



HAL
open science

Classification of Buried Objects From Ground Penetrating Radar Images by Using Second-Order Deep Learning Models

Douba Jafuno, Ammar Mian, Guillaume Ginolhac, Nickolas Stelzenmuller

► **To cite this version:**

Douba Jafuno, Ammar Mian, Guillaume Ginolhac, Nickolas Stelzenmuller. Classification of Buried Objects From Ground Penetrating Radar Images by Using Second-Order Deep Learning Models. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 2025, 18, pp.3185-3197. 10.1109/JSTARS.2024.3524424 . hal-04902720

HAL Id: hal-04902720

<https://hal.science/hal-04902720v1>

Submitted on 21 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Classification of Buried Objects from Ground Penetrating Radar Images by using Second Order Deep Learning Models

Douba Jafuno, Ammar Mian, Guillaume Ginolhac, *Senior Member, IEEE*, Nickolas Stelzenmuller

Abstract—In this paper, a new classification model based on covariance matrices is built in order to classify buried objects. The inputs of the proposed models are the hyperbola thumbnails obtained with a classical Ground Penetrating Radar (GPR) system. These thumbnails are then inputs to the first layers of a classical CNN, which then produces a covariance matrix using the outputs of the convolutional filters. Next, the covariance matrix is given to a network composed of specific layers to classify Symmetric Positive Definite (SPD) matrices. We show in a large database that our approach outperform shallow networks designed for GPR data and conventional CNNs typically used in computer vision applications, particularly when the number of training data decreases and in the presence of mislabeled data. We also illustrate the interest of our models when training data and test sets are obtained from different weather modes or considerations.

Index Terms—Ground Penetrating Radar, covariance matrices, buried objects classification, Symmetric Positive Definite matrix networks

I. INTRODUCTION

The Ground Penetrated Radar (GPR) is a RADAR system that provides an image of the underground [1]–[3]. In particular, it can be used to image buried objects such as mines, pipes (metal, plastic, cast iron, etc.) and even cavities. The main drawback of GPR images is that they are very noisy, in particular due to the clutter that is the sum of all contributions from micro-scatterers in the ground but also because of the strong answer of the different layers of the ground. The buried objects are therefore often difficult to detect/locate and even more difficult to classify. One solution is to use complex systems such as stepped frequency RADAR [4], Multiple Input Multiple Output (MIMO) [5], or polarimetric sensors. But the high cost of these devices is not always attractive for industrial or civil engineering applications. In this paper, we consider classical GPR systems emitting a single wave, called Ricker, and whose image is created by a displacement on one axis of the transmitting/receiving system. In the normal configuration, the RADAR is positioned very close to the ground. In our case, we will study the possibility of placing the RADAR at a certain height above the ground. This study will enable us to assess the robustness of our approach in the case of using GPR placed on a drone.

Douba Jafuno, Ammar Mian and Guillaume Ginolhac are with LISTIC (EA3703), University Savoie Mont-Blanc, FRANCE. Nickolas Stelzenmuller is with Geolithe, FRANCE. This work has been done thanks to the facilities offered by the Univ. Savoie Mont-Blanc - CNRS/IN2P3 MUST computing center.

As noticed previously, the bad quality of the GPR image requires the use of various signal processing, image processing or machine learning techniques to achieve just the right detection and localization performance. In machine learning, it is possible to use deep learning techniques for denoising or inversion [6]–[9], auto-encoders for detection [10] or pattern recognition approaches for localization of buried objects [11]. In signal processing, it is possible to use statistical methods normally used in detection [12]–[15], particle filtering [16], Markov fields [17] or algebraic algorithms [18], [19]. In image processing, most of methods are based on inversion [20], [21], compressive sensing [22] or dictionary learning [23]. A robust inversion method has been proposed in [24], which achieves good performance whatever the type of soil or buried object. All these works are essential for good object detection and localization, but will not suffice if we wish to classify them and thus determine their physical properties. In this paper, we are interested in this last step. Before introducing the proposed approach and those of the literature, we give the different needed assumptions: firstly all the buried objects are detected and correctly localized and secondly we have a certain amount of labeled data at our disposal, enabling us to develop supervised approaches.

For the classification of buried objects from GPR images, a number of studies already exist, based either on classical signal processing techniques [25], [26], machine learning [27], [28] or deep networks [29]–[32]. In all these algorithms, the classification is based on the shape of the hyperbola (in both axes), which is partly related to the shape of the buried object and its electromagnetic properties. Unfortunately, the shape of the hyperbola also depends on elements completely independent of the object. In particular, the technical characteristics of the GPR (frequency, elevation) as well as the type of soil and the number of layers between the object and the ground have an enormous influence on this shape. The aim of this paper is to propose a high-performance and robust approach that will remain effective in as many experimental configurations as possible. It seems complicated to achieve this robustness by developing a shallow CNN model. One solution is to use a very deep network, as is the case in computer vision. But unfortunately, these networks require a large amount of training data to achieve good performance in terms of robustness. And it is well known that this large amount of data is not readily available for applications such as GPR.

One solution is then to change of features before the

classification step. Instead of the image of the hyperbola, a suitable transformation can achieve the performance and robustness objectives sought in this study. For example, it is possible to construct a covariance matrix from this image by using the method proposed in [33]. It is known that second order can improve classification performance, as for example in computer vision [34], [35]. Moreover, we have shown in [36] that this kind of feature brings robustness when a shift is present between the training and the test data. By using a similar approach of [33], we have shown in a previous work [37] that this operation achieves better performance with classical machine learning algorithms than using the raw image. This covariance matrix is constructed from the outputs of the first layers of a deep network. It measures the correlations between these different layers for a given image. To achieve a certain richness in this covariance matrix, it is often useful to use a certain number of layers (around 8-10). In this case, however, the covariance matrix is very large, making classification impossible caused by singularities issues. To solve this issue, a preliminary work [38] has proposed specific layers for Symmetric Positive Definite (SPD) matrices, properties of covariance matrices, to solve this problem of classifying from large covariance matrices. Similar models have also been proposed for EEG data analysis [39], classification of RADAR data [40] or classification in polarimetric SAR images [41].

In this paper, we propose a new classification model based on a transformation of the raw image into a covariance matrix and specific subsequent layers adapted to this SPD matrix. We also propose a different approach to that proposed in [33] for constructing our covariance matrix, which saves memory space while preserving correlation information. To train and to test our new model, we have from Geolithe, a sufficient database containing 4 types of buried objects with different GPR configurations (frequency and elevation) as well as several terrains (dry and wet sand and gravel). We compare our approach with shallow networks and conventional deep networks used in computer vision. We will show good performance and, above all, robustness of our pipeline to different experiments.

The outline of the paper is the following. First, section II introduces the GPR principle as well as some physical considerations allowing to better understand how the shape of the hyperbolas and the buried object are linked. Section III presents the new model to classify buried objects from GPR images. Next, section IV gives some details on the used database for the training and the steps. Finally, our approach is tested and compared to other algorithms in the section V.

II. GROUND PENETRATING RADAR (GPR)

A. GPR Principle

Ground Penetrating Radar (GPR) is a radar system consisting of an antenna that is typically placed on the surface of the ground¹. For a wide range of systems, the electromagnetic wave transmitted by the GPR is a simple wavelet, known in the community as a Ricker. An example of this type of signal is shown in figure 1. The frequency of the wave, from 10MHz

¹In our paper, we will also consider different elevations but the principle presented applies still.




Fig. 1: Waveform emitted by classical GPR, called Ricker.




Fig. 2: GPR Principle with acquisition (left), creation of the A-scan (middle) and the B-scan (right).

to 2GHz, depends on the application. In fact, this frequency is linked to the depth that can be reached. For example, for mine detection, a high frequency will be chosen, as there is no need to reach great depths, whereas the opposite will be chosen if you want to know the composition of the ground over several tens of meters.

The GPR is moved along an axis, as shown on the left of the figure 2. All acquisitions, i.e. the amplitude of the signal over time at a given point, are combined to form an image known as the A-scan. From this A-scan, it is then possible to construct an image of the ground called a B-scan or radargram. All processing, from detection to classification, are classically based on this image. If a buried object is present, it is then seen several times by the GPR, leading to a hyperbola in the B-scan image. Soils are often composed of several layers of different types. In this case, the B-scan shows some lines to represent these layers. On the right-hand side of the figure 2, the B-scan shows the image of a buried object and two layers. In this simple simulation, we do not take noise into account. In reality, the signal-to-noise ratio of GPR images is very low. In particular, there is a lot of clutter due to the Ricker reflecting off small scatterers, such as rocks.

In the next section, we will look at how object type and other parameters influence the shape of the hyperbola.

B. Influence of the physical parameters on the hyperbola shape

Several factors influence the shape of the hyperbola in the radargram. Firstly, GPR parameters, such as the frequency of the transmitted wave or the elevation of the system relative to the ground, strongly affect the resulting hyperbola. In this

article, we will consider several frequencies as well as different GPR elevations. This last point is useful in the case of an airborne GPR, which cannot then be used too close to the ground.

Another factor influencing the shape is obviously the soil and its composition. Each type of soil has its own dielectric permittivity and electrical conductivity, both of which influence the speed of the emitted wave. For example, higher dielectric permittivity generally slows down the propagation of electromagnetic waves or deflects radar waves more than soils with lower permittivity. Furthermore, interfaces between different layers also have an important influence, in part when the two layer parameters are very different from each other. In this case, the Ricker wavelet will undergo several deformations (attenuation, spectral broadening, etc.). For more details on these deformations, please refer to [1].

Finally, the shape of the hyperbola obviously depends on the buried object reflecting the transmitted wave back to the RADAR. Firstly, the size and shape of the object will influence its shape, particularly in the axis of motion of the RADAR. Finally, the electromagnetic properties of the buried object have an impact on the hyperbola, but more in the time axis. Conductive materials, such as metals, absorb more energy and attenuate signals faster than non-conductive materials. To help distinguish between metallic and non-metallic objects, we can also look at the polarization of the hyperbola. Polarity reversal occurs when radar waves reflect off objects whose permittivity is higher than that of the surrounding medium (soil has a lower permittivity than metallic objects). To define polarity, we define that the black areas of the hyperbola correspond to a negative polarity '-' and the lighter or white areas to a positive polarity '+'. So, depending on the polarity of the incident wave, a metal object can appear as a positive (+ - +) or negative (- + -) reflection.

We will show some examples of buried object images in the next section.

C. Examples

On the radargrams in Figure 3, we have used the 200 MHz GSSI antenna which have a positive polarity (+ - +) on wet sand. So a change in polarity will be detected as negative or reversed (- + -). It is therefore more appropriate to speak in terms of normal polarity when the reflection polarity is the same as the incident wave, or reversed polarity when these polarities are different.

On the radargram (a), visible in Figure 3, the signature of the shelter consists of:

- A wide hyperbola corresponding to the roof of the shelter (flat surface). The width of the flattened, high-intensity part corresponds to the width of the shelter (2 m). This reflection specifically corresponds to the soil-air interface (rather than soil-wood) due to the significant contrast in dielectric permittivity between air and soil and the thinness of the wood. The polarity is normal ((+ - +) here) same like GSI antenna : there is no change in polarity at the soil-air contact (decrease in dielectric permittivity).




Fig. 3: Examples of preprocessed GPR images with the 200 MHz antenna in wet sand (after direct wave suppression and histogram correction thanks to GPRpy which is an open-source Ground Penetrating Radar processing and visualization software available in <https://github.com/NSGeophysics/GPRpy>) for different buried objects: Wooden shelter, dummy shell (Metal) and wooden board coated with rubber (Non-Metal). Noticed that the 3 radargrams have different scales.

- Two adjacent hyperbolas beneath the first reflection. They correspond to reflections on the corners of the shelter. Their polarity is inverse.
- A wide hyperbola below the previous hyperbolas, corresponding to the base of the shelter.

On the radargram (b), we observe that the signature of a dummy shell gives

- hyperbolas which are easily recognized by their high intensity, well-defined hyperbolic shape and polarity opposite that of GSSI antennas (- + -). The polarity is reversed for metallic objects.

On the radargram (c), we observe that the signature of a wooden board coated with rubber gives a slightly flattened hyperboles identified by a low intensity. Moreover, the polarity are reversed (- + -) here.

D. Goal of the classification model

The 3 previous examples show that the shape of the hyperbola is linked to the buried object. In this article, we therefore propose to build a classification model that takes as input a thumbnail of each hyperbola. We assume that localization and classification have been carried out in the previous steps. Some pre-processing steps are given in [37].

One of the main concerns regarding the classification strategy based on GPR images is robustness. Indeed, we noted in the previous sections that the shape of the hyperbola also depends on parameters other than those of the buried object, such as the GPR and the soil. In this case, the classification must be invariant to these deformations of the hyperbola independently of the object.




Fig. 4: Architecture of model CNN1 [30]

The proposed strategy in the next section consists in imposing a transformation on the thumbnail image in order to classify it after this treatment. In particular, we study the interest of the second-order model for achieving good classification performance with good robustness, and all this with relatively little training data.

III. SECOND ORDER MODEL

In our task of classifying our radargrams, we rely on neural networks that have shown their effectiveness in image classification tasks.

Let us first describe some of the models that we used as a baseline and illustrate their respective limitations. From this we will then describe our proposed approach that combines several aspects of those models.

A. Shallow CNN network

As we are considering GPR images thumbnails of pre-localized hyperbolas, the first classes of models that are natural to consider are 2D Convolutional Neural Networks (CNNs) that have been successfully applied in computer vision tasks [42]. Rather than consider large models, and given that we consider very few classes compared to the general image classification problem, shallow CNNs have been considered in [30], where several of such small scale architectures are presented. From this study, we selected the best reported model in terms of overall accuracy that is denoted **S-CNN** (Shallow CNN) whose architecture is reported in Figure 4.

It is constructed as a succession of 3 embedding layers with help of standard 2D convolutions, batch normalization, ReLU non-linearity and max-pooling layers. Then a fully connected layer is followed by a softmax for classes probabilities. While this model allows for classification of the hyperbolas, the obtained accuracies we obtained in practice were not as satisfactory on our dataset², than in the one used in [30], who was trained on synthetic dataset and with a different number of classes. Since our database consists of a great number of real images more difficult to interpret than synthetic ones, it can be intuited that the size of this model is not sufficient for our task.

B. Computer-vision models

In order to circumvent the lower generalization capabilities of shallow networks, we consider models that are successful in computer vision tasks. In this work we focused on the ResNet

architecture [43]. When using models from the literature, there are two possible approaches:

- Using the pre-trained weights from another task, to benefit from the rich embedding representations and associated classification layers learned on a much bigger dataset. In this case, the weights can be fine-tuned by using the pre-trained values as initialization.
- Given a sufficient enough database size for the task, it is possible to train from scratch. This is useful in situations where the task is very different than traditional computer visions tasks like the GPR classification problem.

In the following, we denote the first model as **RFT** (ResNet Fine-Tuned) while the second is **RRT** (ResNet Re-Trained from scratch).

Such models are appropriate in handling various classification tasks. However, a problem lies in the very high number of parameters used for the task at hand, making the inference costly compared in terms of amount of labeled data to the previous shallow model. To address this issue, one can consider taking advantage of recent approaches based upon second-order statistics, which have shown promising results in computer-vision as well as in other applications [39]–[41].

C. SPD models

Covariance representations have been shown to be a relevant description when dealing with noisy signals coming from radar systems [44]. Notably, statistical hypothesis testing over second-order modeling have been successful in target detection in GPR [15]. For classification tasks, building upon [45], a preliminary study over a binary classification problem has been done in [46] using a covariance pooling approach. The idea is to take advantage of the standard convolutional embedding layers that are learned from a computer-vision dataset while employing second-order statistics as a mean to both reduce the dimension of the embedded feature space while also providing additional spatial invariance and better noise handling.

As described in [47], we propose to refine the standard convolutional integration layers by backpropagating the loss gradient onto the covariance estimation layers (input tensor construction and estimation) using the matrix backpropagation calculation given in [48]. This new strategy adopted in the paper should enable us to better adapt the various layers to GPR images, which are profoundly different from computer vision data.

D. Proposed models

Based on those previous works we propose two architectures names **SRCNet** (Stacked Residual Covariance Network) and **RCNet** (Residual Covariance Network) that take advantage of the covariance pooling approach while adding some layers from the recent model [38] that are specifically designed to handle covariance matrices. The idea behind those layers is to reduce the dimension of the covariance matrices by non-linear dimension reduction while preserving only the information that is relevant for the classification task. The various pipelines used in this paper are illustrated in Figure 5. Let's describe the different steps of the proposed model:

²as will be seen in section V.

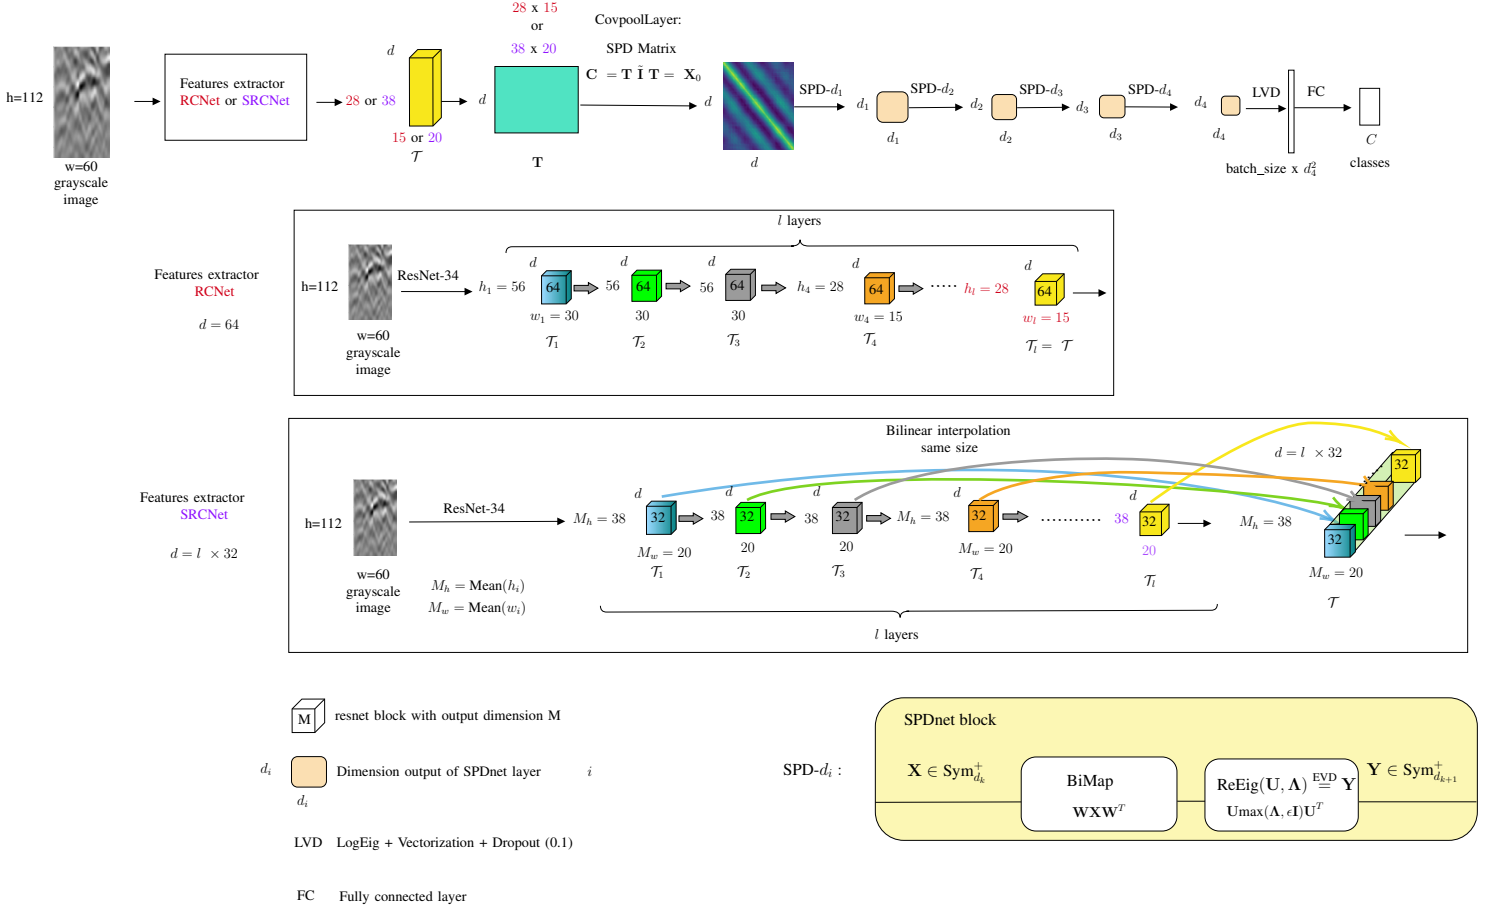


Fig. 5: Illustration of the two architectures used in this paper. In **RCNet** (Residual Covariance Network), we take only the last output of the ResNet blocks while in **SRCNet** (Stacked Residual Covariance Network), we stack the first 32 outputs (to save memory space) of the outputs features by interpolating them to a common size of 38×20 . With **RCNet** (Residual Covariance Network), we have $h_i = 54$ and $w_i = 30$ for the first 3 layers and $h_i = 28$ and $w_i = 15$ for the others.

Resnet34 with l layers: First, we use the first l layers of a non-pretrained ResNet-34 model³. In this paper, we use $l = 8$, with each layer producing 64 output filters. We start with a grayscale image $\mathbf{I} \in \mathbb{R}^{h \times w}$ whose initial image size $h = 112$ and $w = 60$ is taken from [30], to provide a basis for comparison in the remainder of this article. We have considered two scenarios (SRCNet and RCNet), in which we retain a portion of the output from each layer or only the last one.

SRCNet: Starting from our grayscale image, the steps are as follows:

- The first l layers of the ResNet34 model each contain $\bar{d} = 64$ filters.
- Only $\bar{d} = 32$ filters will be retained for each layer to save memory space.
- Therefore, the total number of filter outputs is given by $d = l \cdot \bar{d}$.
- Let $\mathbf{I}_i \in \mathbb{R}^{h_i \times w_i}$ with $i \in [1, d]$ represent the set of filter outputs.

³Preliminary results have shown that performance is better with a non-pretrained network than with one learned with computer vision data, surely because the first layers generally contain simple features.

⁴Note that each images has a different size for each i which explains the next step to resize all images with an identical size.

- Define M_h as the mean of h_i and M_w as the mean of w_i .
- We need to resize the filter outputs \mathbf{I}_i before stacking. For this purpose, we resize \mathbf{I}_i to $\tilde{\mathbf{I}}_i \in \mathbb{R}^{M_h \times M_w}$.
- We then stack the resized filter outputs $\tilde{\mathbf{I}}_i$ into a tensor \mathcal{T} . Thus, $\mathcal{T} = \{\tilde{\mathbf{I}}_i\}_{i \in [1, d]} \in \mathbb{R}^{d \times M_h \times M_w}$.
- Finally, we reshape \mathcal{T} into \mathbf{T} where $\mathbf{T} \in \mathbb{R}^{d \times M}$ with $M = M_h \times M_w$.

RCNet: In this alternative approach, the covariance will be calculated without stacking the output filters, but by taking the layers sequentially with all the output filters (and therefore not just 32 in this configuration). In this way, the covariance matrix will be estimated using only the characteristics of the last l layer. In this case, the total number of filters is $d = 64$, which leads to $\mathcal{T} = \{\tilde{\mathbf{I}}_i\}_{i \in [1, d]} \in \mathbb{R}^{d \times h_l \times w_l}$. Next, we transform \mathcal{T} into \mathbf{T} where $\mathbf{T} \in \mathbb{R}^{d \times M}$ with $M = h_l \times w_l$. The strategy of this second option aims to explore the potential benefits of residual connections and optimize performance in terms of computation time and memory space.

After these steps to build the tensor data \mathcal{T} by either **SRCNet** or **RCNet**, we have the following steps:

Covariance Pooling (CovPool) Layer: Then a covariance pooling layer is added. From \mathbf{T} of **SRCNet** or **RCNet** we

calculate⁵:

$$\mathbf{C} = \bar{\mathbf{T}}\bar{\mathbf{T}}^T,$$

where $\bar{\mathbf{I}} = \frac{1}{N}(\mathbf{I} - \mathbf{1}_N\mathbf{1}_N^T)$ and $\mathbf{C} \in \mathbb{R}^{d \times d}$ is a SPD Matrix (denoted SPD for Symmetric Positive Definite).

SPD Net Layer [38]: The obtained covariance matrix \mathbf{C} can then be of very high dimension, and it seems interesting to reduce the data space for better performance. Therefore, we suggest adding convolutional layers adapted to covariance matrices using the framework proposed in [38], denoted as SPD Net. SPD Net is a model that, for $k \geq 1$, takes as input a matrix $\mathbf{X}_{k-1} \in \text{Sym}_{d_{k-1}}^+$. In our case, we start with $\mathbf{X}_0 = \mathbf{C} \in \mathbb{R}^{d \times d}$, a symmetric positive definite matrix of size d_{k-1} . The dimensionality of the space is reduced through several BiMap convolution layers and a ReEig regularization layer based on an Eigenvalue Decomposition (EVD). Finally, there is a LogEig layer to perform measurements between covariance matrices in a common space (tangent plane to the Riemannian manifold of SPD matrices taken at the identity). The unknowns in the problem are the convolution matrices $\mathbf{W}_k \in \mathbb{R}_*^{d_k \times d_{k-1}}$, which are low-rank and belong to a Stiefel manifold.

Here are some details about the different layers at each step of the forward propagation:

- BiMap Layer (to generate more compact and discriminative SPD matrices):

$$\mathbf{X}_k = f_b^{(k)}(\mathbf{X}_{k-1}; \mathbf{W}_k) = \mathbf{W}_k \mathbf{X}_{k-1} \mathbf{W}_k^T,$$

where \mathbf{X}_{k-1} is the input SPD matrix of the k -th layer, $\mathbf{W}_k \in \mathbb{R}_*^{d_k \times d_{k-1}}$, ($d_k < d_{k-1}$) is the orthonormal transformation matrix (connection weights), $\mathbf{X}_k \in \mathbb{R}^{d_k \times d_k}$ is the resulting matrix and $f_b^{(k)}$ is the function for the k -th layer.

- ReEig Layer (to improve discriminative performance, inspired by ReLU):

$$\mathbf{X}_k = f_r^{(k)}(\mathbf{X}_{k-1}) = \mathbf{U}_{k-1} \max(\epsilon \mathbf{I}, \boldsymbol{\Sigma}_{k-1}) \mathbf{U}_{k-1}^T,$$

where eigenvalue decomposition (EIG) of $\mathbf{X}_{k-1} = \mathbf{U}_{k-1} \boldsymbol{\Sigma}_{k-1} \mathbf{U}_{k-1}^T$, ϵ is a rectification threshold, \mathbf{I} is an identity matrix and $\max(\epsilon \mathbf{I}, \boldsymbol{\Sigma}_{k-1})$ is a diagonal matrix of the corrected eigenvalues in order to stay on the SPD manifold.

Before the final steps, the BiMap and ReEig layers could be repeated several times in order to find the best representation for the classification.

In the final stage of the model, we aim to classify a lower-dimensional discriminant SPD (Symmetric Positive Definite) matrix. To enable the application of traditional fully connected (FC) layers, we first transform the SPD matrix into a feature in Euclidean space. This transformation is performed using the LogEig operator, which calculates the matrix logarithm of the input:

$$\mathbf{X}_k = f_l^{(k)}(\mathbf{X}_{k-1}) = \mathbf{U}_{k-1} \log(\boldsymbol{\Sigma}_{k-1}) \mathbf{U}_{k-1}^T,$$

⁵Actually this covariance matrix is the classical Sample Covariance Matrix (SCM) which assumes that the data distribution is Gaussian.




Fig. 6: Matrix backpropagation principle. Solid lines correspond to forward pass and dashed lines to backward pass. θ are the learned parameters: $\theta_l = \mathbf{W}_l$ for BiMap layers while $\theta = \emptyset$ otherwise. \mathbf{X} is the tensor after ResNet layers and \hat{y} is the classification output.

where $\log(\boldsymbol{\Sigma}_{k-1})$ is the diagonal matrix of logarithms of the eigenvalues of the SPD matrix, from the EVD of \mathbf{X}_{k-1} as with ReEig layer. Finally, we vectorize the resulting matrix before the FC layers and we introduce a dropout mechanism to mitigate the risk of overfitting.

Back-propagation Steps: Since the main steps are matrix operations, the backward is not classical like in most of deep learning models. In particular, the gradients of the operations based on SVD have been firstly derived in [49]. A more stable formula is given in [40] and will be used. An illustration of the backpropagation steps can be found in Figure 6. Calculations concerning \mathbf{W}_k differ, as they are based on Riemannian gradient descent on the Stiefel manifold, where \mathbf{W}_k represents an orthonormal matrix. Detailed explanations of this approach can be found in [38]. Let us summarize hereafter the steps of the backpropagation for the proposed model.

Given loss function $\mathcal{L} : \mathbb{R}^K, \mathbb{R}^K \rightarrow \mathbb{R}$, where K is the number of classes. By denoting f_l the l -th layer in the network, we can define the loss starting at this layer as

$$\mathcal{L}^l = \mathcal{L} \circ \underbrace{f_\ell \circ \dots \circ f_{l-1}}_{\mathcal{L}^{l-1}} \circ f_l,$$

where ℓ is the number of total layers in the network.

Given those definitions, the backpropagation steps are as follows:

- *CovPool Layer:* Let us consider this layer as some f_l in the network, meaning $\mathbf{C} = f_l(\mathbf{T}) = \mathbf{T}\bar{\mathbf{I}}\bar{\mathbf{T}}^T$ and suppose that we have already calculated $\frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{C}}$. Since there is no parameter to learn, the matrix chain-rule [49] on $\mathcal{L}^l(\mathbf{T}) = \mathcal{L}^{l+1} \circ f_l(\mathbf{T})$ yields:

$$\left\langle \frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{C}}, d\mathbf{C} \right\rangle = \left\langle \frac{\partial \mathcal{L}^l}{\partial \mathbf{T}}, d\mathbf{T} \right\rangle,$$

where $\langle \cdot, \cdot \rangle$ is the Frobenius inner product. We have that $d\mathbf{C} = (d\mathbf{T})\bar{\mathbf{I}}\mathbf{T}^T + \mathbf{T}\bar{\mathbf{I}}(d\mathbf{T})^T$ and using symmetries on $\bar{\mathbf{I}}$ yields:

$$\frac{\partial \mathcal{L}^l}{\partial \mathbf{T}} = 2\bar{\mathbf{I}}\mathbf{T}^T \frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{C}}.$$

- *BiMap Layer*: for this layer, we have two gradients to propagate : $\frac{\partial \mathcal{L}^l}{\partial \mathbf{X}_{l-1}}$, the gradient towards the input of the layer and $\frac{\partial \mathcal{L}^\dagger}{\partial \mathbf{W}}$, the gradient to update the weights \mathbf{W} of the bilinear mapping to be learned. We then use the same matrix back-propagation principles as those presented for the *CovPool* layer, but applied to $f_l : \mathcal{S}^{d_{l-1}} \times \mathcal{O}_{d_l \times d_{l-1}} \rightarrow \mathcal{S}^{d_l}$, where $\mathcal{O}_{m \times n}$ is the set of orthonormal matrices of size $m \times n$, and $f_l(\mathbf{X}_l) = \mathbf{W}\mathbf{X}_{l-1}\mathbf{W}^T$. The gradients are then given by:

$$\begin{aligned} \frac{\partial \mathcal{L}^l}{\partial \mathbf{X}_{l-1}} &= \mathbf{W}^T \frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{X}_l} \mathbf{W} \\ \frac{\partial \mathcal{L}^l}{\partial \mathbf{W}} &= 2 \frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{X}_l} \mathbf{W} \mathbf{X}_{l-1}. \end{aligned}$$

One additional thing to take into account is that \mathbf{W} has a special structure. Thus, while doing the gradient step, it is necessary to use a Riemannian retraction operator [50] to keep the weights on the Stiefel manifold.

- *ReEig and LogEig Layer*: Since both operations operate on eigenvalue decomposition, we can decompose f_l as $f_l = f'_l \circ \text{eig}$ where eig means doing the EVD of the input matrix and f'_l is the operation on the eigenvalues and reconstructing the matrix. The gradient towards input is given by [40]:

$$\begin{aligned} \frac{\partial \mathcal{L}^l}{\partial \mathbf{X}_{l-1}} &= 2\mathbf{U}_{l-1} \left(\mathbf{P}^T \odot \left(\mathbf{U}_{l-1}^T \frac{\partial \mathcal{L}^{l+1} \circ f'_l}{\partial \mathbf{U}_{l-1}} \right)_{sym} \right) \mathbf{U}_{l-1}^T \\ &+ \mathbf{U}_{l-1} \left(\frac{\partial \mathcal{L}^{l+1} \circ f'_l}{\partial \boldsymbol{\Sigma}_{l-1}} \right)_{diag} \mathbf{U}_{l-1}^T, \end{aligned}$$

where \mathbf{P} is a square matrix given by

$$\mathbf{P}(i, j) = \begin{cases} \frac{1}{\sigma_i - \sigma_j} & \text{if } i \neq j, \\ 0 & \text{otherwise} \end{cases},$$

and σ_i are the eigenvalues of \mathbf{X}_{l-1} .

For the ReEig layer, the sub-gradients are given by:

$$\begin{aligned} \frac{\partial \mathcal{L}^{l+1} \circ f'_l}{\partial \mathbf{U}_{l-1}} &= 2 \left(\frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{X}_l} \right)_{sym} \mathbf{U}_{l-1} \max(\epsilon \mathbf{I}, \boldsymbol{\Sigma}_{l-1}), \\ \frac{\partial \mathcal{L}^{l+1} \circ f'_l}{\partial \boldsymbol{\Sigma}_{l-1}} &= \mathbf{Q} \mathbf{U}_{l-1}^T \left(\frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{X}_l} \right)_{sym} \mathbf{U}_{l-1}, \end{aligned}$$

where \mathbf{Q} is a diagonal matrix with elements:

$$\mathbf{Q}(i, i) = \begin{cases} 1 & \text{if } \boldsymbol{\Sigma}_{l-1}(i, i) > \epsilon, \\ 0 & \text{otherwise.} \end{cases}$$

For the LogEig layer, the sub-gradients are given by:

$$\begin{aligned} \frac{\partial \mathcal{L}^{l+1} \circ f'_l}{\partial \mathbf{U}_{l-1}} &= 2 \left(\frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{X}_l} \right)_{sym} \mathbf{U}_{l-1} \log(\boldsymbol{\Sigma}_{l-1}), \\ \frac{\partial \mathcal{L}^{l+1} \circ f'_l}{\partial \boldsymbol{\Sigma}_{l-1}} &= \boldsymbol{\Sigma}_{l-1}^{-1} \mathbf{U}_{l-1}^T \left(\frac{\partial \mathcal{L}^{l+1}}{\partial \mathbf{X}_l} \right)_{sym} \mathbf{U}_{l-1}. \end{aligned}$$

Thanks to all these steps, there is a backpropagation for all the layers from the final fully-connected layers to the ResNet34 convolution layers. This means that contrarily to previous works [34], [38] we make advantage of both the bilinear mapping with learnable weights to obtain a discriminative SPD matrix for classification but also the best convolutions and thus embedding space for the task at hand rather than keeping pre-trained weights on another task.

Implementation: we provide a pyTorch implementation of those steps available at https://github.com/ammarmian/another_spdnet.

Computation complexity : Regarding the complexity of the proposed methods (RCNet and SRCNet), the main bottleneck is the computation of the SVD of the covariance input in the BiMap layers which is $O(m^3)$, m being the size of input⁶. Intuitively, this makes the RCNet method more computationally attractive than its SRCNet counterpart. With regards to classic methods, it is difficult to conclude on complexity since deriving time complexity is a difficult for deep learning models and highly dependent on the number of parameters as well as the set of hyperparameters used in training. Given the size of model **S-CNN** though, it is evident that this is the faster method. In order to give an idea of computational cost associated with our models, we consider hereafter a comparison on time of training and inference of the different models on the same hardware.

In figure 7, we show the average run time (training and validation steps) as a function of the ratio of the training dataset for the methods described above and two classical methods (SVM and Random Forests). Full details and the chosen parameters are given in the next section. As might be expected, the classic ML algorithms and the shallow network have the lowest computation times. On the other hand, the SRCNet computation time is clearly the highest. This approach, which consists in stacking all filter outputs, is not at all interesting in terms of computation time. On the other hand, RCNet's computation time is really interesting, since it is most of the time lower than that of ResNet34 -based deep approaches for most training data ratios.

⁶ $m = 64$ for RCNet and $m = 32 \times l$ for SRCNet at the first BiMap.




Fig. 7: Average runtime w.r.t. ratio of training data set over 100 different seeds. For each method, the line corresponds to the mean of accuracy over all the seeds, and the filled area corresponds to 5-th and 95-th quantiles.

Let us now present the experimental results obtained with the proposed models in the next sections.

IV. DATASET DESCRIPTION

The database provided by the Geolithe company include 699 medium-sized radargrams of $(R_x, R_y) = (4000, 800)$ pixels, along with all the necessary acquisition information such as radar frequency (200MHz or 350MHz for Geolithe), radar elevation (0cm, 25cm, 50cm, 75cm, 100cm, 150cm), and soil type (wet sand, dry sand, gravel, dry gravel). Each radargram is associated with a mask of the same size. Main preprocessing steps [37] consists to determine the rectangle around each hyperbola in the radargram by using the mask image in order to build a thumbnail image for each target. An example of this preprocessing for one radargram is shown in 8. All thumbnails are resized to $(w, h) = (60, 112)$ before to be treated by the different classification methods. We have also created a class denoted "empty" (in yellow in the figure 8). To create this class "empty", we selected 5 rectangles of same size for each radargram of our dataset. The centers of the rectangle are randomly chosen and we check if the rectangle does not overlap with another rectangle containing an object.

The final database then consists of 1584 thumbnails, classified into four categories: Metallic, Non-Metallic, Wooden Shelters, and Empty. Tables I and II provides more details on the distribution of the four categories according to elevation, soil and frequency. As used classically in supervised approaches, we divide our dataset into three distinct parts: a training set consisting of 1108 images, a validation set with 238 images, and finally, a test set also comprising 238 images.

The latter will be used to evaluate the accuracy of the different tested models.

V. NUMERICAL EXPERIMENTS

A. List of models

As described in section III, we consider in this paper three models with small variation which are recalled in the following:

- **S-CNN**: the shallow architecture developed in [30] for GPR image classification
- **Resnet34**: a deep network proposed in [43] initially applied to computer vision applications. In this section we will consider two models from this architecture:
 - Resnet Re-Trained (**RRT**): the model is then trained from scratch, which we initialize the weights randomly.
 - Resnet Fine-Tuned (**RFT**): the model is fine-tuned which consists in trained by using the pre-trained weights. In this specific model, they are pre-trained from the ImageNet database.
- Our proposed models in two configurations:
 - **SRCNet**: all the l first layers (only 32 output filters are selected for each layers to save memory space) are considered to build our tensor. In this case the first size of the SPD matrix is $d_0 = 256$.
 - **RCNet**: only the last layer with these 64 output filters are used to build the tensor. In this case the first size of the SPD matrix is $d_0 = 64$.

For these both models, the number of layers of Resnet34 to build the tensor \mathcal{T} is $l = 8$. Moreover, SPD Net is designed with 4 consecutive BiMap and ReEig layers where the sizes of each SPD matrices are specified in Table III. The choice of 4 layers is common when using SPD Net [38] and in particular it is shown in different experiments [51] that it is useless to take more than 4 consecutive BiMap and ReEig layers. We also considered models with less than 4 layers as well as other set of output size and settled on this choice by cross-validation with regards to performance.

All models are used with the same parameters each time. The chosen optimizer is SGD with a momentum of 0.9, a batch size of 8, and a learning rate of 0.007.

In the following subsections, we compare all these models with our database described in IV, first by measuring the influence of the number of training data on the classification results, then the robustness in the presence of mislabeled data in the training set, and finally the robustness to data shifts between the training set and the test set. In all experiments, 100 different random generator seeds are used to construct the partition between training and testing data sets. This allows to obtain more representative results irrespective of the initialization. To showcase results, we decide to show the 5-th, 50-th and 95-th quantiles of the performance metric.

B. Influence of the number of training data

First, we study the influence of the number of training data since we want to design classification models which can




Fig. 8: Example for one radargram (bottom) and the corresponding image of the masks. The positions of the rectangles detected on the masks were projected onto the radargrams. The yellow rectangles correspond to a thumbnail of the empty class.

| Soil/Elevation (cm) | Wooden Shelter | | | | Metallic | | | | Non Metallic | | | | Empty | | | | Total |
|---------------------|----------------|----|----|-----|----------|----|----|-----|--------------|----|----|-----|-------|----|----|-----|-------|
| | 25 | 50 | 75 | 100 | 25 | 50 | 75 | 100 | 25 | 50 | 75 | 100 | 25 | 50 | 75 | 100 | |
| grave | 20 | 16 | 12 | 4 | 13 | 17 | 6 | 11 | 16 | 14 | 8 | 4 | 20 | 12 | 18 | 12 | 203 |
| dry grave | 24 | 18 | 17 | 19 | 18 | 14 | 16 | 6 | 14 | 15 | 14 | 10 | 15 | 19 | 19 | 20 | 258 |
| sand | 37 | 44 | 48 | 54 | 48 | 40 | 55 | 50 | 40 | 44 | 44 | 48 | 45 | 50 | 45 | 43 | 735 |
| wet sand | 18 | 21 | 22 | 22 | 20 | 28 | 22 | 32 | 29 | 26 | 33 | 37 | 19 | 18 | 17 | 24 | 388 |
| Total | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 1584 |

TABLE I: Distribution of our database for each soil and depending on the elevation of GPR.

| Frequency/Elevation (cm) | Wooden Shelter | | | | Metallic | | | | Non Metallic | | | | Empty | | | | Total |
|--------------------------|----------------|----|----|-----|----------|----|----|-----|--------------|----|----|-----|-------|----|----|-----|-------|
| | 25 | 50 | 75 | 100 | 25 | 50 | 75 | 100 | 25 | 50 | 75 | 100 | 25 | 50 | 75 | 100 | |
| 200 MHz | 53 | 46 | 54 | 49 | 24 | 13 | 19 | 18 | 26 | 28 | 26 | 20 | 40 | 46 | 45 | 40 | 547 |
| 350 MHz | 46 | 53 | 45 | 50 | 75 | 86 | 80 | 81 | 73 | 71 | 73 | 79 | 59 | 53 | 54 | 59 | 1037 |
| Total | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 99 | 1584 |

TABLE II: Distribution of our database for each frequency and depending on the elevation of GPR.

| | d_0 | d_2 | d_4 | d_6 | d_8 |
|--------|-------|-------|-------|-------|-------|
| SRCNet | 256 | 235 | 217 | 179 | 128 |
| RCNet | 64 | 58 | 54 | 44 | 32 |

TABLE III: Output dimensions of SPDNet layers (BiMap and ReEig).

provide good performance with a limited training dataset. The results of test accuracy w.r.t training ratio are shown in Figure 9. We can observe that our proposed approaches, **SRCNet** and **RCNet** both outperform the classical approaches at any given training ratio and have lower variability over seeds. The **Resnet34** approaches perform better than [30] especially when the number of training data increases. The retraining approach, **RRT**, gives better results compared to the fine-

tuning approach, **RFT**, in particular when the training ratio becomes smaller.

In order to showcase the benefits of using CNNs compared to more classical machine learning methods, we also compare these results to Support Vector Machine (SVM) and Random Forests which are reported to be used in GPR applications [28], [52]. To that end, we used standard image processing pipeline : we vectorize the images and standardize to have zero mean and variance 1. Then Principal Component Analysis (PCA) is used to reduce dimensionality with a threshold of 95% of explained variance. Concerning the methods, we use:

- SVM classifier using a radial basis function (RBF) kernel whose hyperparameters are optimized through grid search.
- RF classifier : We used Gini index [53] for node splitting. Again the hyperparameters are tuned using a grid search




Fig. 9: Results of test accuracy w.r.t to training dataset percentage over 100 different seeds. For each method, the line corresponds to the mean of accuracy over all the seeds, and the filled area corresponds to 5-th and 95-th quantiles.

approach.

The results are presented in Figure 10 where they are compared to our best performing model RCNet. This comparison allows to show that deep learning models are better performing in our scenario.




Fig. 10: Comparison of results compared to classical machine learning models. Test accuracy w.r.t to training dataset percentage over 100 different seeds. For each method, the line corresponds to the mean of accuracy over all the seeds, and the filled area corresponds to 5-th and 95-th quantiles.

C. Robustness to mislabeled Data

The second simulation studies the effect to have mislabeled data in the training data set. Actually it can be difficult to correctly labeled the hyperbola in particular because of the low Signal to Noise Ratio (SNR) in radargram. For this simulation, we introduce a variable level of mislabeled data, 0 to 20% of error, in the training set. The results of test accuracy w.r.t mislabelling percentage are presented in Figure 11. We have comparable dynamic than the previous experiment with better performance and robustness of **RCNet** and **SRCNet**

(with a better result when considering only the last layer for the tensor construction). The case of [30] is interesting as it shows that shallow models performance decreases quickly as soon as there are a few mislabelled data. Like in the previous simulation, we observe that our approaches have the least variability over the seeds. In conclusion, this result is in line with the analysis made in [54], which also shows that covariance matrix utilization brings great robustness to mislabeled data in metric learning methods.




Fig. 11: Results of test accuracy w.r.t training dataset mislabelling percentage. For each method, the line corresponds to the mean of accuracy over all the seeds, and the filled area corresponds to 5-th and 95-th quantiles. The legends are the same as in Fig. 9.

D. Robustness to data shift

In the case of GPR data, data are collected in different configurations: different frequencies and elevations for RADAR, different soils and weather conditions. As it is difficult (if not impossible) to have labeled data for all these configurations, it is very important to build models with the property of being robust to possible changes between training and test data. These transformations are also known in many applications and are referred to as data shifts. To study the behavior of the models proposed in this article, we consider 4 scenarios:

- * Scenario A: the training and validation sets include images obtained using a RADAR located at an altitude of 75 cm or 100 cm, while the test set is made up of images acquired at an altitude of 50 cm.
- * Scenario B: the frequency is chosen at 200 MHz for the training and validation sets, while the test set uses only data obtained with a frequency of 350 MHz.
- * Scenario C: the training and validation sets are made up of data acquired in dry gravel, whereas the test set uses gravel.
- * Scenario D: same scenario as C, with wet sand for the training and validation sets, and dry sand for the test set.

The distribution and number of elements in the sets for each scenario are detailed in table IV.

1) *Scenario A*: In Figure 12 we represent five boxplots given the accuracy performances of the different models. We easily concluded that our models, in particular **RCNet** performance is almost unchanged compared to the classical case,

| | Train | Val | Test |
|-----------------------------|-------|-----|------|
| A: ELV 75,100 vs 50 | 634 | 79 | 79 |
| B: FRQ 350 vs 200 | 394 | 49 | 49 |
| C: GRD Dry Gravel vs Gravel | 394 | 49 | 49 |
| D: GRD Wet Sand vs Dry Sand | 672 | 84 | 84 |

TABLE IV: Distribution of the training, validation and test sets for the 4 scenario.

are particular robust to this transformation. In this case, main transformations are scaling but the shape of the hyperbola slightly changes. As expected, the shallow model is the less robust. We also noticed that the variability over the seeds is very small with the models built from covariance matrix. This result is also in line with those obtained in [36] where covariance matrix are used to classify crops in Satellite Image Times Series.




Fig. 12: Results of test accuracy results for scenario A (ELV 75,100 vs 50).

2) *Scenario B*: The box diagrams are now shown in figure 13. As in the previous scenario, the models we propose outperform the classical models. But in this case, the best algorithm is **SRCNet** and we notice a sharp degradation in the performance of all approaches. The number of transformations between the two sets is then too high. This is because the frequency is linked to the penetration of the wave into the ground, as well as to the possible resolution which results in very different radargrams.




Fig. 13: Results of test accuracy results for scenario B (FRQ 350 vs 200 394 49 49).

3) *Scenario C and D*: The boxplots for the scenario C are shown in 14 while those for the scenario D are given in 15. In both cases, the second order deep learning models give better results and lower variabilities. We can conclude that if we want to obtain robust performances, it is clear that our

approaches are better suited than shallow models or classical deep learning models as **Resnet34**.




Fig. 14: Results of test accuracy results for scenario C (GRD Dry Gravel vs Gravel).




Fig. 15: Results of test accuracy results for scenario D (GRD Wet Sand vs Dry Sand).

VI. CONCLUSION

In this paper, we have proposed a new deep learning model based on second-order moments to classify buried objects from the hyperbola thumbnails obtained with a classical GPR system. The proposed model is the concatenation of several models: the first is composed of the first layers of a classical CNN and is used to obtain a covariance matrix from the outputs of convolutional filters, while the second is composed of specific layers to classify SPD matrices. These models are tested on a database composed of several radargrams and compared with shallow models and conventional CNNs typically used in computer vision applications. Our approach gives better results, particularly when the number of training data decreases and in the presence of mislabeled data. We also illustrated the value of second-order deep learning models when training data and test sets are obtained from different weather modes or considerations.

Thanks to this last analysis, we believe that our algorithms will perform well for other GPR datasets. In addition, the proposed architecture will be interesting to test on other RADAR data classification applications (e.g. remote sensing), but also in medical applications (e.g. EEG). As in GPR, these fields of application use very noisy data and do not have a large number of labeled data.

The main extension of this work will be to develop a robust, high-performance model that combines the detection,

localization and classification stages. In particular, we want to integrate second-order layers into the classical fast-RCNN network [55]–[57]. Indeed, the first blocks of this network is very close to our models, since it constructs a data tensor from convolutional layers. As in the proposed work in the current paper, these new layers could make it possible to obtain a more robust and efficient network with a low number of labeled data.

REFERENCES

- [1] David J Daniels, *Ground Penetrating Radar*, IEE, 2004.
- [2] Harry M. Jol, Ed., *Ground Penetrating Radar: Theory and Applications*, Elsevier, 2009.
- [3] Andrea Benedetto, Fabio Tosti, Luca Bianchini Ciampoli, and Fabrizio D’Amico, “An overview of ground-penetrating radar signal processing techniques for road inspections,” *Signal Processing*, vol. 132, pp. 201–209, 2017.
- [4] A.C. Gurbuz, J.H. McClellan, and W.R. Scott, “A compressive sensing data acquisition and imaging method for stepped frequency GPRs,” *Signal Processing, IEEE Transactions on*, vol. 57, no. 7, pp. 2640–2650, July 2009.
- [5] Zhaofa Zeng, Jing Li, Ling Huang, Xuan Feng, and Fengshan Liu, “Improving target detection accuracy based on multipolarization MIMO GPR,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 53, no. 1, pp. 15–24, Jan 2015.
- [6] Bin Liu, Yuxiao Ren, Hanchi Liu, Hui Xu, Zhengfang Wang, Anthony G. Cohn, and Peng Jiang, “GPRInvNet: Deep learning-based ground-penetrating radar data inversion for tunnel linings,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 10, pp. 8305–8325, 2021.
- [7] Xingkun He, Can Wang, Rongyao Zheng, Zhibin Sun, and Xiwen Li, “Gpr image denoising with NSST-UNET and an improved BM3D,” *Digital Signal Processing*, vol. 123, pp. 103402, 2022.
- [8] Qiqi Dai, Yee Hui Lee, Hai-Han Sun, Genevieve Ow, Mohamed Lokman Mohd Yusof, and Abdulkadir C. Yucel, “Dmrf-UNET: A two-stage deep learning scheme for gpr data inversion under heterogeneous soil conditions,” *IEEE Transactions on Antennas and Propagation*, pp. 1–1, 2022.
- [9] Zhi-Kang Ni, Jun Pan, Cheng Shi, Shengbo Ye, Di Zhao, and Guangyou Fang, “DI-based clutter removal in migrated gpr data for detection of buried target,” *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022.
- [10] Paolo Bestagini, Federico Lombardi, Maurizio Lualdi, Francesco Picetti, and Stefano Tubaro, “Landmine detection using autoencoders on multipolarization gpr volumetric data,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 1, pp. 182–195, 2021.
- [11] Christian Maas and Jörg Schmalzl, “Using pattern recognition to automatically localize reflection hyperbolas in data from ground penetrating radar,” *Computers and Geosciences*, vol. 58, pp. 116–125, 2013.
- [12] H. Brunzell, “Detection of shallowly buried objects using impulse radar,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, no. 2, pp. 875–886, 1999.
- [13] A.M. Zoubir, I.J. Chant, C.L. Brown, B. Barkat, and C. Abeynayake, “Signal processing techniques for landmine detection using impulse ground penetrating radar,” *Sensors Journal, IEEE*, vol. 2, no. 1, pp. 41–51, Feb 2002.
- [14] K.C. Ho and P.D. Gader, “A linear prediction land mine detection algorithm for hand held ground penetrating radar,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 40, no. 6, pp. 1374–1384, 2002.
- [15] Q. Hoarau, G. Ginolhac, A. M. Atto, and J. M. Nicolas, “Robust adaptive detection of buried pipes using GPR,” *Signal Processing*, vol. 132, pp. 293–305, Mar. 2017.
- [16] W. Ng, T. Chan, H.C. So, and K.C. Ho, “Particle filtering based approach for landmine detection using Ground Penetrating Radar,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 46, no. 11, pp. 3739–3755, Nov 2008.
- [17] A. Manandhar, P.A. Torrione, L.M. Collins, and K.D. Morton, “Multiple-instance hidden markov model for gpr-based landmine detection,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 53, no. 4, pp. 1737–1745, April 2015.
- [18] V. Kovalenko, A.G. Yarovoy, and L.P. Ligthart, “A novel clutter suppression algorithm for landmine detection with GPR,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 45, no. 11, pp. 3740–3751, Oct 2007.
- [19] Jing Li, Cai Liu, Zhaofa Zeng, and Lingna Chen, “GPR signal denoising and target extraction with the CEEMD method,” *IEEE Geoscience and Remote Sensing Letters*, vol. 12, no. 8, pp. 1615–1619, 2015.
- [20] Jianping Wang, Pascal Aubry, and Alexander Yarovoy, “Efficient implementation of GPR data inversion in case of spatially varying antenna polarizations,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 4, pp. 2387–2396, 2018.
- [21] Guillaume Terrasse, Jean-Marie Nicolas, Emmanuel Trouvé, and Émeline Drouet, “Sparse decomposition of the GPR useful signal from hyperbola dictionary,” in *2016 24th European Signal Processing Conference (EUSIPCO)*, 2016, pp. 2400–2404.
- [22] Michele Ambrosanio and Vito Pascazio, “A compressive-sensing-based approach for the detection and characterization of buried objects,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 7, pp. 3387–3395, 2015.
- [23] Fabio Giovanneschi, Kumar Vijay Mishra, Maria Antonia Gonzalez-Huici, Yonina C. Eldar, and Joachim H. G. Ender, “Dictionary learning for adaptive GPR landmine classification,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 12, pp. 10036–10055, 2019.
- [24] Matthieu Gallet, Ammar Mian, Guillaume Ginolhac, Esa Ollila, and Nickolas Stelzenmuller, “New robust sparse convolutional coding inversion algorithm for ground penetrating radar images,” *IEEE Transactions on Geoscience and Remote Sensing*, 2023.
- [25] Wenbin Shao, Abdesselam Bouzerdoum, and Son Lam Phung, “Sparse representation of gpr traces with application to signal classification,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 7, pp. 3922–3930, 2013.
- [26] Fok Hing Chi Tivive, Abdesselam Bouzerdoum, and Canicious Abeynayake, “Gpr signal classification with low-rank and convolutional sparse coding representation,” in *2017 IEEE Radar Conference (Radar-Conf)*, 2017, pp. 1352–1356.
- [27] W. Shao, A. Bouzerdoum, S. L. Phung, L. Su, B. Indraratna, and C. Rujikiatkamjorn, “Automatic classification of GPR signals,” in *Proceedings of the XIII International Conference on Ground Penetrating Radar*, June 2010, pp. 1–6.
- [28] Haoqiu Zhou, Xuan Feng, Yan Zhang, Enehdelihai Nilot, Minghe Zhang, Zejun Dong, and Jiahui Qi, “Combination of Support Vector Machine and H-Alpha Decomposition for Subsurface Target Classification of GPR,” in *2018 17th International Conference on Ground Penetrating Radar (GPR)*, June 2018, pp. 1–4, ISSN: 2474-3844.
- [29] 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*, May 2015, vol. 9454, pp. 385–394, SPIE.
- [30] Maha Almamani, Dalei Wu, Yu Liang, Li Yang, Dryver Huston, and Tian Xia, “Classifying GPR Images Using Convolutional Neural Networks,” in *11th EAI International Conference on Mobile Multimedia Communications*, Jan. 2018.
- [31] Xisto L. Travassos, Sérgio L. Avila, and Nathan Ida, “Artificial Neural Networks and Machine Learning techniques applied to Ground Penetrating Radar: A review,” *Applied Computing and Informatics*, vol. 17, no. 2, pp. 296–308, Jan. 2020, Publisher: Emerald Publishing Limited.
- [32] Mostafa Elsaadouny, Jan Barowski, and Ilona Rolfes, “ConvNet Transfer Learning for GPR Images Classification,” in *2020 German Microwave Conference (GeMiC)*, Mar. 2020, pp. 21–24, ISSN: 2167-8022.
- [33] Sara Akodad, Lionel Bombrun, Junshi Xia, Yannick Berthoumieu, and Christian Germain, “Ensemble learning approaches based on covariance pooling of CNN features for high resolution remote sensing scene classification,” *Remote Sensing*, vol. 12, no. 20, pp. 3292, 2020.
- [34] Peihua Li, Jiangtao Xie, Qilong Wang, and Wangmeng Zuo, “Is second-order information helpful for large-scale visual recognition?,” in *Proceedings of the IEEE ICCV*, 2017, pp. 2070–2078.
- [35] Peihua Li, Jiangtao Xie, Qilong Wang, and Zilin Gao, “Towards faster training of global covariance pooling networks by iterative matrix square root normalization,” in *Proceedings of the IEEE CVPR*, 2018, pp. 947–955.
- [36] Antoine Collas, Arnaud Breloy, Chengfang Ren, Guillaume Ginolhac, and Jean-Philippe Ovarlez, “Riemannian optimization for non-centered mixture of scaled gaussian distributions,” *IEEE Transactions on Signal Processing*, vol. 71, pp. 2475–2490, 2023.
- [37] Matthieu Gallet, Ammar Mian, Guillaume Ginolhac, and Nickolas Stelzenmuller, “Classification of GPR signals via covariance pooling on CNN features within a riemannian framework,” in *IGARSS 2022*, IEEE, 2022, pp. 365–368.

- [38] Zhiwu Huang and Luc Van Gool, "A riemannian network for SPD matrix learning," in *Proceedings of the AAAI conference on artificial intelligence*, 2017, vol. 31.
- [39] Reinmar Kobler, Jun-ichiro Hirayama, Qibin Zhao, and Motoaki Kawanabe, "Spd domain-specific batch normalization to crack interpretable unsupervised domain adaptation in eeg," in *Neurips 2022*, 2022.
- [40] Daniel Brooks, Olivier Schwander, Frédéric Barbaresco, Jean-Yves Schneider, and Matthieu Cord, "Riemannian batch normalization for SPD neural networks," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [41] Junfei Shi, Wei Wang, Haiyan Jin, Mengmeng Nie, and Shanshan Ji, "Riemannian complex matrix convolution network for polar image classification," 2023.
- [42] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, 2012.
- [43] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [44] E. Ollila, D. E. Tyler, V. Koivunen, and H. V. Poor, "Complex elliptically symmetric distributions: Survey, new results and applications," *IEEE Transactions on Signal Processing*, vol. 60, no. 11, pp. 5597–5625, 2012.
- [45] Sara Akodad, Lionel Bombrun, Maria Puscasu, Junshi Xia, Christian Germain, and Yannick Berthoumiu, "Deep Ensemble Learning Model Based on Covariance Pooling of Multi-Layer CNN Features," in *2022 IEEE International Conference on Image Processing (ICIP)*, Bordeaux, France, Oct. 2022, pp. 1081–1085, IEEE.
- [46] Matthieu Gallet, Ammar Mian, Guillaume Ginolhac, and Nickolas Stelzenmuller, "Classification of GPR signals via covariance pooling on CNN features within a riemannian framework," in *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*, 2022, pp. 365–368.
- [47] Peihua Li, Jiangtao Xie, Qilong Wang, and Wangmeng Zuo, "Is second-order information helpful for large-scale visual recognition?," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2070–2078.
- [48] Catalin Ionescu, Orestis Vantzos, and Cristian Sminchisescu, "Matrix backpropagation for deep networks with structured layers," in *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 2965–2973.
- [49] Catalin Ionescu, Orestis Vantzos, and Cristian Sminchisescu, "Training deep networks with structured layers by matrix backpropagation," 2015.
- [50] Nicolas Boumal, *An introduction to optimization on smooth manifolds*, Cambridge University Press, 2023.
- [51] Rui Wang, Xiao-Jun Wu, Ziheng Chen, Tianyang Xu, and Josef Kittler, "DreamNet: A deep Riemannian manifold network for SPD matrix learning," in *Proceedings of the 16th Asian Conference on Computer Vision (ACCV2022)*, 2022.
- [52] Nairit Barkataki, Sharmistha Mazumdar, Rajdeep Talukdar, Priyanka Chakraborty, Banty Tiru, and Utpal Sarma, "Prediction of size of buried objects using ground penetrating radar and machine learning techniques," in *2020 International Conference on Computational Performance Evaluation (ComPE)*, 2020, pp. 781–785.
- [53] Leo Breiman, "Random forests," *Machine learning*, vol. 45, pp. 5–32, 2001.
- [54] Antoine Collas, Arnaud Breloy, Guillaume Ginolhac, Chengfang Ren, and Jean-Philippe Ovarlez, "Robust geometric metric learning," in *Proceedings of EUSIPCO 2022*, Aug. 2022.
- [55] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 580–587.
- [56] Shaoqing Ren, Kaiming He, Ross B Girshick, and Jian Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [57] Ross Girshick, "Fast R-CNN," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1440–1448.