DeepGPR: Learning to Identify Moisture Defects in Building Envelope Assemblies from Ground Penetrating Radar

Bilal Sher and Chen Feng

Tandon School of Engineering, New York University, USA

bas9876@nyu.edu, cfeng@nyu.edu

Abstract -

Conventionally used moisture detection equipment such as infrared scanners and capacitance meters require a trained interpreter to understand moisture issues on rooftops. Additionally, conventional sensors can only provide reliable results in specific environmental conditions. In this paper, we will discuss the various methods used for roof moisture scans and their limitations. We will then provide an in-depth analysis of GPR paired with deep segmentation neural networks for roof moisture scans, including its advantages, limitations, and potential applications. We will also present a case study demonstrating the effectiveness of this approach in detecting moisture damage in a real-world scenario. Our preliminary experiments find that deep neural networks are effective in segmenting GPR radargrams and finding moisture, with particular neural networks more effective than others.

Keywords -

Ground penetrating radar; moisture detection; building envelope analysis; rooftop moisture survey

1 Introduction

Roof moisture scans are essential for maintaining the structural integrity of buildings by identifying moisture damage to the roof. Traditional methods such as visual inspection, infrared thermography, nuclear moisture gauges, and capacitance meters have limitations that can affect the accuracy and reliability of results. However, Ground Penetrating Radar (GPR) paired with Artificial Intelligence (AI) has emerged as a promising solution for conducting roof moisture scans.

GPR is a non-destructive testing technique that uses electromagnetic waves to detect and image subsurface features of materials. The integration of GPR and AI enables a comprehensive and accurate assessment of moisture levels in the roof structure. GPR provides high-resolution imaging of the subsurface, while AI can process the data and identify patterns that may not be visible to the human eye. This combination allows for early detection of moisture damage, reducing repair costs and increasing the lifespan of the roof.

GPR technology captures data by measuring the dielectric properties of materials, making it an effective tool for detecting moisture within building envelopes and roofs [1, 2]. The Proceq GP8800 SFCW handheld GPR sensor is a popular device for capturing GPR data, recording data at a fixed distance interval of 1 cm [3]. However, interpreting the data can be challenging, which has limited its widespread adoption [4].

To overcome this challenge, a deep segmentation neural network was tested to pair with GPR for moisture detection applications on building envelopes, particularly roof assemblies. The preliminary experiments found that deep neural networks are effective in segmenting GPR radargrams and finding moisture, with particular neural networks being more effective than others. By leveraging deep learning techniques, GPR can provide more accurate and reliable results in detecting moisture, enhancing its potential for building maintenance and inspection.

The potential benefits of GPR for detecting moisture issues on rooftops and building facades are significant. By identifying moisture issues early on, building owners and maintenance professionals can address them before they become costly problems. Additionally, GPR can help ensure the safety and longevity of structures by detecting potential structural issues caused by moisture. While more research is needed to fully realize the potential of GPR in this application, it is clear that the technology offers a valuable tool for enhancing building maintenance and inspection.

2 Related Works

Researchers have explored the use of artificial intelligence (AI) and neural networks in analyzing ground penetrating radar (GPR) scans. This has included using AI to assess moisture content in various materials, such as concrete and soil. One example is Kilic and Unluturk [5], which used a simple artificial neural network to analyze a GPR scan and classify it as wet or dry. Others, such as Zhang et al. [6], have used more advanced techniques, like Resnet and YoloV2, to draw bounding boxes around areas of suspected moisture. Researchers have also attempted to estimate soil moisture content using GPR scans, such as with Qiao et al. [7]'s radial basis function neural net-

work. Zheng et al. [8] improved on this method by using a convolutional neural network (CNN) connected to a fully connected layer to analyze soil moisture content. Other authors have used similar CNN and fully connected layer set-ups to find object representations within GPR scans [9, 10, 11, 12, 13]. Hou et al. [14] implemented a Mask R-CNN to segment hyperbolic signatures of rebar in GPR scans of a bridge deck. Our work improves upon these methods in a number of ways. First, it presents a new approach to simulating moisture in building envelope assemblies. Second, it tests various segmentation models equipped with a line scan conversion block to determine if a portion of the GPR scan is wet or dry. Finally, it adds five additional data channels for analysis, including a max-amplitude normalization channel, a time-gain channel, and three additional channels that are power spectral density images based on the raw GPR scan channel, the max amplitude normalized channel, and the time-gain channel.

3 Preparation and Data Collection

3.1 GPR Technology

GPR is a non-destructive testing technique that uses electromagnetic waves to detect and image subsurface features of materials. It works by measuring the dielectric properties of materials, making it an effective tool for detecting moisture within building envelopes and roofs, as the dielectric constant of water is much higher than other materials commonly used in construction [15, 16, 17, 18, 19]. The GPR technology can be captured using different methods, with impulse and stepped frequency continuous wave (SFCW) being the most common. In impulse GPR, a fixed frequency pulse is sent into a medium, and the reflected signal is detected [4]. SFCW GPR, on the other hand, sends a continuous signal with a modulated frequency into the medium and listens for the reflectances from various wavelengths. Some studies have shown that SFCW is the superior configuration for capturing data of smaller, shallow targets, which is particularly relevant in the structural and building analysis use case of GPR [20, 4].

3.2 Data Collection in Lab Setting

A novel testbed was created to test building assemblies with simulated moisture contents (See Figure 1). The base of the test bed had a 5/16" thick 4' x 8' standard-size OSB sheathing board. Moisture was simulated by placing moistened paper towels inside a plastic Ziploc bag on various portions of the test bed. The base of the test bed had a standard-size OSB sheathing board. Normal OSB moisture content can vary from 11.5% to 12.5% in New York City [21]. A base moisture content of 11.7% was assumed, and the amount of water needed to be added to the paper towels was calculated accordingly. Due to size and spacing requirements on the test kit as well as the range of moisture content required to cause mold growth, the following breakdown in tests was chosen:

Due to size and spacing requirements on the test kit as well as the range of moisture content required to cause mold growth, the following breakdown in tests was chosen:

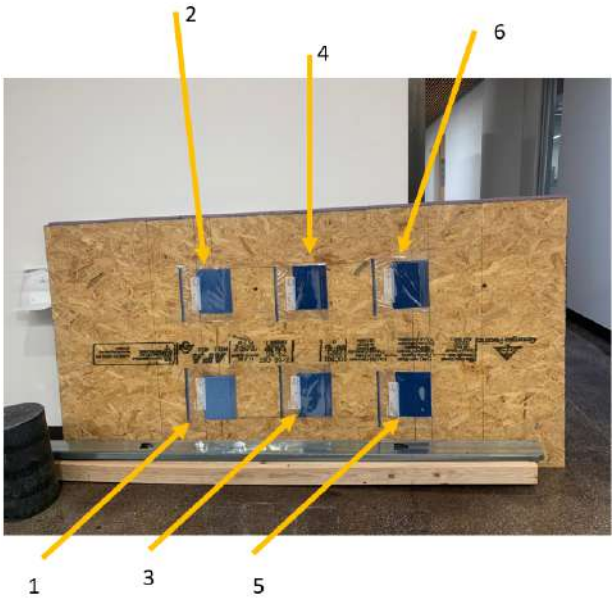

Figure 1. Moisture Testbed

Packet	1	2	3	4	5	6
MC (%)	11.7	18.7	25.8	32.8	39.8	46.8

Various building materials were added to create a dataset of GPR scans (See Figure 2). 48 different combinations of building material were captured, with each configuration producing 6 different scans. Each configuration produced 6 different scans making it equivalent to 288 different GPR scans conducted with varying moisture content and different superimposed building envelope materials. The building materials superimposed onto the testbed included:

1. $\frac{7}{16}in$ OSB sheet
2. $1in$ Rigid Foam Insulation
3. Timber 2×4 studs
4. Metal 2×4 studs
5. $8in \times 4in \times 2in$ clay masonry wall bricks
6. $4in \times 16in \times 8in$ hollow concrete blocks
7. R19 Batt Insulation

Thin materials like vapor barriers, air barriers, and waterproofing membranes were not included in the test as

Figure 2. Various testbed configurations

they don't affect GPR readings. The test simulated moisture condensing on both exterior and interior sides of a wall assembly, depending on the climate and time of year. Research from other authors confirms that moisture is not evenly distributed throughout a wall or roof assembly, and it is more concentrated at either the interior or exterior side of the assembly [22]. This test was also designed based on research that shows that moisture content at or above 19% will catalyze mold growth [23]. Figure 3 shows the testbed set up with layers of interior and exterior insulation to simulate an EIFS assembly. The test bed was simply supported over two tables, and the tables were placed at least 30 cm from the edges of the moisture areas to ensure that the table legs did not affect the GPR readings. Overall, the test bed provided a comprehensive dataset of GPR scans conducted with varying moisture content and different building envelope materials, which can help researchers and practitioners in developing effective moisture control strategies for building assemblies.

Figure 3. Testbed configuration for multiple layers of insulation

3.3 Pre-Processing of the Collected GPR Data

During the GPR testing process, data was collected by moving the Proceq GPR 8800 unit over the moist sections of the testbed, which were premarked to ensure ease of testing. As the wheel attached to the GPR unit moves, it records data. However, discontinuities may occur in the radargram when the unit gets caught on surfaces being scanned. To address this, scans were redone more slowly when the unit was caught on a brick. In some cases, the GPR unit was raised slightly to allow it to clear obstacles such as slightly protruding bricks. The scans were recorded on the Proceq GPR Live app and uploaded to the Proceq servers, where they were exported as SEG-Y files. Different GPR manufacturers have different data exporting methods, but the benefit of using equipment that exports data in an SEG-Y format is that we can easily process it in Python through the SEGY-SAK python library. The Software Underground community routinely performs machine learning and deep learning on GPR-acquired datasets. The SEG-Y file is processed with the SEGY-SAK API to extract a numpy array of the raw, unfiltered GPR scan data. Each column in the numpy array represents a trace. This library is supported by the Software Underground community, which routinely performs machine learning and deep learning on GPR-acquired datasets.

$$\boldsymbol{G_{i,j}} = \begin{bmatrix} t_{r_{1,1}} & \dots & t_{r_{1,n}} \\ \vdots & \ddots & \vdots \\ t_{r_{m,1}} & \dots & t_{r_{m,n}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{T_{r_1}} & \dots & \boldsymbol{T_{r_n}} \end{bmatrix} \quad (1)$$

The data obtained from the GPR was raw radargram data, which was then processed by maximum amplitude normalizing the data. This is a noise removal technique that is valid on flat/level surfaces or surfaces where the GPR unit is a constant distance from the surface being measured. Maximum amplitude trace normalization finds the average peak amplitude across all traces and scales each trace so that its maximum amplitude is now the average peak amplitude.

$$A_p = \frac{1}{n} \sum_{i=1}^{n} \max(\boldsymbol{T_{r_i}}) \quad (2)$$

$$\tilde{\boldsymbol{G}}_{\boldsymbol{i,j}} = \begin{bmatrix} t_{r_{1,1}} \cdot \frac{A_p}{\max(\boldsymbol{T_{r_1}})} & \dots & t_{r_{1,n}} \cdot \frac{A_p}{\max(\boldsymbol{T_{r_n}})} \\ \vdots & \ddots & \vdots \\ t_{r_{m,1}} \cdot \frac{A_p}{\max(\boldsymbol{T_{r_1}})} & \dots & t_{r_{m,n}} \cdot \frac{A_p}{\max(\boldsymbol{T_{r_n}})} \end{bmatrix} \quad (3)$$

In addition to applying maximum amplitude trace normalization, a temporal signal gain was also applied to the images. There are many ways of applying a temporal signal gain, but the linear and exponential methods were used in the tests below. For linear signal gains, $\boldsymbol{Gn^l_{m\times n}}$ later signals are enhanced by multiplying the traces with a linearly increasing gain vector that is the same length as the trace. For exponential signal gains, $\boldsymbol{Gn^e_{m\times n}}$, the signal can be enhanced by multiplying the traces with an exponentially increasing gain vector.

$$\boldsymbol{Gn}^{\boldsymbol{l}}_{\boldsymbol{m\times n}} = C \cdot \begin{bmatrix} 1 & \dots & 1 \\ \vdots & \ddots & \vdots \\ m & \dots & m \end{bmatrix} \tag{4}$$

$$\tilde{\boldsymbol{G}}'_{\boldsymbol{i,j}} = \boldsymbol{G}_{\boldsymbol{i,j}} \odot \boldsymbol{Gn}^{\boldsymbol{l}}_{\boldsymbol{m\times n}} \tag{5}$$

$$\boldsymbol{Gn}^{\boldsymbol{e}}_{\boldsymbol{m\times n}} = \begin{bmatrix} 1^c & \dots & 1^c \\ \vdots & \ddots & \vdots \\ m^c & \dots & m^c \end{bmatrix} \tag{6}$$

$$\tilde{\boldsymbol{G}}'_{\boldsymbol{i,j}} = \boldsymbol{G}_{\boldsymbol{i,j}} \odot \boldsymbol{Gn}^{\boldsymbol{e}}_{\boldsymbol{m\times n}} \tag{7}$$

This resulted in 3 different GPR scans: the raw data scan $\boldsymbol{G}_{\boldsymbol{i,j}}$, the maximum amplitude trace normalized scan $\tilde{\boldsymbol{G}}_{\boldsymbol{i,j}}$, and the temporal signal gain scan $\tilde{\boldsymbol{G}}'_{\boldsymbol{i,j}}$. These were further enhanced by finding a power spectral density image associated with each scan. The power spectral density image was generated by finding the PSD of each individual trace and then concatenating them into a 2D image.

$$PSD(\boldsymbol{G}_{\boldsymbol{i,j}}, \tilde{\boldsymbol{G}}_{\boldsymbol{i,j}}, \tilde{\boldsymbol{G}}'_{\boldsymbol{i,j}}) = \boldsymbol{P}_{\boldsymbol{i,j}}, \tilde{\boldsymbol{P}}_{\boldsymbol{i,j}}, \tilde{\boldsymbol{P}}'_{\boldsymbol{i,j}} \tag{8}$$

There are other GPR scan normalization techniques that are not applicable to our current scanning data set and would not produce additional usable information. As a result, each GPR scan could be represented as a 6 channel tensor $\boldsymbol{I}^{(m\times n\times 6)} = \left[\boldsymbol{G}_{\boldsymbol{i,j}}, \tilde{\boldsymbol{G}}_{\boldsymbol{i,j}}, \tilde{\boldsymbol{G}}'_{\boldsymbol{i,j}}, \boldsymbol{P}_{\boldsymbol{i,j}}, \tilde{\boldsymbol{P}}_{\boldsymbol{i,j}}, \tilde{\boldsymbol{P}}'_{\boldsymbol{i,j}}\right]$. This 6 channel tensor served as the input to the deep learning network. Data was annotated by denoting sections, i.e., multiple consecutive whole traces, of the radargram as being either moist or dry. The locations of the simulated moisture in the test bed was used as a guide because there were minor changes in the starting or ending position of the GPR scan from run to run. This was used to produce a 2D image mask, $\boldsymbol{M}_{1\times n}$, with a pixel height of 1 and a width that represented the distance the GPR moved.

4 GPR Segmentation by Deep Learning

Data augmentation is a technique used to increase the diversity of data in a dataset for training machine learning models. In this study, random horizontal flips and random resizing were used for data augmentation. Horizontal flipping had a 50% chance of occurring while resizing had a 75% chance of occurring. Resizing involved expanding or shrinking the horizontal portions of the scans by up to 40% compared to the original horizontal scan length. After resizing, the binary 0 or 1 representation of the masks would no longer hold, and any mask value above 0 was set to 1.

To ensure that the scans could be easily processed by deep learning segmentation algorithms, the height and width of the batches were fixed to 672 and 128, respectively. Scans were top-padded by copying the first row of the input tensor. If the scan width was less than 128, scans were left-padded, and if the scan width was greater than 128, scans were left-cropped. The left padding was a copy of the leftmost trace.

Every segmentation model had a line scan conversion head at the end. This conversion head was a block consisting of a 2D convolutional layer with a 3x3 kernel and 1x1 padding, a 2D bilinear upsampling layer that brought the output size back up to the original input size followed by ReLU activation, and a 2D convolutional layer with a kernel of input height x 3 and padding applied width-wise, but not height-wise, and a channel reduction to 1.

A series of experiments were conducted to determine the optimal model for accurately segmenting a raw GPR scan. Various hyperparameters and model configurations were tested. The results of these tests will be used to guide future tests on data obtained from real-world field tests.

A number of hyperparameters were kept the same during all tests. Each individual input in the batch was standardized per channel using the following formula:

$$\boldsymbol{I}^{(m\times n\times 6)}_{std} = \left[\quad \dots \quad \left[\frac{\boldsymbol{C}_{**k} - \mathbf{mean}(\boldsymbol{C}_{**k})}{\mathbf{std}(\boldsymbol{C}_{**k})}\right] \quad \dots \quad \right] \tag{9}$$

For all tests, the learning rate started off at 1e−4 and was reduced to 5e−5 after 125 epochs. All tests were run for 250 epochs.

Three metrics were used to evaluate the success of the model:

1. Intersection of Union
2. Dice Score
3. Pixel Accuracy (Accuracy)

IoU and dice score are crucial metrics for evaluating line scan segmentation algorithms, with IoU being the most important as this is a segmentation problem.

Deep Supervision: UNet is a widely used image segmentation model, with nearly 60,000 citations as of April 2023. [24] However, it struggles with segmenting fine details, which can represent a problem for GPR scans as they can contain finer details than typical images. To address this, Zhou et al created UNet++, which adds dense skip connections from higher and lower levels of the segmentation encoder and includes deep supervision. This helps to train earlier layers and ensure that finer details are detected. The effectiveness of UNet++ was tested to evaluate whether it can improve segmentation accuracy.

Encoder Type: Segmentation models have an encoder and decoder. The encoder converts data into a form that the decoder can process. VGG16 is the standard encoder for UNet and UNet++. Other encoder networks like ResNets and Inceptionv4 can be used instead of VGG16. Different

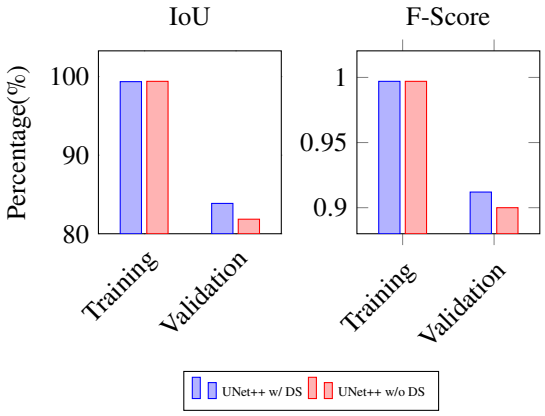

Figure 4. Effect of Deep Supervision on Model Accuracy

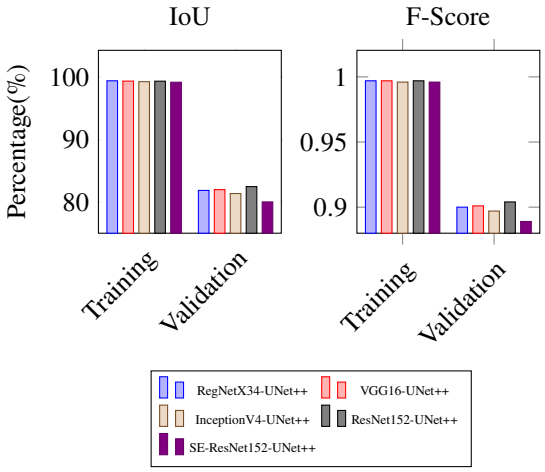

Figure 5. Effect of Encoder Type on Model Accuracy

networks have different capabilities. A test was conducted to see the effect of different encoder types on segmentation.

Encoder Depth: Deep Learning Encoders extract features from images to create a feature map, with deeper encoders extracting more features up to a limit. Using deeper encoder networks can improve models. To test encoder depth effects, increasingly deep ResNet encoders were added to a UNet++ decoder.

Decoder Type: Various segmentation decoders exist beyond the standard UNet and UNet++ algorithms, including Feature Pyramid Networks, DeepLabV3, DeepLabV3+, and Pyramid Attention Networks. Feature Pyramid Networks are similar to UNets but use skip connections and lateral connections passed through a 1x1 convolution. DeepLabV3 concatenates dilated convolutions over

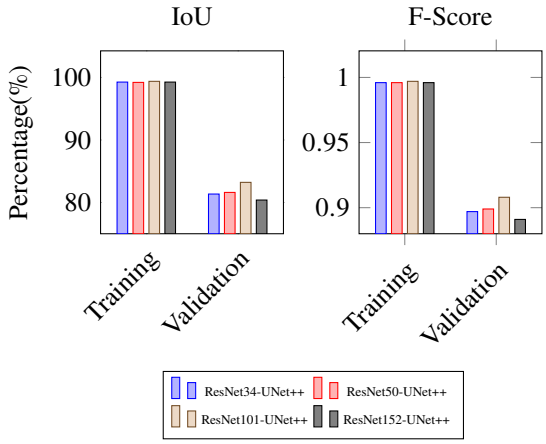

Figure 6. Effect of Encoder Depth on Model Accuracy

an encoded feature map to obtain global features, while DeepLabV3+ is an improved version of this model. Pyramid Attention Networks combine high and low level features using a feature pyramid attention module and global upsampling attention module. To evaluate the effectiveness of these decoders in segmenting GPR scans, tests were conducted using different encoder types. A ResNet34 encoder was used in the first set of tests, followed by an InceptionV4 encoder in the second set, and a RegNetX32 encoder in the third set. These tests aimed to determine the most effective decoder for the GPR scan segmentation task, while also ensuring that the encoder did not significantly affect the model's accuracy.

Decoder Depth: The UNet and UNet++ models use a VGG16 network as their standard decoder. This network decodes by progressively lowering the number of channels within layers in each consecutive block of the decoder network until eventually it the number of channels has been reduced to the number of output channels. To train a CNN to detect more features, the decoder can be modified to have significantly more channels. In testing, a UNet++ algorithm was trained with various encoders and different decoder depths to evaluate their performance.

Type 1 (T1) Decoders had a [16, 32, 64, 128, 256] channel structure. Type 2 (T2) Decoders had a [32, 64, 128, 256, 512] channel structure.

ROC Curves & False Positive vs. False Negative Rate To assess the performance of moisture detection models, select models were analyzed using receiver operating characteristic (ROC) curves (See Figure 9). Current moisture detection methods are criticized for high false positive rates, which can be addressed by thresholding, but this increases false negatives. Detection error tradeoff (DET)

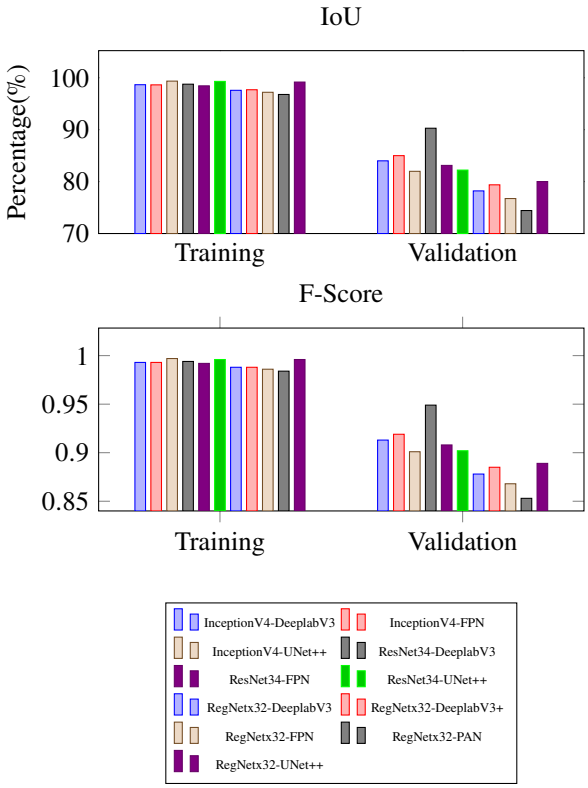

Figure 7. Effect of Decoder Type on Model Accuracy

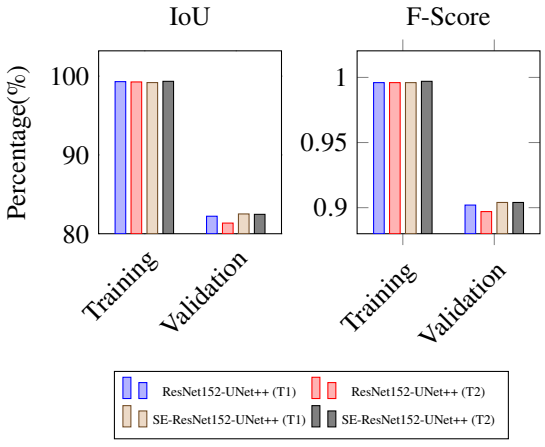

Figure 8. Effect of Decoder Depth on Model Accuracy

curves explore this trade-off and provide insight into segmentation models' performance at different scales.

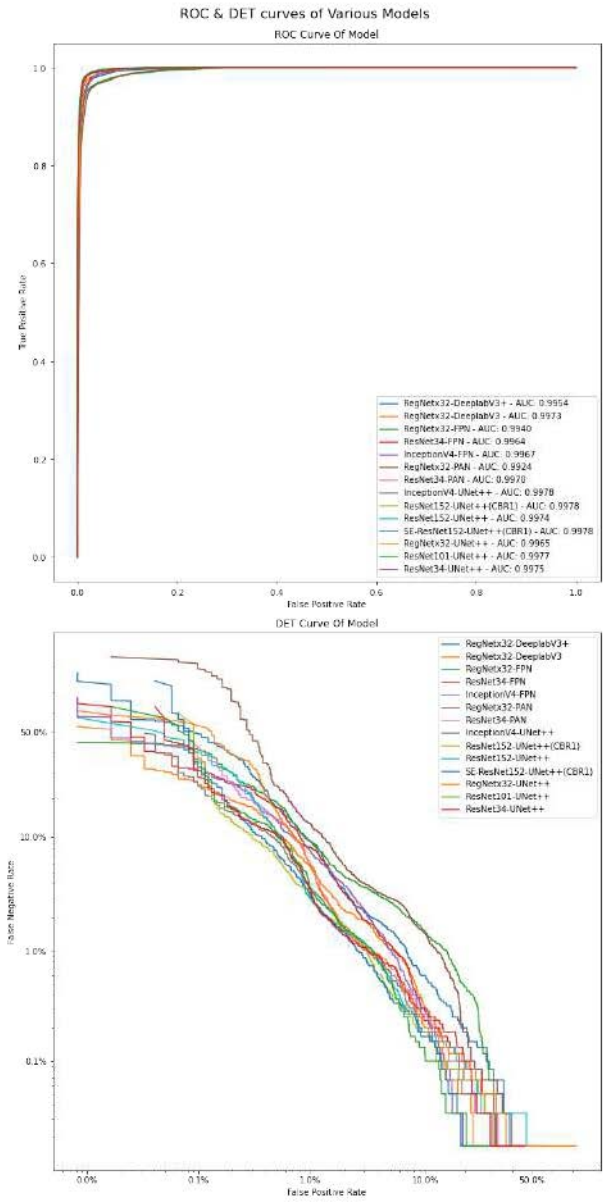

Figure 9. ROC & DET curves of various evaluated models

5 Conclusion

5.1 Testbed

The testbed simulates moisture in building envelopes using various construction materials. The conditions simulated are in some aspects more rigorous than real-world

non-destructive testing for moisture intrusion. Moisture is located under multiple dry layers of building material, making electrical impedance testing and IRT the only viable testing methods. Additionally, scanning was conducted when moisture was at a depth of up to 2 feet, which is beyond the range of non-destructive moisture testing methods.

Future works to improve upon this testbed would be to:

1. Soak various materials to artificially increase their internal moisture content.
2. Add additional building materials such as different kinds of rigid and non-rigid insulation, conductive waterproofing membranes like EPDM, and exterior wall finishings like stucco.
3. To conduct tests at various temperatures

However, the current testbed setup establishes a standard protocol for conducting future tests. Field testing will further confirm the efficacy of this testbed setup.

5.2 Data Analysis

Providing a six-channel tensor containing scans and PSD images gave the network additional data. Another data point could be a 3D tensor created from converting GPR traces into a spectrogram or a Mel-spectrogram. Combining the 2D and 3D data could present a problem, but it can be fused within the model.

5.3 Deep Learning Experiments

The ResNet34 encoder with DeeplabV3 decoder network achieved the highest validation IoU while maintaining high training accuracy. The InceptionV4-FPN model also performed well in terms of validation IoU. UNet++ decoders achieved the highest AUC scores but underperformed on validation tests, indicating overfitting. The best model had the smallest discrepancy between training and validation accuracies and achieved over 9% pixel accuracy on the validation set. However, the model struggled to classify the edges of moisture as wet or dry. The relatively small size of the training and validation dataset may have impacted the model's performance.

6 Future Work

The predicted line scan mask could then be fused with the path of the GPR to create a moisture survey map over a rooftop. This allows for the creation of automated moisture survey maps using robotics, SLAM, and deep learning.

Acknowledgment. This work is supported by NSF CNS-2228568 and TI-2232494, and DOE E-ROBOT challenge.

References

[1] Andrew Horsley and David S. Thaler. Microwave detection and quantification of water hidden in and on building materials: Implications for healthy buildings and microbiome studies. 19 (1):67. ISSN 1471-2334. doi:10.1186/s12879-019-3720-1. URL https://doi.org/10.1186/s12879-019-3720-1.

[2] Iván Garrido, Mercedes Solla, Susana Lagüela, and Norberto Fernández. IRT and GPR Techniques for Moisture Detection and Characterisation in Buildings. 20(22):6421. ISSN 1424-8220. doi:10.3390/s20226421. URL https://www.mdpi.com/1424-8220/20/22/6421.

[3] Proceq GP8800 Structural imaging and inspection with SFCW ground penetrating radar technology. URL www.screeningeagle.com/en/products/proceq-gp8800-ultraportable-concrete-gpr-radar.

[4] Erica Carrick Utsi. *Ground Penetrating Radar: Theory and Practice*. Elsevier Science & Technology. ISBN 978-0-08-102217-7. URL http://ebookcentral.proquest.com/lib/nyulibrary-ebooks/detail.action?docID=4844133.

[5] Gokhan Kilic and Mehmet S. Unluturk. Performance evaluation of the neural networks for moisture detection using GPR. 29(4):283–296. ISSN 1058-9759. doi:10.1080/10589759.2014.941839. URL https://doi.org/10.1080/10589759.2014.941839.

[6] Jun Zhang, Xing Yang, Weiguang Li, Shaobo Zhang, and Yunyi Jia. Automatic detection of moisture damages in asphalt pavements from GPR data with deep CNN and IRS method. 113:103119. ISSN 0926-5805. doi:10.1016/j.autcon.2020.103119. URL https://www.sciencedirect.com/science/article/pii/S0926580519311537.

[7] Xu Qiao, Feng Yang, and Xianlei Xu. The prediction method of soil moisture content based on multiple regression and RBF neural network. In *Proceedings of the 15th International Conference on Ground Penetrating Radar*, pages 140–143. doi:10.1109/ICGPR.2014.6970402.

[8] Jing Zheng, Xingzhi Teng, Jie Liu, and Xu Qiao. Convolutional Neural Networks for Water Content Classification and Prediction With Ground Penetrating Radar. 7:185385–185392. ISSN 2169-3536. doi:10.1109/ACCESS.2019.2960768.

[9] Jun Sonoda and Tomoyuki Kimoto. Object Identification form GPR Images by Deep Learning. In *2018 Asia-Pacific Microwave Conference (APMC)*, pages 1298–1300. doi:10.23919/APMC.2018.8617556.

[10] Zhongming Xiang, Abbas Rashidi, and Ge (Gaby) Ou. An Improved Convolutional Neural Network System for Automatically Detecting Rebar in GPR Data. In *Computing in Civil Engineering 2019*, pages 422–429. American Society of Civil Engineers. ISBN 978-0-7844-8243-8. doi:10.1061/9780784482438.054. URL http://ascelibrary.org/doi/10.1061/9780784482438.054.

[11] Lance E. Besaw and Philip J. Stimac. Deep convolutional neural networks for classifying GPR B-scans. In *Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XX*, volume 9454, pages 385–394. SPIE. doi:10.1117/12.2176250.

[12] Takahiro Yamaguchi, Tsukasa Mizutani, and Tomonori Nagayama. Mapping Subsurface Utility Pipes by 3-D Convolutional Neural Network and Kirchhoff Migration Using GPR Images. 59(8):6525–6536. ISSN 1558-0644. doi:10.1109/TGRS.2020.3030079.

[13] F. Ponti, F. Barbuto, P. P. Di Gregorio, F. Frezza, F. Mangini, R. Parisi, P. Simeoni, and M. Troiano. GPR radargrams analysis through machine learning approach. 35(12):1678–1686. ISSN 0920-5071. doi:10.1080/09205071.2021.1906329. URL https://doi.org/10.1080/09205071.2021.1906329.

[14] Feifei Hou, Wentai Lei, Shuai Li, Jingchun Xi, Mengdi Xu, and Jiabin Luo. Improved Mask R-CNN with distance guided intersection over union for GPR signature detection and segmentation. 121:103414. ISSN 0926-5805. doi:10.1016/j.autcon.2020.103414. URL https://www.sciencedirect.com/science/article/pii/S0926580520309948.

[15] Stanislav Stefanov Zhekov, Ondrej Franek, and Gert Frolund Pedersen. Dielectric Properties of Common Building Materials for Ultrawideband Propagation Studies [Measurements Corner]. 62(1):72–81. ISSN 1558-4143. doi:10.1109/MAP.2019.2955680.

[16] Yosef Pinhasi, Asher Yahalom, and Sergey Petnev. Propagation of ultra wide-band signals in lossy dispersive media. In *2008 IEEE International Conference on Microwaves, Communications, Antennas and Electronic Systems*, pages 1–10. doi:10.1109/COMCAS.2008.4562803.

[17] M. Uematsu and E. U. Frank. Static Dielectric Constant of Water and Steam. 9 (4):1291–1306. ISSN 0047-2689, 1529-7845. doi:10.1063/1.555632. URL http://aip.scitation.org/doi/10.1063/1.555632.

[18] Qingqing Cao and Imad L. Al-Qadi. Effect of Moisture Content on Calculated Dielectric Properties of Asphalt Concrete Pavements from Ground-Penetrating Radar Measurements. 14(1):34.

[19] Goran Stojanović, Milan Radovanović, Mirjana Malešev, and Vlastimir Radonjanin. Monitoring of Water Content in Building Materials Using a Wireless Passive Sensor. 10(5):4270–4280. ISSN 1424-8220. doi:10.3390/s100504270. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3292119/.

[20] Guido Tronca, Isaak Tsalicoalou, Samuel Lehner, and Gianluca Catanzariti. Comparison of pulsed and stepped frequency continuous wave (SFCW) GPR systems. In *2018 17th International Conference on Ground Penetrating Radar (GPR)*, pages 1–4. doi:10.1109/ICGPR.2018.8441654.

[21] W. T. Simpson. Specific gravity, moisture content, and density relationship for wood. URL https://www.fs.usda.gov/treesearch/pubs/5914.

[22] Andre Desjarlais. A New Look at Moisture Control in Low-Slope Roofing. In *Fourth International Symposium on Roofing Technology*. URL https://nrcawebstorage.blob.core.windows.net/files/TechnicalLibraryNRCA/5855.pdf.

[23] Daniela Tudor, Sara C Robinson, and Paul A Cooper. The influence of moisture content variation on fungal pigment formation in spalted wood. 2:69. ISSN 2191-0855. doi:10.1186/2191-0855-2-69. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3577483/.

[24] [1505.04597] U-Net: Convolutional Networks for Biomedical Image Segmentation. URL https://arxiv.org/abs/1505.04597.