notable example is a two-step deep CNN model, which initially applies image processing techniques such as migration and normalized cross-correlation to pinpoint potential rebar locations, followed by classification using a trained CNN [18]. This method, along with other DL algorithms such as YOLO-based architectures [19,20], has demonstrated excellent results in identifying rebars in GPR scans with high accuracy.

2.2. Reinforcement Corrosion

The assessment of rebar corrosion within bridge decks is a critical aspect of structural health monitoring. Ground-penetrating radar (GPR) plays a pivotal role in this process. Chloride-induced corrosion, a common threat to bridge integrity, begins when environmental moisture and chloride ions penetrate the concrete. This infiltration leads to the formation of expansive rust products, which can increase in volume by four to six times, as identified in studies by Zaki et al. [21]. This expansion causes significant stress within the concrete, leading to micro-cracking, which can escalate to more severe damage like

macrocracking, debonding, delamination, and spalling. Such deterioration compromises the load-bearing capacity and overall structural integrity of the bridge.

Early detection and assessment of such corrosion, particularly in regions prone to high chloride ion concentration, are crucial for maintaining structural health. GPR methodologies have traditionally employed two main approaches for evaluating chloride-induced corrosion. The first utilizes empirical or theoretical models to correlate changes in the bulk dielectric constant with volumetric alterations pre- and post-corrosion [22]. This method is rooted in electromagnetic mixing theory, similar to principles applied in soil mechanics. The second approach involves full waveform inversion, which uses inverse GPR attributes, such as amplitude, velocity, and dielectric constant, to map chloride content and moisture distribution within the concrete structure. Recent advancements have seen the integration of machine learning (ML) and deep learning (DL) techniques in evaluating reinforcement corrosion using GPR. These methods focus on interpreting the blurred hyperbolic reflection amplitudes in GPR scans, which indicate corrosion near rebars. For instance, Zhang et al. [23] developed a sophisticated single-shot multi-box detector model that automates the detection of rebar corrosion. This model, trained on an extensive dataset of B-scans, has shown remarkable accuracy in identifying corrosion, rivaling traditional in-house software.

Moreover, ML and DL methods are increasingly being used to enhance the clarity of GPR data, replacing conventional signal processing techniques. For example, the use of denoising convolutional neural networks and automatic encoders has proven effective in reducing noise in GPR data [24]. Additionally, innovative DL-based methods have been designed to eliminate rebar-induced clutter and amplify weaker defect signals [25]. Such methods employ sophisticated neural network architectures, incorporating elements like residual-inception blocks and attention modules to improve the accuracy of defect detection under complex conditions. These developments mark a significant leap in the capability of GPR technology to identify and assess rebar corrosion accurately, thus aiding in the timely maintenance and rehabilitation of bridge structures.

2.3. Challenges and Limitations in Current Practices

The efficacy of machine learning (ML) and deep learning (DL) models in interpreting GPR data heavily relies on the availability of comprehensive datasets that are both extensive and accurately labeled. The creation of such datasets presents significant challenges within the field. Not only is it a time-consuming process, but it also demands a high degree of expertise, particularly in accurately distinguishing and labeling the diverse array of features and anomalies present in GPR scans. Compounding this challenge is the labor-intensive nature of manual labeling. This process is not only tedious but also prone to inconsistencies and human error. Additionally, the variability in GPR data formats, resulting from the use of equipment from different manufacturers, further complicates the development of universally compatible and high-quality datasets. These issues collectively create significant barriers to the creation of standardized datasets that are crucial for the effective training and validation of ML/DL models.

Another major impediment in the field is the conspicuous absence of open-source GPR datasets, particularly those that are tailored for ML/DL applications in bridge deck inspection. This gap severely restricts the potential for collaborative research and hampers the development and validation of sophisticated data interpretation models. The lack of such datasets not only slows down the pace of technological advancements but also limits the scope of research in this critical area. The proposition of an open-source GPR dataset designed specifically for ML/DL applications could serve as a significant resource in this domain. It would empower researchers to develop, test, and refine data interpretation models more effectively, thereby enhancing bridge deck inspections' accuracy and

efficiency. The introduction of such a dataset has the potential to revolutionize current inspection practices, leading to more precise and expedited assessments. Moreover, it would encourage collaborative research efforts, fostering broader innovation and progress in the field.

3. Methodology

3.1. Simulated GPR Data

To create a diverse ground-penetrating radar (GPR) dataset for bridge deck analysis, this study produced a substantial number of synthetic GPR radargrams using gprMax 3.1.7 [10], an open-source GPR simulation tool. gprMax has gained recognition for its capability to replicate accurate GPR data for various subsurface investigations. In an effort to mimic the conditions of actual concrete bridge decks as closely as possible, the research incorporates several bridge deck design standards into the simulation process.

Asphalt overlays are often applied to bridge decks as a preservation measure, shielding the concrete below from further wear and environmental elements, thereby prolonging the structure's service life. The use of overlays is particularly common in refurbishing aging concrete bridges where the existing deck may have deteriorated. To capture the range of possible scenarios in service, the simulation includes bridge decks with and without asphalt overlay in a randomized fashion, ensuring that the generated dataset encapsulates the variability encountered in real-world bridge deck inspections. Table 1 delineates the parameters used in the simulation of bridge decks.

Table 1. Simulated bridge deck parameters.

Bridge Type	Thickness (cm)	Rebar Layers	Overlay Presence
Small pedestrian	127–152.4	1	Randomized
Highway	152.4–228.6	2	Randomized
Heavy-duty	228.6–304.8	2	Randomized

Note: 'Overlay Presence' in Table 1 specifies whether an asphalt overlay is included in the simulation for each bridge type. This is randomized to reflect the actual variations found in bridge decks, where some may have overlays and others may not.

In our efforts to enhance the ground-penetrating radar (GPR) dataset, we have meticulously outlined simulation parameters to capture the varied structural elements and materials characteristic of bridge deck construction. Table 2 displays our simulation matrix, which provides a range of values for critical layers and rebar dimensions crucial for authentic GPR modeling. The top concrete cover, ranging from 50 to 75 mm thick [26], mimics the uppermost protective layer of the bridge deck. This layer shields reinforcing steel bars from environmental elements and delays the onset of corrosive processes. Similarly, the bottom concrete cover varies between 25 and 50 mm to reflect the underlying support layer of the rebar, contributing to the bridge deck's overall impact resistance and structural rigidity. We also included an asphalt overlay, with a thickness ranging from 12.7 to 76.2 mm [27], which mirrors the refurbishment techniques applied to extend the functional lifespan of bridge structures through the application of a protective surface layer.

Table 2. Bridge deck component dimensions.

Component	Thickness/Dimension (mm)
Top concrete cover	50–75
Bottom concrete cover	25–50
Asphalt overlay	12.7–76.2
Rebar diameter	12–26
Rebar spacing	127–330

In order to accurately represent the diverse load-bearing requirements of bridges, the rebar diameter ranges from 12 to 26 mm [28]. Additionally, the spacing between the rebar varies from 127 to 330 mm [29], a key factor in assessing structural integrity and analyzing GPR signal dispersion. These parameters are randomized within the specified ranges, ensuring that the resulting simulations closely resemble real-world bridge decks. By incorporating this variability, the GPR radargrams produced can serve as a solid foundation for developing advanced models for interpreting GPR data. This will make a significant contribution to the progress of non-destructive evaluation techniques for civil infrastructure assessment. Figure 1 presents examples of simulated bridge deck models for different bridge types.

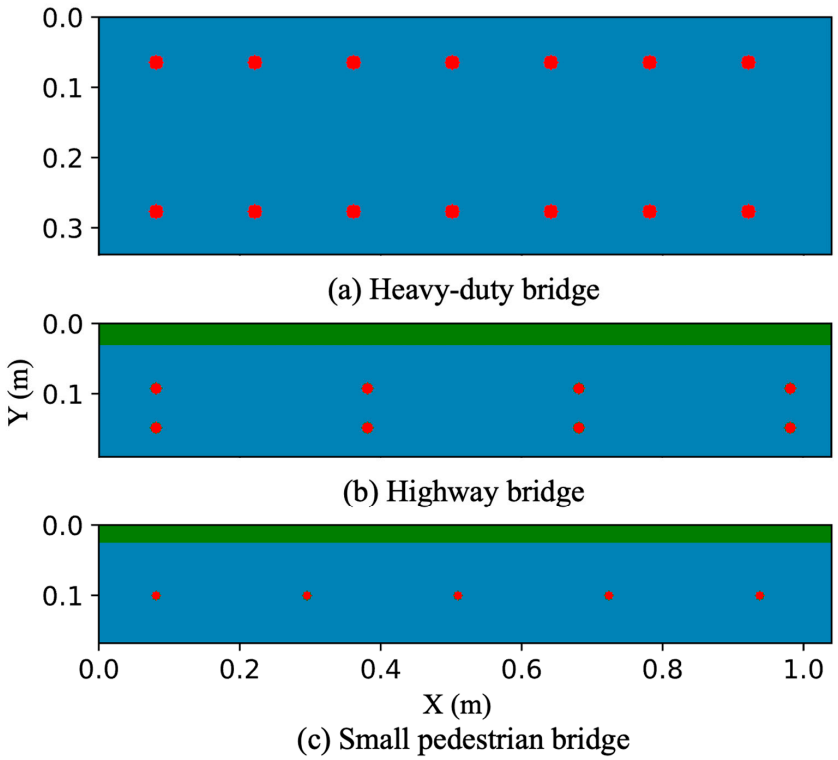


Figure 1. Examples of simulated bridge deck models for different bridge types.

Table 3 outlines the dielectric constant and conductivity for various materials commonly encountered in bridge deck construction and their respective properties for GPR simulation purposes. The dielectric constant and conductivity parameters for concrete and asphalt are presented with a range of values, indicating that these properties are randomized within the specified limits to more accurately reflect the variability of real-world conditions. Concrete exhibits a dielectric constant ranging from 6 to 12 and a conductivity between 0.0001 and 0.01 mS/m [30]. Asphalt, a material with different electromagnetic properties, has a dielectric constant that varies from 4 to 10 and a similar conductivity range as concrete [31]. Rebar, a critical component in reinforced concrete structures, is assigned a dielectric constant of 1 and a significantly higher conductivity of 1,000,000 mS/m, characteristic of its highly conductive metallic nature.

Table 3. Electromagnetic properties of bridge deck materials.

Material	Dielectric Constant	Conductivity (mS/m)
Concrete	6–12	0.0001–0.01
Asphalt	4–10	0.0001–0.01
Rebar	1	1,000,000

3.2. GPR Simulation Parameters

In the detailed simulation environment essential for our GPR dataset creation, we have carefully selected the operational frequency range to accommodate the nuanced resolution and penetration depth required for comprehensive bridge deck analysis. The chosen frequency spectrum—1.2 GHz, 1.6 GHz, 2 GHz, and 2.6 GHz—ensures a balanced trade-off between resolution and penetration. Lower frequencies, such as 1.2 GHz, provide deeper penetration, while higher frequencies, like 2.6 GHz, offer finer resolution and are crucial for detecting smaller anomalies and near-surface defects. The Ricker wavelet, a widely utilized waveform in GPR simulations, is employed due to its symmetric properties and single dominant frequency component, making it an ideal choice for resolving distinct boundaries within the bridge deck, such as interfaces between different materials or layers. This waveform allows for a clearer interpretation of the reflected signals, aiding in the accurate identification of structural features and potential areas of concern within the scanned materials. Spatial resolution during the simulation is set at 0.002 m, a meticulous choice that strikes a balance between computational efficiency and the level of detail that can be resolved. Figure 2 shows simulated GPR scans with different frequencies.

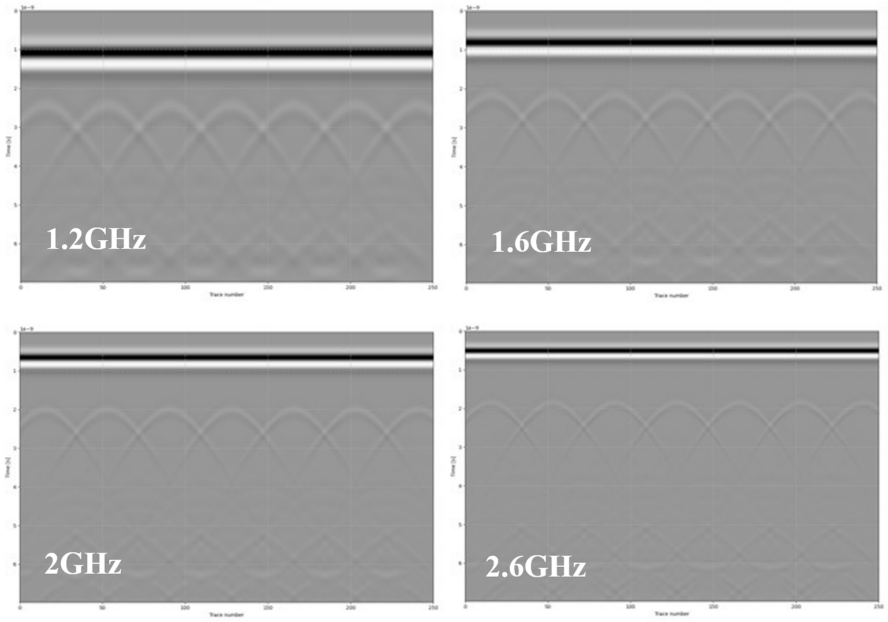


Figure 2. Simulated GPR scans with different frequencies.

3.3. GPR Data Augmentation

In total, we simulated 20,000 GPR scans with unique bridge deck specifications, with 5000 scans generated for each antenna frequency: 1.2 GHz, 1.6 GHz, 2 GHz, and 2.6 GHz. To enhance the dataset's diversity and ensure the model's robustness against real-world variations, three key augmentation techniques were applied, as illustrated in Figure 3: direct wave suppression, random noise addition, and Gaussian blur. These techniques aim to replicate the challenges encountered in real-world GPR data acquisition and to improve the model's ability to interpret subsurface information accurately.

The first technique, direct wave suppression, was used to remove the strong, early-arriving direct wave signals that typically dominate raw GPR data. These direct waves, caused by surface reflections and antenna coupling, often overshadow the weaker reflections originating from subsurface objects or anomalies, such as rebar corrosion or voids in the concrete. Suppressing these waves enhances the visibility of meaningful subsurface reflections, creating data that better represents processed field scans where similar signal processing techniques have been applied. The second augmentation method involved the

addition of random noise to the data. In real-world applications, GPR signals are frequently contaminated by environmental noise, equipment-related interference, or electromagnetic disturbances, all of which can degrade data quality. Introducing random noise into the simulated dataset ensures the model can learn to differentiate relevant subsurface features from noise, making it more resilient and effective in handling noisy field data collected in challenging inspection environments.

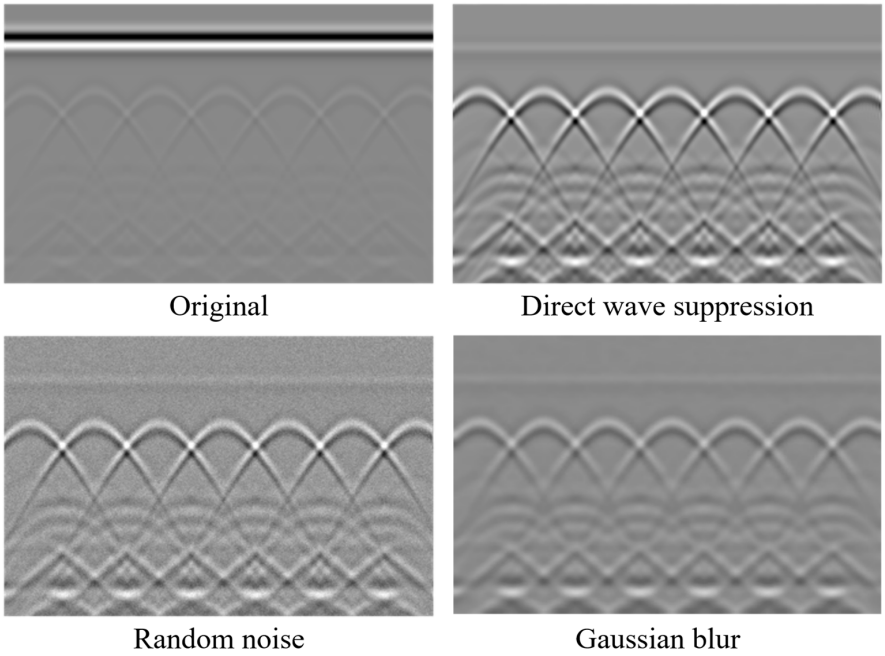


Figure 3. Simulated GPR data augmentation achieved through direct wave suppression, addition of random noise, and application of Gaussian blur.

Finally, Gaussian blur was applied to the data to mimic the effects of signal dispersion and attenuation as electromagnetic waves propagate through heterogeneous concrete bridge decks. This augmentation technique smooths abrupt transitions in the data, simulating the natural interactions of GPR waves with varying material properties such as concrete density, embedded aggregates, and moisture content. By simulating these real-world effects, Gaussian blur helps create a dataset that reflects the complexities of field data more accurately. By combining these augmentation techniques with the simulated GPR scans of concrete bridge decks, this study establishes a diverse and realistic dataset, enabling the development of robust models capable of supporting accurate and reliable bridge inspection processes under diverse operational conditions.

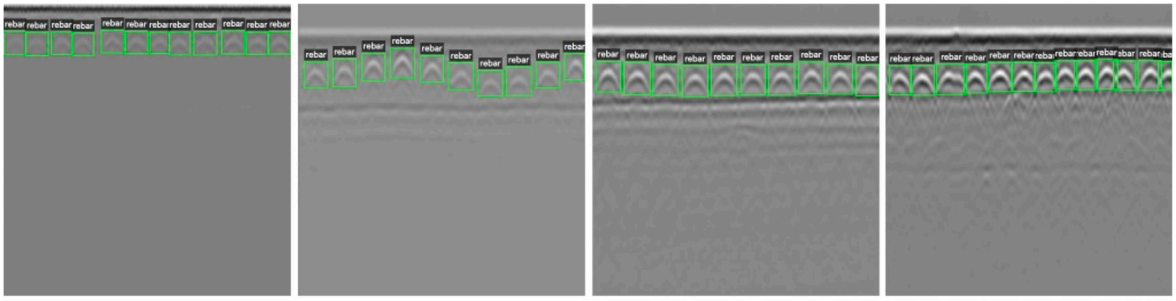
3.4. GPR Data Annotation

The annotation of GPR data is a critical step in preparing the dataset for model training and evaluation, particularly for the task of identifying subsurface features such as rebar. In this study, the annotation focused on labeling bounding boxes around hyperbolic features in the GPR scans, which correspond to the signatures of rebar within the concrete bridge deck. Accurate identification and annotation of these hyperbolic features are essential, as they directly correlate to the presence and positioning of rebar, a key aspect of bridge inspection.

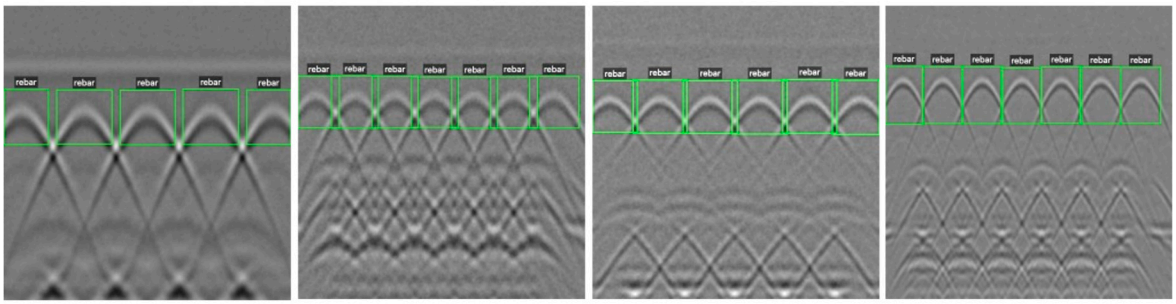
To perform the annotation, we used Labelme (version 5.5), a widely adopted graphical image annotation tool that supports efficient and accurate labeling. Each hyperbolic feature in the GPR data was enclosed in a bounding box, ensuring that the annotations captured the full extent of the hyperbola while excluding unnecessary background noise. This process required careful attention to detail, as hyperbolic patterns can vary in size

and shape depending on factors such as antenna frequency, rebar depth, and subsurface material properties. To ensure the quality and consistency of the annotations, the labeled data underwent a rigorous review process. After the initial annotation was completed, a second team member, trained in identifying hyperbolic features in GPR data, reviewed the annotations for accuracy and completeness. This cross-checking process helped minimize errors and ensured that the bounding boxes correctly captured the hyperbolic features of interest. Furthermore, annotations were standardized across the entire dataset to maintain uniformity. This involved defining clear guidelines for bounding box placement, such as ensuring that the bounding box tightly enclosed the hyperbolic feature and avoided excessive overlap with adjacent features or noise artifacts. This standardized approach ensured that the annotations provided high-quality input for training machine learning models to recognize rebar signatures in GPR scans.

A comprehensive dataset of GPR scans was annotated for this study, including both simulated and real-world data. For the simulated GPR scans, a total of 20,000 scans were annotated, focusing on bounding boxes around hyperbolic features corresponding to rebar signatures. Additionally, 2255 real GPR images were annotated, representing scans from four different bridges. Among these, two pedestrian bridges in Knoxville, Tennessee, were scanned using a 2 GHz antenna with the GSSI SIR 4000 system (Geophysical Survey Systems Inc., Nashua, NH, USA), while a local roadway bridge in Warren County, New Jersey, was scanned using a 1.6 GHz antenna with the GSSI SIR-20 system (Geophysical Survey Systems Inc., Nashua, NH, USA) [32]. The final dataset was collected from a highway bridge on U.S. Route 15 in Haymarket, Virginia, using a 2 GHz antenna with the IDS RIS Hi-Bright GPR system (IDS GeoRadar, Pisa, Italy) [32]. Figure 4 presents some representative examples of annotated and simulated GPR scans.



(a) Real GPR scans



(b) Simulated GPR scans

Figure 4. Examples of annotated real and simulated GPR scans.

4. Experiment and Results

4.1. Algorithm Selection and Configuration

In our experiments, we evaluated the performance of several YOLO (You Only Look Once) object detection models—specifically YOLOv8, YOLOv9, YOLOv10, and YOLOv11—along with their small, medium, and large backbones. We selected these models for their

advancements in real-time object detection, which is critical for automated, on-site GPR-based bridge inspections. Real-time detection methods enable faster scanning and analysis, thereby reducing both time and resources during field work.

Each YOLO version introduces architectural and algorithmic improvements to balance accuracy, inference speed, and computational complexity. These enhancements typically involve more efficient backbone networks, refined feature aggregation strategies (such as PANet, FPN, and CSP connections), and either anchor-based or anchor-free detection mechanisms. YOLOv8 offers a lightweight backbone and optimized feature pyramid refinements; YOLOv9 includes improved feature fusion layers and better loss functions; YOLOv10 leverages next-generation backbone designs and advanced data augmentation; and YOLOv11 further streamlines the detection pipeline, often achieving higher precision at greater computational cost.

We tested each model with multiple variant sizes to account for different resource constraints and to identify the best fit for our application. Smaller variants are ideal for real-time, edge-device scenarios where GPU memory and processing power are limited, whereas larger variants are more suitable for powerful workstations capable of handling heavier models. Hyperbolic signatures in GPR data can be subtle and vary significantly in size, so balancing accuracy and speed is essential for detecting critical structural features. This multi-model, multi-variant approach provides a comprehensive assessment of each YOLO architecture's strengths and weaknesses on a diverse set of simulated and real GPR scans. Our findings can serve as a benchmark for future researchers seeking to optimize rebar detection models for ground-penetrating radar applications, helping guide the selection of an optimal trade-off between detection precision, inference speed, and resource usage.

4.2. Implementation and Metrics

The present study was conducted using a workstation running Ubuntu 20.04, equipped with an AMD® Ryzen Threadripper Pro 5955WX CPU, 128 GB RAM, and dual NVIDIA RTX A6000 GPUs for model training and testing. The deep learning network was trained using the stochastic gradient descent (SGD) optimizer, with carefully tuned hyperparameters to enhance performance. The initial learning rate was set to 0.01, and the final learning rate was calculated as a fraction (0.01) of the initial value. The batch size was set to 32, and the number of epochs was 50. The optimizer's momentum was configured at 0.937, serving as the momentum parameter for SGD, while a weight decay of 0.0005 was applied to regularize the model and minimize overfitting. To ensure a smooth start to training, a warmup phase spanning 3 epochs was employed, during which the learning rate and momentum increased gradually. The warmup momentum began at 0.8, and the initial bias learning rate was set at 0.1.

Data augmentation techniques were employed to enhance the diversity and robustness of the training dataset, thereby improving the model's generalization capabilities. A Mosaic augmentation was applied with a probability of 100%, combining four images into a single composite image during training. This approach introduced significant variability in image composition, allowing the model to learn more generalized features. Additionally, geometric augmentations were applied, including translation and scaling. The translation factor of 10% enabled the images to shift horizontally or vertically by up to 10% of their dimensions, while the scaling factor of 50% allowed objects to be resized by up to half their original size. These transformations helped the model become invariant to minor positional and size variations, further strengthening its robustness. Furthermore, color augmentations were performed by adjusting the HSV (hue, saturation, and value) color space. The hue was modified by 1.5%, the saturation by 70%, and the value (brightness) by 40%. These

adjustments simulated variations in color intensity, making the model more resilient to real-world scenarios with diverse visual appearances.

We evaluated the performance of our YOLO-based rebar detection models using the mean Average Precision (mAP50-95) metric, a widely recognized benchmark in object detection. This metric assesses both detection accuracy and localization quality, providing a comprehensive measure of how well the models identify and precisely locate hyperbolic signatures corresponding to rebar in GPR scans. The mAP50-95 metric is calculated by averaging the Average Precision (AP) scores across multiple Intersection over Union (IoU) thresholds, ranging from 0.50 to 0.95 in increments of 0.05. This ensures that the evaluation captures model performance across varying degrees of localization strictness. At lower IoU thresholds, detections with some misalignment still contribute to the score, reflecting general detection capability. At higher IoU thresholds, only highly accurate bounding boxes are counted, emphasizing precise localization. By averaging over this range, mAP50-95 provides a balanced assessment of both detection robustness and localization accuracy, making it a reliable metric for evaluating rebar detection in GPR scans.

4.3. Results

To assess the effectiveness of various YOLO models for hyperbolic feature detection in GPR scans, we conducted experiments comparing model performance when trained from scratch and when using transfer learning. The dataset was split into 70% training (1578 images) and 15% validation (338 images) for the real GPR dataset. For transfer learning, we first trained models on a simulated GPR dataset, using 14,000 images for training and 3000 images for validation. The best-performing model on the validation set was then used as pretrained weights for fine-tuning on the real GPR dataset. The results of our experiments are summarized in Table 4. The mean Average Precision (mAP50-95) values are reported for both training scenarios, along with the observed improvement due to transfer learning.

Table 4. Model performance comparison (training from scratch vs. transfer learning).

Model	No. of Params (Millions)	Training from Scratch (mAP50-95)	Transfer Learning (mAP50-95)	Improvement
YOLOv8n	3.2	0.886	0.941	0.055
YOLOv8s	11.2	0.899	0.964	0.065
YOLOv8m	25.9	0.907	0.971	0.064
YOLOv8l	43.7	0.911	0.977	0.066
YOLOv8x	68.2	0.917	0.977	0.06
YOLOv9t	2.0	0.880	0.944	0.064
YOLOv9s	7.1	0.892	0.963	0.071
YOLOv9m	20.0	0.902	0.972	0.07
YOLOv9c	25.3	0.905	0.974	0.069
YOLOv9e	57.3	0.873	0.976	0.103
YOLOv10n	2.3	0.869	0.950	0.081
YOLOv10s	7.2	0.878	0.968	0.09
YOLOv10m	15.4	0.881	0.975	0.094
YOLOv10l	29.5	0.869	0.977	0.108
YOLOv10x	56.9	0.024	0.978	0.954
YOLOv11n	2.6	0.880	0.943	0.063
YOLOv11s	9.4	0.887	0.966	0.079
YOLOv11m	20.1	0.898	0.975	0.077
YOLOv11l	25.3	0.886	0.976	0.09
YOLOv11x	56.9	0.892	0.978	0.086

The results demonstrate that transfer learning significantly improves performance across all model variants. For most models, the mAP50-95 score increased by approximately 6% to 10%, highlighting the effectiveness of pretraining on simulated GPR data before fine-tuning on real GPR scans. Notably, YOLOv10x showed an exceptionally large improvement (+0.954 mAP50-95), likely due to instability in training from scratch, where the model struggled to converge. Smaller models such as YOLOv8n, YOLOv9t, and YOLOv10n, while computationally efficient, benefited from transfer learning but still achieved slightly lower overall accuracy compared to larger variants. Larger models such as YOLOv8x, YOLOv9e, YOLOv10x, and YOLOv11x exhibited the highest accuracy after transfer learning, demonstrating their ability to leverage complex features extracted from the pretraining stage.

To evaluate the robustness of the models under data-constrained scenarios, we conducted an experiment using only 339 images for training, while keeping the same validation set (338 images). The objective was to analyze how reducing the training dataset size affects model performance, and whether transfer learning could mitigate the performance drop typically observed when training with limited data. The results in Table 5 below highlight the model performance in terms of mAP50-95, comparing training from scratch and training with pretrained weights obtained from the simulated GPR dataset.

Table 5. Model performance with limited training data (training from scratch vs. transfer learning).

Model	No. of Params (Millions)	Training from Scratch (mAP50-95)	Transfer Learning (mAP50-95)	Improvement
YOLOv8n	3.2	0.763	0.898	0.135
YOLOv8s	11.2	0.787	0.920	0.133
YOLOv8m	25.9	0.775	0.938	0.163
YOLOv8l	43.7	0.769	0.951	0.182
YOLOv8x	68.2	0.789	0.952	0.163
YOLOv9t	2.0	0.109	0.902	0.793
YOLOv9s	7.1	0.104	0.920	0.816
YOLOv9m	20.0	0.653	0.945	0.292
YOLOv9c	25.3	0.539	0.940	0.401
YOLOv9e	57.3	0.000	0.952	0.952
YOLOv10n	2.3	0.565	0.908	0.343
YOLOv10s	7.2	0.331	0.934	0.603
YOLOv10m	15.4	0.069	0.946	0.877
YOLOv10l	29.5	0.001	0.954	0.953
YOLOv10x	56.9	0.000	0.952	0.952
YOLOv11n	2.6	0.714	0.900	0.186
YOLOv11s	9.4	0.654	0.928	0.274
YOLOv11m	20.1	0.695	0.946	0.251
YOLOv11l	25.3	0.065	0.947	0.882
YOLOv11x	56.9	0.032	0.951	0.919

The results highlight the critical role of simulated GPR data in improving model performance, particularly in data-limited scenarios. When training from scratch with only 339 real GPR images, many models struggled to learn meaningful patterns, with some—particularly larger models—failing to converge at all. This is a common challenge in deep learning applications where real-world data is scarce, as models require a sufficient volume of labeled examples to effectively capture complex features, such as hyperbolic patterns indicative of rebar in bridge decks. Without enough training samples, models trained from scratch exhibited low accuracy and poor generalization, making them unreliable for real-world deployment.

However, when transfer learning was applied using pretrained weights from the simulated GPR dataset, the models exhibited substantial performance improvements, with gains ranging from 13.5% to over 95% in mAP50-95. This demonstrates that the simulated dataset successfully captured essential GPR signal characteristics, allowing models to learn fundamental feature representations before fine-tuning on real data. As a result, models pretrained on simulated scans were able to generalize significantly better, despite being trained on a small real-world dataset. Notably, larger models that initially failed to converge (e.g., YOLOv10l, YOLOv10x, and YOLOv11x) achieved near-perfect accuracy when initialized with simulated GPR-trained weights, proving that pretraining on synthetic data can compensate for the lack of extensive real-world annotations.

One of the key advantages of using simulated GPR data for pretraining is that it enables models to learn robust feature representations without requiring extensive manual labeling efforts. Generating high-quality simulated GPR scans with controlled variations in material properties, rebar spacing, and antenna frequencies provides a diverse dataset that mirrors real-world conditions. By training models on this rich synthetic dataset first, we can significantly reduce the dependency on large real-world datasets, which are often costly and labor-intensive to collect. This is particularly valuable for bridge inspection applications, where acquiring high-quality GPR scans with labeled rebar locations can be both time-consuming and impractical.

Additionally, the simulated dataset allows for systematic control over data characteristics, ensuring that the model is exposed to a wide range of potential bridge deck conditions, including different subsurface materials, noise levels, and rebar configurations. This structured variability improves model generalization, enabling it to adapt more effectively to unseen real-world conditions. The success of transfer learning in this study underscores how simulated GPR data serves as a scalable and cost-effective alternative to manually collected real-world datasets, offering a practical solution for training high-performance deep learning models for GPR-based infrastructure inspection.

5. Discussion

5.1. Comparison of Transfer Learning from COCO vs. Simulated GPR Dataset

To assess the effectiveness of domain-specific pretraining, we compared the performance of models pretrained on the COCO dataset versus those pretrained on simulated GPR data. The results, summarized in Table 6, show marginal differences in performance, with only minor variations in mAP50-95 scores between the two pretraining approaches. While some models demonstrated slight improvements when using simulated GPR weights (e.g., YOLOv10n: +0.014, YOLOv11n: +0.009), others showed negligible differences or even slight declines (YOLOv9c, YOLOv9e, YOLOv10m, YOLOv10l: −0.001). These variations suggest that the impact of pretraining on simulated GPR data, while slightly favorable, is not significantly distinct from COCO-based pretraining for the given dataset size.

The relatively small differences in mAP50-95 indicate that both COCO and simulated GPR pretraining provide a good initialization for model weights, helping models converge more effectively than training from scratch. However, COCO pretraining did not show clear advantages despite being widely used for general-purpose object detection. This is likely because COCO-based pretraining focuses on learning edge-based object features from natural images, whereas GPR data consists of distinct hyperbolic signatures that are fundamentally different from conventional visual patterns.

Although the improvements from simulated GPR pretraining are small, the results suggest that domain-specific pretraining does not negatively impact performance and may still be beneficial in scenarios where real-world data is extremely limited. The fact that models pretrained on GPR data achieved comparable results to COCO-pretrained

models indicates that synthetic GPR scans can serve as a viable alternative for initializing deep learning models, particularly when large-scale labeled GPR datasets are unavailable. Further refinement of the simulated dataset, including increased variability in subsurface conditions, rebar configurations, and noise characteristics, may enhance its effectiveness in transfer learning applications.

Table 6. Comparison of transfer learning performance using COCO vs. simulated GPR pretrained weights.

Model	No. of Params (Millions)	Transfer Learning (COCO Dataset)	Transfer Learning (Simulated GPR Dataset)	Improvement
YOLOv8n	3.2	0.934	0.941	0.007
YOLOv8s	11.2	0.960	0.964	0.004
YOLOv8m	25.9	0.973	0.971	−0.002
YOLOv8l	43.7	0.977	0.977	0
YOLOv8x	68.2	0.977	0.977	0
YOLOv9t	2.0	0.937	0.944	0.007
YOLOv9s	7.1	0.963	0.963	0
YOLOv9m	20.0	0.972	0.972	0
YOLOv9c	25.3	0.975	0.974	−0.001
YOLOv9e	57.3	0.977	0.976	−0.001
YOLOv10n	2.3	0.936	0.950	0.014
YOLOv10s	7.2	0.968	0.968	0
YOLOv10m	15.4	0.976	0.975	−0.001
YOLOv10l	29.5	0.978	0.977	−0.001
YOLOv10x	56.9	0.979	0.978	−0.001
YOLOv11n	2.6	0.934	0.943	0.009
YOLOv11s	9.4	0.966	0.966	0
YOLOv11m	20.1	0.976	0.975	−0.001
YOLOv11l	25.3	0.976	0.976	0
YOLOv11x	56.9	0.977	0.978	0.001

To further evaluate the impact of domain-specific pretraining, we compared models pretrained on COCO versus simulated GPR data, using a smaller training dataset (339 images) and the same validation set. The goal was to determine whether GPR-specific pretraining offers an advantage in data-limited scenarios. As shown in Table 7, models pretrained on simulated GPR data consistently outperformed those pretrained on COCO, with improvements ranging from +0.003 to +0.029 mAP50-95. The largest gains were observed in smaller models, such as YOLOv10n (+0.029) and YOLOv9e (+0.020), suggesting that lightweight models benefit more from domain-specific feature extraction when data is scarce. Larger models, like YOLOv10x (−0.002), showed negligible differences, indicating that their capacity allows them to compensate for the lack of GPR-specific pretraining. While the improvements are small, the results reinforce the value of synthetic GPR data for transfer learning. Unlike COCO, which contains natural images, simulated GPR data better captures the unique hyperbolic patterns of subsurface objects, enhancing model adaptation to real-world GPR scans. This suggests that pretraining on synthetic GPR scans can reduce reliance on large, annotated datasets, making deep learning more practical for infrastructure inspection.

Table 7. Transfer learning performance comparison using COCO vs. simulated GPR pretrained weights (limited training data).

Model	No. of Params (Millions)	Transfer Learning (COCO Dataset)	Transfer Learning (Simulated GPR Dataset)	Improvement
YOLOv8n	3.2	0.882	0.898	0.016
YOLOv8s	11.2	0.915	0.920	0.005
YOLOv8m	25.9	0.923	0.938	0.015
YOLOv8l	43.7	0.932	0.951	0.019
YOLOv8x	68.2	0.939	0.952	0.013
YOLOv9t	2.0	0.893	0.902	0.009
YOLOv9s	7.1	0.916	0.920	0.004
YOLOv9m	20.0	0.930	0.945	0.015
YOLOv9c	25.3	0.933	0.940	0.007
YOLOv9e	57.3	0.932	0.952	0.02
YOLOv10n	2.3	0.879	0.908	0.029
YOLOv10s	7.2	0.926	0.934	0.008
YOLOv10m	15.4	0.943	0.946	0.003
YOLOv10l	29.5	0.947	0.954	0.007
YOLOv10x	56.9	0.954	0.952	−0.002
YOLOv11n	2.6	0.887	0.900	0.013
YOLOv11s	9.4	0.914	0.928	0.014
YOLOv11m	20.1	0.926	0.946	0.02
YOLOv11l	25.3	0.931	0.947	0.016
YOLOv11x	56.9	0.930	0.951	0.021

5.2. Synthetic Training and Real-World Validation

To assess the generalization capability of models trained on synthetic GPR data, we evaluated their performance on a real-world validation set consisting of 338 GPR images. Table 8 shows moderate transferability of the models, with some adapting better than others. YOLOv11l achieved the highest performance (0.535), followed by YOLOv11n (0.482) and YOLOv8m (0.463), suggesting that mid-to-large models can better leverage synthetic training for real-world detection.

Conversely, smaller models like YOLOv9t (0.296) and YOLOv9s (0.267) underperformed, likely due to limited capacity to learn complex radargram features. Interestingly, some larger models (e.g., YOLOv8l and YOLOv10x) showed lower-than-expected performance, indicating possible overfitting to synthetic data features that do not fully align with real-world conditions. These results highlight both the potential and limitations of using synthetic data alone. While it provides a scalable alternative for training, a performance gap remains when transitioning to real data. Future improvements could include fine-tuning on real data, domain adaptation strategies, or hybrid training approaches to enhance generalization.

Table 8. Performance of YOLO-based models trained on synthetic GPR data and validated on real GPR data.

Model	No. of Params (Millions)	mAP50-95
YOLOv8n	3.2	0.408
YOLOv8s	11.2	0.362
YOLOv8m	25.9	0.463
YOLOv8l	43.7	0.322
YOLOv8x	68.2	0.454
YOLOv9t	2.0	0.296
YOLOv9s	7.1	0.267
YOLOv9m	20.0	0.275
YOLOv9c	25.3	0.384
YOLOv9e	57.3	0.353
YOLOv10n	2.3	0.448
YOLOv10s	7.2	0.349
YOLOv10m	15.4	0.386
YOLOv10l	29.5	0.348
YOLOv10x	56.9	0.340
YOLOv11n	2.6	0.482
YOLOv11s	9.4	0.353
YOLOv11m	20.1	0.446
YOLOv11l	25.3	0.535
YOLOv11x	56.9	0.356

5.3. Generalization to Unseen Bridges

To evaluate the generalization capability of the trained models, we conducted a cross-bridge validation experiment. Specifically, one bridge dataset containing 209 labeled GPR radargrams was set aside as an unseen test set, while the remaining three bridge datasets comprising 2046 radargrams were used for training. This setup allows us to assess how well models trained with real-world data, with and without transfer learning, perform when applied to structurally and environmentally different bridge conditions. The results, summarized in Table 9, demonstrate that incorporating transfer learning with simulated GPR data notably enhances model generalization. Across all tested YOLO variants, models pretrained on the simulated dataset consistently outperformed those trained from scratch. The improvement in mAP50-95 ranged from 0.045 (YOLOv8x) to as high as 0.745 (YOLOv10x), indicating the robustness and adaptability introduced by pretraining on a large and diverse set of synthetic GPR radargrams. Notably, compact models such as YOLOv9t and YOLOv10n exhibited substantial gains of 0.129 and 0.13, respectively, which is particularly encouraging for deployment scenarios with limited computational resources.

These findings reinforce the effectiveness of domain-specific synthetic pretraining, especially in settings where real-world data are limited or exhibit structural variability. By learning generic features from simulated scenarios—designed to capture various rebar configurations, material properties, and noise patterns—the models were better equipped to recognize similar patterns in new, unseen bridge data. This highlights the potential of using high-fidelity synthetic datasets not just as a supplement, but as a strategic foundation to build more generalizable deep learning models for GPR-based bridge inspections.

Table 9. Model performance on an unseen bridge (training from scratch vs. transfer learning with simulated GPR data).

Model	No. of Params (Millions)	Training from Scratch (mAP50-95)	Transfer Learning (mAP50-95)	Improvement
YOLOv8n	3.2	0.816	0.881	0.065
YOLOv8s	11.2	0.821	0.891	0.071
YOLOv8m	25.9	0.833	0.900	0.067
YOLOv8l	43.7	0.823	0.889	0.065
YOLOv8x	68.2	0.846	0.891	0.045
YOLOv9t	2.0	0.759	0.888	0.129
YOLOv9s	7.1	0.791	0.874	0.082
YOLOv9m	20.0	0.817	0.911	0.094
YOLOv9c	25.3	0.822	0.912	0.090
YOLOv9e	57.3	0.779	0.895	0.116
YOLOv10n	2.3	0.766	0.896	0.130
YOLOv10s	7.2	0.785	0.908	0.124
YOLOv10m	15.4	0.821	0.904	0.083
YOLOv10l	29.5	0.775	0.911	0.135
YOLOv10x	56.9	0.163	0.909	0.745
YOLOv11n	2.6	0.783	0.890	0.106
YOLOv11s	9.4	0.782	0.901	0.119
YOLOv11m	20.1	0.806	0.905	0.099
YOLOv11l	25.3	0.775	0.905	0.130
YOLOv11x	56.9	0.803	0.909	0.106

5.4. Limitation and Future Studies

Despite the promising results demonstrated in this study, several limitations remain that merit further exploration. First, this research exclusively focused on YOLO-based models for rebar detection. Although YOLO architectures were chosen due to their strong balance of accuracy and computational efficiency, particularly important in resource-constrained field applications, this limited scope does not provide insights into how YOLO compares with other prominent object detection frameworks, such as Faster R-CNN or RetinaNet. Future studies should benchmark these alternative models alongside YOLO to gain a comprehensive understanding of each architecture's relative strengths and limitations in ground-penetrating radar (GPR) applications. Such comparisons should also examine computational complexity, inference speed, and resource requirements, all critical factors for practical deployment. Additionally, the publicly available simulated dataset created in this study facilitates such benchmarking, allowing researchers to conduct comparative evaluations and make informed decisions about suitable detection models for bridge inspections. Second, this study specifically targets rebar detection, representing only one aspect of bridge deck integrity. Critical indicators of structural deterioration, including cracks, voids, corrosion, and delaminations, were not explicitly considered. Although our simulation framework is flexible, the current simulated dataset does not include these additional defect types, restricting immediate application to comprehensive, multi-class defect detection scenarios. Future research should therefore aim to generate and annotate synthetic datasets encompassing these broader defect categories. Developing versatile detection models capable of simultaneously identifying multiple defect types would significantly enhance the practical utility and comprehensiveness of automated bridge inspections.

Third, although the simulated dataset was designed to reflect diverse conditions, including variations in rebar spacing, concrete properties, antenna frequencies, and noise

levels, discrepancies between simulated and real-world data continue to pose challenges. These differences arise from simplifications in the simulation environment and can affect model performance when applied to field data. While pretraining on simulated data yielded strong results, particularly in data-limited scenarios, additional work is needed to close the domain gap. Future studies should investigate domain adaptation techniques such as adversarial learning, feature alignment, self-supervised methods, or hybrid fine-tuning to improve generalization and robustness. The simulated dataset released with this study provides a foundation for testing and refining such approaches. Fourth, multipath scattering and the resulting ghost artifacts, especially in scenes with closely spaced rebars, remain an unresolved challenge. These artifacts can generate false targets that impair detection accuracy. Although this study did not implement specific ghost suppression techniques, recent advancements in deep learning have shown potential in addressing this issue. For instance, Esposito et al. [33] applied a microwave tomography-based framework to filter multipath ghosts, and Lei et al. [34] proposed a signal enhancement method to suppress mutual interference. Integrating similar strategies into object detection pipelines could improve the reliability of GPR-based systems under complex subsurface conditions.

Finally, despite our real-world validation involving GPR data collected from four different bridges, the structural types, material conditions, environmental factors, and subsurface complexities represented within this limited dataset remain insufficiently comprehensive. The variability encountered in real bridge inspection scenarios far exceeds that captured by these four structures alone, which potentially limits model generalization and robustness assessments. Therefore, future research should expand the collection of real-world GPR datasets across a broader range of bridge types, geographical locations, and environmental conditions, providing more rigorous evaluations of model generalizability. Moreover, to address the practical challenges associated with the manual annotation of large real-world datasets, researchers might explore advanced labeling methodologies such as semi-supervised or active learning. These approaches could significantly streamline the annotation process, facilitating more efficient scaling of annotated real-world GPR data. Ultimately, addressing these limitations through targeted dataset expansion and methodological innovation will further improve the accuracy, reliability, and applicability of deep learning-driven GPR inspections for practical infrastructure monitoring.

6. Conclusions

This study investigated the application of deep learning models for automated rebar detection in GPR scans, addressing critical challenges in non-destructive bridge inspection. By leveraging both real and simulated GPR datasets, we examined the impact of domain-specific transfer learning and demonstrated its effectiveness in improving detection accuracy. The results indicate that pretraining on simulated GPR data provides more significant performance gains compared to conventional pretraining on general-purpose datasets like COCO, emphasizing the importance of domain-relevant feature learning for GPR-based object detection. A major contribution of this work is the development and utilization of a comprehensive dataset that includes 2255 real GPR images collected from four bridges and a large-scale simulated GPR dataset generated using gprMax. The synthetic dataset was designed to capture variations in material properties, rebar configurations, antenna frequencies, and subsurface conditions, providing a scalable and diverse training resource. The findings confirm that combining real and simulated data enhances model generalization, leading to more accurate and reliable rebar detection in bridge deck inspections. The study demonstrates that pretraining on simulated GPR data significantly improves detection accuracy, particularly when real training data are limited. In contrast, conventional COCO-based pretraining offers minimal benefits for GPR-based tasks, rein-

forcing the need for domain-specific datasets. The results further suggest that simulated GPR data provide a scalable and effective alternative to real-world training data, reducing reliance on extensive manual annotations while maintaining high detection performance.

Funding: This research received no external funding.

Data Availability Statement: The original data presented in the study are openly available on GitHub at https://github.com/InfraSmartLab/Open_GPR_Dataset_for_Bridge_Deck (accessed on 25 June 2025).

Acknowledgments: This research was funded in part by Kennesaw State University (KSU). The authors gratefully acknowledge the support from KSU. Any opinions, findings, recommendations, and conclusions in this paper are those of the author, and do not necessarily reflect the views of Kennesaw State University.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. ARTBA Bridge Report. Available online: https://artbabridgereport.org (accessed on 11 February 2025).
2. Sabellano, R.; Bandpey, Z.; Shokouhian, M. Evaluating and Predicting Deterioration of Bridges Using Machine Learning Applications. In Proceedings of the Structures Congress 2023, New Orleans, LA, USA, 3–6 May 2023; pp. 150–163. [CrossRef]
3. Wai-Lok Lai, W.; Dérobert, X.; Annan, P. A Review of Ground Penetrating Radar Application in Civil Engineering: A 30-Year Journey from Locating and Testing to Imaging and Diagnosis. *NDT E Int.* **2018**, *96*, 58–78. [CrossRef]
4. Persico, R. *Introduction to Ground Penetrating Radar: Inverse Scattering and Data Processing*; John Wiley & Sons: Hoboken, NJ, USA, 2014; ISBN 978-1-118-30500-3.
5. Parrillo, R.; Roberts, R.; Haggan, A. Bridge Deck Condition Assessment Using Ground Penetrating Radar. In Proceedings of the ECNDT Conference Proceeding, Berlin, Germany, 25–29 September 2006; pp. 25–29.
6. Hu, D.; Chen, J.; Li, S. Reconstructing Unseen Spaces in Collapsed Structures for Search and Rescue via Deep Learning Based Radargram Inversion. *Autom. Constr.* **2022**, *140*, 104380. [CrossRef]
7. Wang, M.; Hu, D.; Chen, J.; Li, S. Underground Infrastructure Detection and Localization Using Deep Learning Enabled Radargram Inversion and Vision Based Mapping. *Autom. Constr.* **2023**, *154*, 105004. [CrossRef]
8. Hu, D. Deep Learning–Based Framework for Bridge Deck Condition Assessment Using Ground Penetrating Radar. *J. Struct. Des. Constr. Pract.* **2025**, *30*, 4025041. [CrossRef]
9. Sui, X.; Leng, Z.; Wang, S. Machine Learning-Based Detection of Transportation Infrastructure Internal Defects Using Ground-Penetrating Radar: A State-of-the-Art Review. *Intell. Transp. Infrastruct.* **2023**, *2*, liad004. [CrossRef]
10. Warren, C.; Giannopoulos, A.; Giannakis, I. gprMax: Open Source Software to Simulate Electromagnetic Wave Propagation for Ground Penetrating Radar. *Comput. Phys. Commun.* **2016**, *209*, 163–170. [CrossRef]
11. Hu, D.; Li, S.; Chen, J.; Kamat, V.R. Detecting, Locating, and Characterizing Voids in Disaster Rubble for Search and Rescue. *Adv. Eng. Inform.* **2019**, *42*, 100974. [CrossRef]
12. Hu, D.; Chen, L.; Du, J.; Cai, J.; Li, S. Seeing through Disaster Rubble in 3D with Ground-Penetrating Radar and Interactive Augmented Reality for Urban Search and Rescue. *J. Comput. Civ. Eng.* **2022**, *36*, 04022021. [CrossRef]
13. Pashoutani, S.; Zhu, J. Ground Penetrating Radar Data Processing for Concrete Bridge Deck Evaluation. *J. Bridge Eng.* **2020**, *25*, 04020030. Available online: https://ascelibrary.org/doi/abs/10.1061/(ASCE)BE.1943-5592.0001566 (accessed on 11 February 2025). [CrossRef]
14. Pryshchenko, O.A.; Plakhtii, V.; Dumin, O.M.; Pochanin, G.P.; Ruban, V.P.; Capineri, L.; Crawford, F. Implementation of an Artificial Intelligence Approach to GPR Systems for Landmine Detection. *Remote Sens.* **2022**, *14*, 4421. [CrossRef]
15. Dinh, K.; Gucunski, N.; Zayed, T. Automated Visualization of Concrete Bridge Deck Condition from GPR Data. *NDT E Int.* **2019**, *102*, 120–128. [CrossRef]
16. Liu, H.; Lin, C.; Cui, J.; Fan, L.; Xie, X.; Spencer, B.F. Detection and Localization of Rebar in Concrete by Deep Learning Using Ground Penetrating Radar. *Autom. Constr.* **2020**, *118*, 103279. [CrossRef]
17. Asadi, P.; Gindy, M.; Alvarez, M. A Machine Learning Based Approach for Automatic Rebar Detection and Quantification of Deterioration in Concrete Bridge Deck Ground Penetrating Radar B-Scan Images. *KSCE J. Civ. Eng.* **2019**, *23*, 2618–2627. [CrossRef]
18. Dinh, K.; Gucunski, N.; Duong, T.H. An Algorithm for Automatic Localization and Detection of Rebars from GPR Data of Concrete Bridge Decks. *Autom. Constr.* **2018**, *89*, 292–298. [CrossRef]

19. Wang, Y.; Qin, H.; Zhang, D.; Wu, T.; Pan, S. Domain Adaption YOLO Network to Enhance Target Detection in GPR Images. *IEEE Trans. Geosci. Remote Sens.* **2024**, *62*, 5936415. [CrossRef]
20. Park, S.; Kim, J.; Jeon, K.; Kim, J.; Park, S. Improvement of GPR-Based Rebar Diameter Estimation Using YOLO-V3. *Remote Sens.* **2021**, *13*, 2011. [CrossRef]
21. Zaki, A.; Megat Johari, M.A.; Wan Hussin, W.M.A.; Jusman, Y. Experimental Assessment of Rebar Corrosion in Concrete Slab Using Ground Penetrating Radar (GPR). *Int. J. Corros.* **2018**, *2018*, 5389829. [CrossRef]
22. Senin, S.F.; Hamid, R. Ground Penetrating Radar Wave Attenuation Models for Estimation of Moisture and Chloride Content in Concrete Slab. *Constr. Build. Mater.* **2016**, *106*, 659–669. [CrossRef]
23. Zhang, Y.-C.; Yi, T.-H.; Lin, S.; Li, H.-N.; Lv, S. Automatic Corrosive Environment Detection of RC Bridge Decks from Ground-Penetrating Radar Data Based on Deep Learning. *J. Perform. Constr. Facil.* **2022**, *36*, 04022011. [CrossRef]
24. Feng, D.; Wang, X.; Wang, X.; Ding, S.; Zhang, H. Deep Convolutional Denoising Autoencoders with Network Structure Optimization for the High-Fidelity Attenuation of Random GPR Noise. *Remote Sens.* **2021**, *13*, 1761. Available online: https://www.mdpi.com/2072-4292/13/9/1761 (accessed on 11 February 2025). [CrossRef]
25. Liu, H.; Yue, Y.; Lai, S.; Meng, X.; Du, Y.; Cui, J.; Spencer, B.F. Evaluation of the Antenna Parameters for Inspection of Hidden Defects behind a Reinforced Shield Tunnel Using GPR. *Tunn. Undergr. Space Technol.* **2023**, *140*, 105265. [CrossRef]
26. Krauss, P.D.; Rogalla, E.A. *Transverse Cracking in Newly Constructed Bridge Decks*; NCHRP Report; National Academy Press: Washington, DC, USA, 1996.
27. Pashoutani, S.; Zhu, J. *Condition Assessment of Bridge Decks with Asphalt Overlay*; Nebraska Department of Transportation: Lincoln, NE, USA, 2019.
28. Salomon, A.L.; Moen, C.D. *Structural Design Guidelines for Concrete Bridge Decks Reinforced with Corrosion-Resistant Reinforcing Bars*; Virginia Department of Transportation: Richmond, VA, USA, 2014.
29. Elgabbas, F.; Ahmed, E.A.; Benmokrane, B. Experimental Testing of Concrete Bridge-Deck Slabs Reinforced with Basalt-FRP Reinforcing Bars under Concentrated Loads. *J. Bridge Eng.* **2016**, *21*, 4016029. [CrossRef]
30. Davis, J.; Huang, Y.; Millard, S.G.; Bungey, J.H. Determination of Dielectric Properties of Insitu Concrete at Radar Frequencies. In Proceedings of the International Symposium on Non-Destructive Testing in Civil Engineering, Berlin, Germany, 16–19 September 2003.
31. Porubiaková, A.; Komačka, J. A Comparison of Dielectric Constants of Various Asphalts Calculated from Time Intervals and Amplitudes. *Procedia Eng.* **2015**, *111*, 660–665. [CrossRef]
32. Kaur, P.; Dana, K.J.; Romero, F.A.; Gucunski, N. Automated GPR Rebar Analysis for Robotic Bridge Deck Evaluation. *IEEE Trans. Cybern.* **2016**, *46*, 2265–2276. [CrossRef] [PubMed]
33. Esposito, G.; Catapano, I.; Ludeno, G.; Soldovieri, F.; Gennarelli, G. A Deep Learning Strategy for Multipath Ghosts Filtering via Microwave Tomography. *IEEE Trans. Geosci. Remote Sens.* **2023**, *62*, 5100314. [CrossRef]
34. Lei, W.; Tan, X.; Luo, C.; Xue, W. Mutual Interference Suppression and Signal Enhancement Method for Ground-Penetrating Radar Based on Deep Learning. *Electronics* **2024**, *13*, 4722. [CrossRef]

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