Urban objects classification using Mueller matrix polarimetry and machine learning

IRENE ESTÉVEZ,[1,*] **FILIPE OLIVEIRA,**[2] **PEDRO BRAGA-FERNANDES,**[1,3] **MIGUEL OLIVEIRA,**[1] **LUÍS REBOUTA,**[1] **AND MIKHAIL I. VASILEVSKIY**[2,4]

[1] *Centre of Physics, Department of Physics, University of Minho, Azurém Campus, 4800-058 Guimarães, Portugal*
[2] *Centre of Physics, Department of Physics, University of Minho, Gualtar Campus, 4710-057 Braga, Portugal*
[3] *Bosch Car Multimedia SA, Rua Max Grundig 35, 4705-820 Braga, Portugal*
[4] *International Iberian Nanotechnology Laboratory, Avenida Mestre José Veiga, 4715-330 Braga, Portugal*
[*] *irene.estevez@fisica.uminho.pt*

Abstract: Detecting and recognizing different kinds of urban objects is an important problem, in particular, in autonomous driving. In this context, we studied the potential of Mueller matrix polarimetry for classifying a set of relevant real-world objects: vehicles, pedestrians, traffic signs, pavements, vegetation and tree trunks. We created a database with their experimental Mueller matrices measured at 1550 nm and trained two machine learning classifiers, support vector machine and artificial neural network, to classify new samples. The overall accuracy of over 95% achieved with this approach, with either models, reveals the potential of polarimetry, specially combined with other remote sensing techniques, to enhance object recognition.

1. Introduction

Fully autonomous vehicles (AVs) are an increasing reality already being tested on public roads [1,2]. Self-driving vehicles promise to be a disruptive innovation in transportation, as it is predicted that they will have a significant impact on road safety, on transport efficiency and on the environment, at both the individual and population level [1,3–7]. One of the main challenges in developing AVs is to make them aware of their surroundings to accurately detect and recognize different kinds of static or moving objects, such as pedestrians, vehicles, traffic signs and different kinds of obstacles. RADARs, LiDARs, SONARs, video cameras and road condition sensors are examples of exteroceptive sensors that are being developed to solve this problem [7–9]. These sensors generate a massive amount of data that needs to be processed and analyzed in real-time [9]. To this end, artificial intelligence is used to identify the most appropriate driving maneuver to be performed due to their adaptivity, reliability, flexibility and capability of making real-time decisions, conventionally made by humans [10–12].

The sensors referred above are based directly or indirectly on intensity measurements of radiation. However, there are other techniques that allow enhancing information to analyze a scene. In particular, light-matter interactions can modify the polarization of light and its characterization allows expanding the volume and nature of the acquired information. The polarization state of an optical beam interacting with matter can change by varying the amplitudes, the relative phase of its orthogonal field components, or by changing its degree of polarization. One approach used to analyze polarization changes produced by samples is based on the Stokes vector-Mueller matrix formalism. This formalism is the most appropriate representation of polarization when considering not fully polarized light [13]. The Stokes vector describes the state of polarization of any light beam, whereas the Mueller matrix (MM), consisting of 4×4 real elements, provides information related to the samples. The first element (m_{00}) is related to the

irradiance of the sample and the remaining 15 elements encode the polarimetric information, such as diattenuation, retardance and depolarization. Taking into account that different materials could result in different polarimetric responses, the measurement of the Stokes vector of scattered light or the Mueller matrix of samples could provide additional information useful to distinguish materials of relevant objects, which could be an advantage compared to conventional intensity data. There is a vast amount of literature available on polarimetry used in many research fields [14], such as biomedicine [15–17], materials characterization [18–20] and astronomy [21–23].

In the last decades, polarimetry has shown its potential in remote sensing applications [14,24,25]. More recently, due to the enhancement of information detected by this technique, with great potential for autonomous driving, there has been a growing interest in using artificial intelligence with polarimetric measurements for materials identification. Complete interpretation of polarimetric results and their relation to the physical properties of light-scattering objects requires a quantitative electromagnetic model for the latter, which is not feasible for the vast majority of real-world objects (such as, e.g. tree leaves). The use of phenomenological relations between the polarimetric response and the type of the object, which can be established using machine learning algorithms, is a solution for this practical problem.

Depending on the illumination used to characterize the scene polarimetrically, two approaches can be found in the literature. On the one hand, different research works have studied the benefits of active polarimetry (by controlling the state of polarization that illuminates the scene) to classify materials [26–28]. In addition, much work evaluating the potential of polarimetry and, more specifically, MM measurements of real-world materials have been conducted [29,30], enabling material recognition not only by their shape or color, but also by their structural information. On the other hand, by using passive polarization (i.e. polarization modifications of diffuse skylight in sunlit outdoor scenes), several studies have demonstrated that polarimetric images together with deep learning can improve object detection in road scenes [31–33], even in adverse weather conditions. Moreover, passive polarimetry can be combined with other optical techniques and deep learning to effectively detect and classify objects [34,35]. For instance, Brown et al. [36] developed a hybrid LiDAR and passive polarimetric imager combination for material classification. By using machine learning, they classified materials with great accuracy.

The objective of this work was to study the feasibility of Mueller matrix measurements for classifying real-world urban objects, which can be of interest for polarimetric sensors in autonomous driving. In this context, a comprehensive study of the Mueller matrix of several materials for several rotation angles is presented. Then, the performance of two classification models (Support Vector Machine (SVM) and Artificial Neural Network (ANN)), developed to classify the measured materials, is discussed in order to analyze the viability and potential of polarimetry and machine learning to discriminate between real-world urban objects.

The article is organized as follows. In the next section, we provide the methodology, a description of the samples used to develop this work, and a brief review of the studied classification algorithms and the data preparation. Section 3 presents the data analysis of the studied real-world samples based on their Mueller matrix results and their use for classifying them into predefined classes by using SVM and ANN models. Section 4 is devoted to conclusions.

2. Methods

2.1. Mueller matrix imaging polarimeter

A complete Mueller matrix imaging polarimeter, operating in reflection, has been developed for this study at the University of Minho. Considering that the main goal of this work is to demonstrate the potential of Mueller matrix for classifying relevant urban objects in order to provide additional environmental information for AVs, the choice of the wavelength is important. The wavelength used in our MM polarimeter was chosen with the intention of reducing the impact of solar background radiation and atmospheric attenuation [37], as well as complying with

human eye safety. Note that, independently of the application, the highest average optical power of the light source used to illuminate a scene where people can be present will always be limited by the eye safety norms [38]. In this sense, light with wavelengths from 0.4 to 1.4 µm is easily transmitted through the cornea and can be focused on the retina [38,39]. Therefore, we used a laser with a wavelength of 1550 nm, because the eye safety regulations (EN 60825-1 standard) permit a higher laser output power and the solar background radiation is reduced compared to shorter wavelengths [39].

A comprehensive scheme of the MM imaging polarimeter is shown in Fig. 1. Light from the laser (Thorlabs, Model LDM1550) is expanded by the lens system L1. Then, the beam passes through a linear polarizer, LP1, with its transmission axis at 0° with respect to the laboratory vertical and through two liquid crystal variable retarders, LC1 and LC2, from Meadowlark, with their fast axes at 45° and 0° to the vertical, respectively. The combination of these polarization optics forms the polarization state generator, PSG. The generated polarization states are controlled by applying voltages to both LCs, since their retardances depend on the applied voltage [40]. The generated polarized light illuminates the sample at a sample rotation angle, θ, that is controlled by a motorized rotation stage. The samples were rotated to analyze the effect of the materials orientation in the measured Mueller matrix and in the classification performance. The sample rotation angle is measured with respect to the specular reflection condition, denoted by $\theta = 0°$. Note that, when $\theta \neq 0°$, off-specular reflection is being analyzed. Light reflected or scattered by the sample passes through the polarization state analyzer, PSA, consisting of the same elements of the PSG but in the inverse order. Finally, a camera lens, L2, images the sample on an InGaAs SWIR camera (Allied Vision, with a 640(H) × 512(V) resolution, a cell size of 15 µm × 15 µm and temperature stabilization via TEC features, Model GoldEye G-033 TEC1-C).
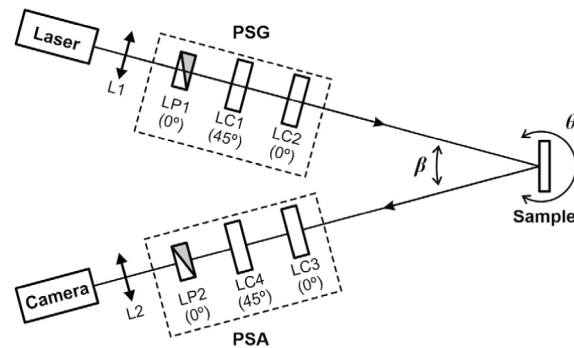

Fig. 1. Scheme of the Mueller matrix imaging polarimeter. The polarization state generator (PSG) consists of a linear polarizer (LP1) with its transmission axis at 0° and two liquid crystal variable retarders (LC1 and LC2) with their fast axis oriented at +45° and 0°, respectively. The polarization state analyzer (PSA) is reciprocal to the PSG. β is the bistatic angle (~9°) between the illumination and detection arms and θ is the sample rotation angle, measured with respect to the specular reflection condition. The picture corresponds to . The lens system L1 expands the laser beam to illuminate a larger area of the samples and the lens system L2 images the sample on the camera.

During the experiment, the detection arm was adjusted to be as close as possible to the illumination arm, in order to mimic backscattering remote sensing studies. The angle between them (β in Fig. 1) was fixed at ~9°. Thirty-six intensity images were acquired for a certain combination of voltages applied to the LCs of the PSG and PSA. Then, the experimental Mueller matrix image was calculated on a pixel-by-pixel basis from the polarimetric data reduction equation described in Ref. [13]. Obtained experimental Mueller matrix images correspond

to an illuminated area on the samples of 40×40 mm². The calibration of the Mueller matrix polarimeter is discussed in Section 1 of Supplement 1.

2.2. Samples description

For this study, the experimental Mueller matrices of several urban objects were measured for a set of angles θ from −50° to 50° by using the MM imaging polarimeter described above. In this context, it is important to note that the bistatic angle β was fixed (Fig. 1), which implies that, except for $\theta = 0°$ (specular reflection), diffuse light was measured.

The studied samples were chosen to be representative of real-world objects that can be found in urban environments. With this in mind, we decided to group real-world objects into six classes, which we consider relevant for being recognized by AVs and for demonstrating the potential of Mueller matrix assisted by artificial intelligence models. Note that, in urban environments, many materials can be typically associated with a certain class (for example, fabrics are typically associated with clothes used by pedestrians). Thus, we carried out measurements and posterior classification of samples from the classes listed in Table 1, with their brief description provided in the second column.

Table 1. Classes of urban objects and associated samples selected to be measured in this study.

Urban Objects	Description of the measured samples
Vehicles	Car paints (with metallic flakes and solid paints)
Pedestrians	Clothes (cotton, polyester, viscose, nylon, wool, . . .)
Pavements	Granite cobblestones
Traffic Signs	Danger warning, mandatory, information, priority traffic signs, . . .
Vegetation	Leaves from: trees (cherry (*Prunus avium*), oak (*Quercus faginea*), *Platanus occidentalis*, ...) and grass (*Eruca vesicaria, Amaranthus blitum, Plantago major*, . . .)
Tree trunks	Tree logs from: pine (*Pinus pinaster* Aiton), eucalipt (*Eucalyptus globulus*), olive tree (*Olea europaea*), loquat (*Eriobotrya japonica*), cherry tree (*Prunus avium*), walnut (Juglans regia), . . .

For the vehicles' class, we studied automotive paint samples which correspond to real paints of diverse car brands, such as Opel, Mercedes, Fiat, Ford, among others. Two kinds of car paints were characterized, solid paints (no sparkle effects) and metallic paints (containing metallic flakes). Moreover, in order to study the behavior of the metallic automotive paints, several samples containing different flake sizes were analyzed. Apart from that, clothes of diverse textiles or fabrics, compositions (e.g. cotton, nylon, wool, polyester, among others) and types, as well as of different seasons and genders were measured. For the pavements' class, granite cobblestones were studied because they are commonly used in low-speed zones in Portugal. Furthermore, different types of traffic signs (such as danger warning, mandatory, information or priority signs) with several inks and sheeting were analyzed. Both the front side and back side of the traffic signs were characterized. Finally, diverse types of vegetation (tree, grass and bush leaves) and tree logs from local gardens and fields were characterized. Some of the studied samples were provided by municipal services in Braga and *"BOSCH Car Multimedia"*. See Table S1 in Supplement 1 for more details on the studied urban objects. Some examples of samples of each one of the six studied classes are shown in Fig. 2.

2.3. Classification algorithms

Artificial intelligence is particularly suitable for recognition applications because of its capacity of analyzing large amounts of data and detecting patterns which can be used, for example, to make predictions. Based on the previous knowledge of data and the final application, there are different learning forms that can be used [41]. In our case, we are interested in supervised
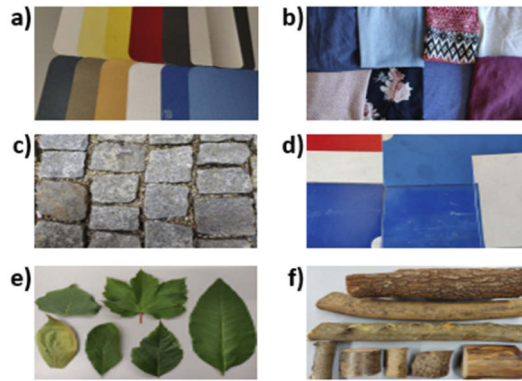
Fig. 2. Images of a set of measured samples of a) metallized and solid car paints, b) clothes, c) granite cobblestones, d) traffic signs, e) leaves and f) tree logs.

learning methods since not only the inputs (or features) are known (MM and sample rotation angle, θ), but also the outputs (the object classes described in Table 1). In particular, we focus on multiclass methods because they are suitable for determining whether a variable belongs to a class when considering more than two possibilities [42].

The classification process consists of three parts. First, the database containing our six classes is prepared by checking for anomalies or missing data, cleaning data, etc. In this part, the best set of inputs are determined. Secondly, the classifier is trained with the training data, tuning the model parameters for the database. This process is repeated until the classification error on the training data is small enough to make accurate predictions. Finally, the performance of the classifier is estimated by processing new data (not present in the training database) and classifying the new samples into one of the previously defined classes. In this work, we created a database with several predictors/outputs (the object classes presented in Table 1) and features/inputs (15 normalized MM elements, all except the normalizing element m_{00}, and the rotation angle of the sample, θ).

For this study, two types of classifiers were trained and tested, namely, the SVM and ANN. These two models were selected because they have proven to be robust and effective in similar classification problems [27,28,36,42].

2.3.1.   Support Vector Machine (SVM)

The objective of SVM algorithm is to find the optimal hyperplane that accurately separates data points belonging to different classes (see Fig. 3). In this sense, multiple hyperplanes can separate data points so, during the training process, the maximum margin hyperplane is selected by using support vectors (i.e., the small subset of the training data that lie the closest to the hyperplane) [41,43]. That is to say, the algorithm looks for the hyperplane that maximizes the perpendicular distance, called margin, between the support vectors and the decision boundary (as shown in Fig. 3). When data points are not separable, they are projected into a higher dimension by using mathematical equations, called kernel functions, transforming original data into a new feature space [43]. Two of the most used kernels are the polynomial and the Gaussian ones (the latter is also called Gaussian Radial Basis Function, RBF).

SVM was conceived to classify two classes (binary classification), however this method can be extended to perform multiclass classifications [44]. Several methods can be employed to use binary classification algorithms for multiclass classification problems. In this work, we have employed the one-vs-one and one-vs-all methods. On the one hand, the one-vs-one (also known as all-pairs) approach compares all pairs of classes to each other by constructing a binary
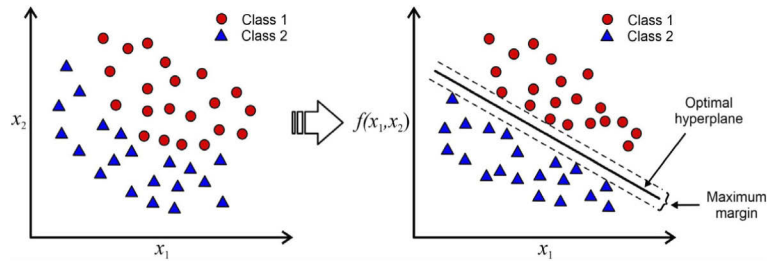
Fig. 3. Illustration of SVM algorithms for classifying 2 classes, red circles and blue triangles. The inputs, x1 and x2, are represented in the axes. The training data (left side) is transformed by a kernel, $f(\text{x1, x2})$ to a new plane, where the optimal hyperplane is obtained.

training sequence. With this method, the problem is divided into $N(N-1)/2$ binary classification models, with N being the number of classes [41,44]. For new classifications, each $N(N-1)/2$ model is used to predict new data. The selected class is the one that has the highest number of predictions. One of the main advantages of the one-vs-one method is that new classes can be added to the existing set of classes without the need to fully perform new computation. This can be useful when working with a significant number of classes. However, the number of binary classifiers will increase significantly as a function of the number of classes. On the other hand, the one-vs-all (also known as one-versus-rest) method creates a binary classification model for each class, discriminating each class from the rest of classes [41]. That is, this method is based on maximizing the margin from each class to the remaining classes. By contrast to one-vs-one, N binary classification models are trained in the one-vs-all method, requiring less training time. More details on the SVM algorithms are given in Refs. [44,45].

2.3.2. Artificial Neural Network (ANN)

An alternative to SVMs are the artificial neural networks. ANNs computation is performed by a mesh of computing nodes/neurons and connections/edges, that operate collectively and simultaneously on the data and inputs [41]. Each node has a defined number n of inputs, x_n, their associated weights, w_n, (strength of a connection) and the bias initialization, w_0, that are summed and evaluated by an activation function, g, to derive the output. These weights can be modified until the result of the activation function is above a threshold, i. e., a desired output is produced. Note that, in contrast to SVM, where the previously defined basis functions (kernels) are centered on the training data, ANNs allow the used basis function (activation function) to be adapted during training [44]. In the case of a single node, the output y is defined by [44]:

$$y = g\left(\sum_{i=0}^{n} w_i x_i\right), \tag{1}$$

where x_0 has the constant value 1. Common activation functions (g) are the logistic (sigmoid), the rectified linear unit (ReLU) and the hyperbolic tangent functions which are chosen depending on the nature of the training data. Nodes are grouped in layers in a way that an ANN can have multiple layers. A simplified diagram of an ANN is shown in Fig. 4. Typically, neural networks are formed by an input layer (including the inputs), a few hidden layers (in our diagram, two hidden layers are shown) and an output layer (containing the output signal).

The ANN is trained iteratively until the classification loss (prediction error) is minimized. For this, a loss function (also known as cost function) is used to compute this error and quantify the model performance. The most popular choice of the loss function is the Mean Squared Error (MSE). Since in multilayer networks the error may contain several local minima, several trainings
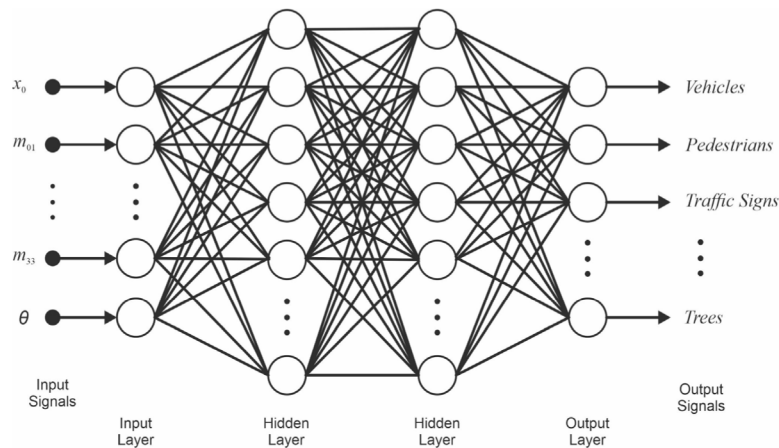
Fig. 4. Diagram of a simplified feed-forward Artificial Neural Network with 2 hidden layers. The input, hidden, and output variables are represented by nodes and the weights are represented by the lines connecting nodes. The inputs are the Mueller matrix elements (m_{01}, m_{02}, ... m_{33}), normalized by the m_{00} element, and the sample rotation angle (θ), and the outputs correspond to the classes described in Table 1.

should be performed in order to obtain the best classifier. When compared with SVM, ANN is more suitable for large datasets, but it has the disadvantage of higher computing cost during training.

2.3.3. Data preprocessing and feature analysis

ANN and SVM models for classification work by determining similarities and differences in the training data, requiring large amounts of data to successfully train the classifiers. The quality, amount and format of the training data as well as its information are important to lead a reliable classifier. For this reason, the database must be prepared in order to improve the quality of the models. Thus, several data preprocessing techniques were performed for our database.

First, the data was cleaned by filtering the experimental measurements to obtain physically realizable MMs [46,47]. Then, the training data was balanced and standardized, to give the same weight to all features (MM, normalized by the m_{00} element, and sample rotation angle, θ). During the training process, one class could be considered less important if it is imbalanced (less represented compared to the other ones). Some classification algorithms, such as ANN, are very sensitive to unequal distributions of classes in the training data. Therefore, our training data contain the same number of measurements per object class. Note that a balanced classifier is a useful solution inside the laboratory to verify the accuracy of individual object class recognition, but for a real-world sensor application, a weighted classification could be a better solution, since the proportion of objects and their importance in an urban environment are not the same.

Moreover, features measured at different scales do not contribute equally to the model, giving more weight to the larger range inputs. Our training data was transformed to improve the performance during the training process and the stability of the model [41]. Typically, data is transformed when the original inputs have significantly different variability. In this study, we have standardized our data in a way that the units of the features are comparable. The data, x_i, is centered around the mean, μ, of each feature and scaled by their standard deviation, σ. The transformation is described as [41]:

$$z_i = \frac{x_i - \mu}{\sigma}. \tag{2}$$

Furthermore, in order to choose the most relevant features for the training process, it is important to analyze the quality of the inputs. In this sense, by selecting only the most important features, the classification accuracy and predicting performance could be improved and effects from noise or irrelevant variables and computation cost could be reduced. The analysis of the features' importance of our training data was made by applying the Boruta algorithm [48,49]. This algorithm is capable of capturing non-linear relationships in the data by running a random forest classifier on an extended information system (created by adding shuffled copies to all features, known as shadow features, and comparing the importance of real inputs with those random shadow features). As a result, the Boruta algorithm quantifies the importance of each feature.

2.3.4.   Multiclass evaluation metrics

Different metrics were calculated to evaluate the performance of the trained ANN and SVM models. In particular, by using additional measurements (different from the training data), the classification performance of the models was analyzed by calculating the overall accuracy and the confusion matrix. The overall classification accuracy is defined as the proportion of correctly classified samples out of all the evaluated samples. The confusion matrix is a table that allows the visualization of the performance of a classifier, summarizing the prediction results of the model (correct and incorrect predictions). The confusion matrices presented in Section 3.2 show the true (in the rows) and the predicted (in the columns) classes of the data. The correct classifications are given by the diagonal elements of these matrices, while the non-diagonal ones are the incorrect predictions (misclassifications).

3.   Results

3.1.   Measurement results

The Mueller matrix images of ~50 samples per urban object class (described in Table 1) were measured with the MM imaging polarimeter described in Section 2.1. In order to consider real situations where the illumination angle could vary as a function of the object orientation, we characterized the samples at multiple rotation angles from −50° to 50°. The experimental images were filtered to obtain physically realizable MMs [46,47]. Some images of the experimental MMs of the studied urban objects are presented on Figs. S3 to S10 in Section 3 of Supplement 1. The obtained results show significant pixel-to-pixel variations due to spatial inhomogeneities and roughness of the inspected samples (for example irregularities in tree barks and granite pavements, or scratches in traffic signs). To mitigate these effects, the 16 elements of each MM image were averaged across the same illuminated area, obtaining several averaged Mueller matrices per measurement. The created database was divided into two parts, a set of balanced data, which was used for training the classification models (training data: 4000 measurements per object class), and another set to be used during the validation process (test data: 1000 measurements per object class). In order to test the classifiers performance, training and test data contains Mueller matrices of different measured samples. That is to say, a new set of samples (not previously analyzed) was measured for creating the test dataset, such as clothes made with other fabrics and colors, leaves from other trees and grass, and car paint of other car brands.

First, training data was visualized in order to observe the differences between the studied classes. Figure 5 represents the relationship between some averaged Mueller matrix elements, normalized by the element m_{00}, of each set of urban objects and sample rotation angles. In Fig. 5(a), m_{11} as a function of the rotation angle, θ, for each object class is shown. The normalized MM elements m_{11} as a function of m_{22} and m_{10} as a function of m_{33} are represented in Fig. 5(b) and (c), respectively. For both representations, the rotation angle θ varies from −50° to 50°.
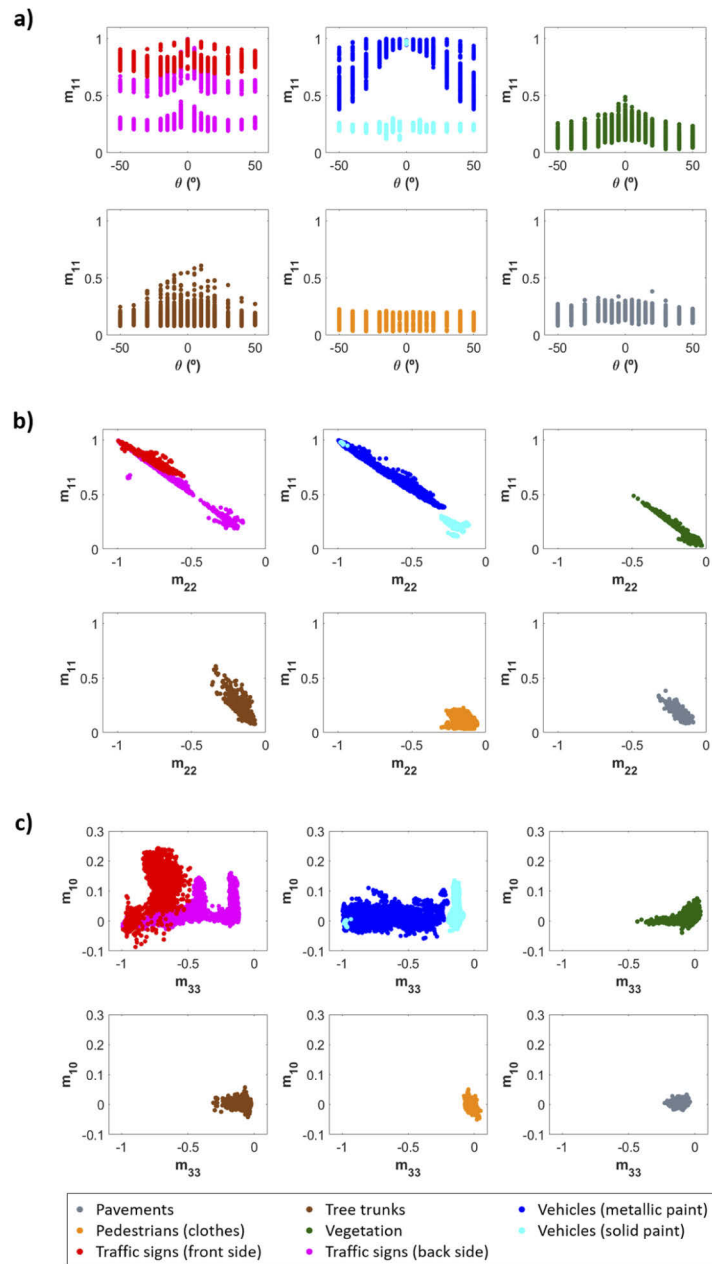
Fig. 5. Representation of various Mueller matrix elements for the measured samples: a) m_{11} as a function of the sample rotation angle (θ), b) m_{11} as a function of m_{22}, and c) m_{10} as a function of m_{33} for six urban object classes: traffic signs (front side (red) and back side (magenta)), vehicles (metallic paint (dark blue) and solid paint (cyan)), vegetation (dark green), tree trunks (brown), clothes of pedestrians (orange) and pavements (grey).

As expected, we observed that the Mueller matrix depends on the rotation angle of the measured sample. In addition, it was noticed that some samples considered of the same real-world object class, but made with different materials, result in different Mueller matrices. For instance, the vehicle class has distinct polarimetric behaviors for metallic paints, with different flake sizes and

orientations (dark blue dots in Fig. 5), and for non-metallic or solid paints (cyan dots in Fig. 5). These two different behaviors for car paints are related with the Umov Effect [50] and agree with the study of linear polarization carried out in Ref. [51]. Different responses can also be observed with the measured parts of the front side and back side of traffic signs (red and magenta dots in Fig. 5, respectively). In particular, the back side of the traffic signs (magenta dots in Fig. 5) shows two behaviors corresponding to the kind of paint used to protect them. Hence, we considered creating subclasses for the vehicle and traffic sign classes. Note that separating classes into subclasses can reduce the complexity of the input data, improving classifiers performance. The criteria for creating subclasses were the materials from which the surface of the urban objects is made and the advantage of recognizing them. By taking into account the composition of the studied automotive paints, we divided the vehicles class into two subclasses named metallic and solid car paints. Nevertheless, for the traffic signs class, we created just two subclasses: the front side and back side of the traffic signs, because only these two subclasses are of interest for autonomous driving.

We have observed in Fig. 5 that the solid paint subclass of vehicles and the back side of the traffic signs have great differences between specular ($\theta = 0°$) and diffuse reflections. Moreover, samples with high roughness, such as clothes of pedestrians and pavements, exhibit significant depolarization and small variations with the rotation angle. As a consequence, their MM elements are distributed in specific small regions (see Fig. 5(b) and (c)). From the obtained results and because different materials have different polarimetric responses, we conclude that some Mueller matrix elements are essential for identifying these predefined classes.

By visualizing our training data, we observed that our inputs have significant scale differences (e.g. m_{10} ranges from -0.087 to 0.243, while the rotation angle θ ranges from $-50°$ to $50°$). Therefore, our database was standardized by using Eq. (2). Then, we studied the importance of our sixteen standardized features to train our models, looking for those features that are uncorrelated and non-redundant. For that the Boruta algorithm [48,49] was applied. Figure 6 shows the importance among our inputs, where green box-plots represent Z scores of confirmed important attributes. The obtained results allow us to infer that although the variations between different studied classes for some normalized MM elements are small, they carry significant amounts of information and could be used to classify new samples into our predefined classes. As expected, the Boruta algorithm classified the normalized MM element m_{00} as not important to train our model, since its value is always 1.

The elements m_{33}, m_{10} (represented in Fig. 5(c)) and m_{01} show the greater importance. In the case of the rotation angle, the Boruta algorithm indicates that it is less important than the MM elements (see Fig. 6). That is because the rotation angle information is parametrically present in the Mueller matrices of the studied samples and implicitly determines the MM-rotation angle correlation. This result is of interest for some remote sensing applications, since to measure the rotation angle of an object could be impractical.

3.2. Classifiers training and implementation

SVM results

Several SVMs models were trained with different kernels and multiclass methods using 5-fold cross-validation. The best performance for the training data was obtained for the RBF kernel function and for the one-vs-all multiclass classification method, when considering as features the whole normalized MM and the rotation angle θ, all of them standardized, and as outputs the classes and subclasses described above. The overall classification accuracy obtained for the test data was 95.25%. Figure 7 presents the normalized confusion matrix of the test data, which shows the excellent performance of this trained model for classifying real-world objects, even for materials with similar polarimetric properties, such as clothes, vegetation, tree trunks and pavements (see Fig. 5). In fact, the most relevant object classes for autonomous driving (the two
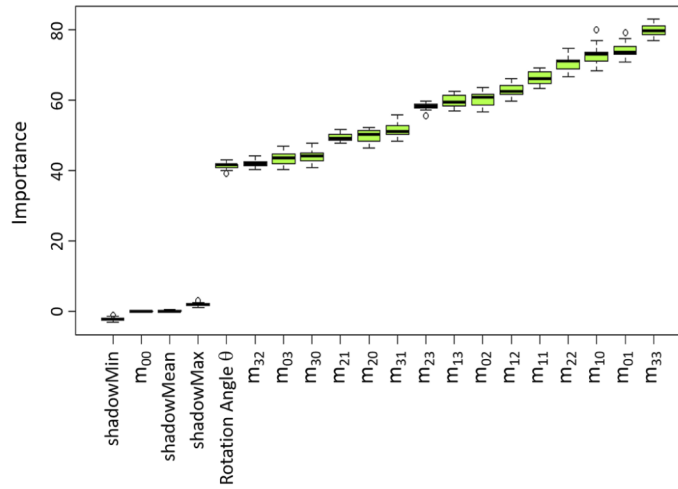
Fig. 6. Boruta algorithm results for our training data. Green box-plots represent Z scores of confirmed important attributes and the blue box-plots correspond to the shadow attributes (shadowMin, shadowMean and shadowMax corresponding to minimal, average and maximum Z score of a shadow attribute [48,49]).

types of car paints, the clothes of the pedestrians and the traffic signs) are classified above 95%. However, there is a misclassification between some urban objects. In particular, 9% of the actual vegetation is misclassified as pavements. From the point of view of the final application, this misclassification error is not as critical as misclassifying clothes of pedestrians, car paints of vehicles or traffic signs. Despite this, there is a good agreement between actual and predicted classes and subclasses.
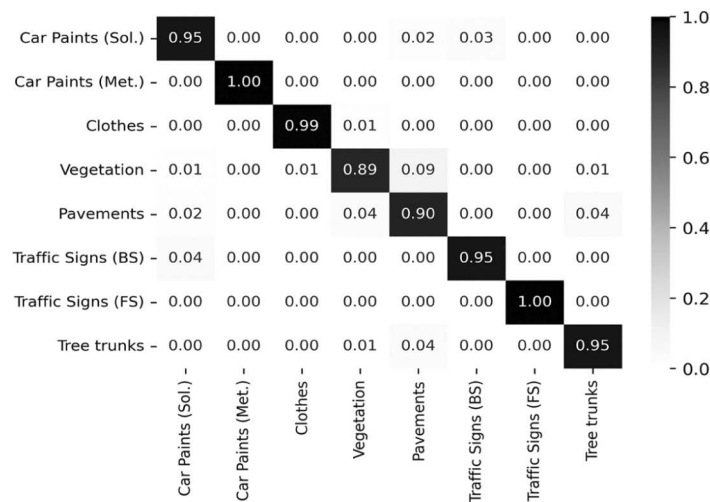

Fig. 7. Normalized confusion matrix showing the performance of our SVM classifier for the previously discussed urban object classes and subclasses. The selected features were: the normalized Mueller matrix and the rotation angle, θ, both of them standardized. The subclasses, solid (Sol.) and metallic (Met.) car paints, and front side (FS) and back side (BS) of traffic signs are presented separately.

ANN results

As stated above, various ANNs were trained in order to determine the optimum set of parameters and hyperparameters for our training data. In this sense, we evaluated a different number of learning rates, activation functions and hidden layers with different numbers of nodes per layer, among others, using 5-fold cross-validation. The most accurate ANN for the training data was obtained for two hidden layers with 36 and 29 nodes per layer respectively, and the *tanh* function as the activation function. The performance of this classifier was analyzed by classifying the test data and calculating its confusion matrix (shown in Fig. 8). Considering as features both the whole normalized MM and the rotation angle θ, all of them standardized, we obtained an overall classification accuracy of 95.18%. The performance of the ANN classifier is similar to the performance of the SVM model discussed above. Small differences between these two models can be observed within each class (see Figs. 7 and 8). In the case of the ANN, all classes and subclasses have an accuracy equal or greater than 90%. Note that vegetation classification is more accurate for our ANN than for our SVM model, although there is also some slight misclassification error between vegetation and pavements. Nevertheless, the relevant urban classes have similar classification accuracy for both models.
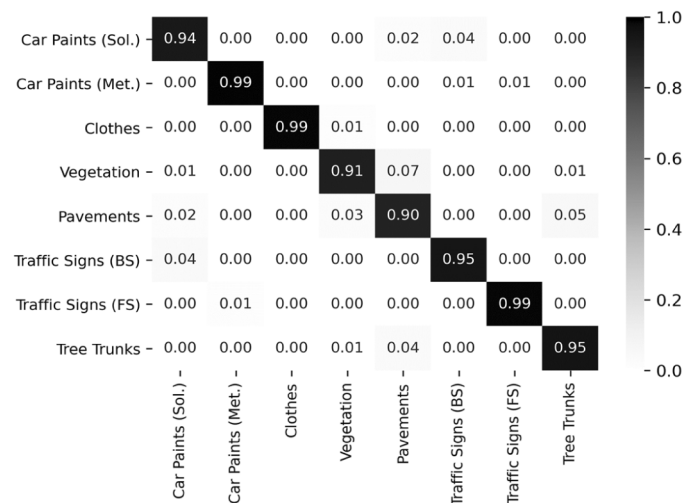

Fig. 8. Normalized confusion matrix showing the performance of our ANN classifier for the previously discussed urban object classes and subclasses. The selected features were: the normalized Mueller matrix and the rotation angle, θ, both of them standardized. The subclasses, solid (Sol.) and metallic (Met.) car paints, and front side (FS) and back side (BS) of traffic signs are presented separately.

3.2.1. Reduced number of features

In addition, by taking into account that the rotation angle, according to the Boruta criterion, is the least important feature (see Fig. 6), new classifiers were trained, reducing the number of inputs to the standardized Mueller matrix. This approach could be useful in remote sensing applications where usually the rotation angle is not measured. As a result, the overall accuracy was slightly reduced compared to the models including θ as an input. In fact, the overall accuracy for the SVM was 94.23%, while for the ANN was 94.50%. Despite some reduction of the accuracy, the relevant classes were already well classified. A comprehensive comparison of the performance of the SVM and ANN models with both sets of features for classifying the previously discussed classes is shown in Fig. 9. The observed decrease in some classes' accuracy (blue circles in

Fig. 9) is attributed to the elimination of the sample's orientation in the training and validation processes.
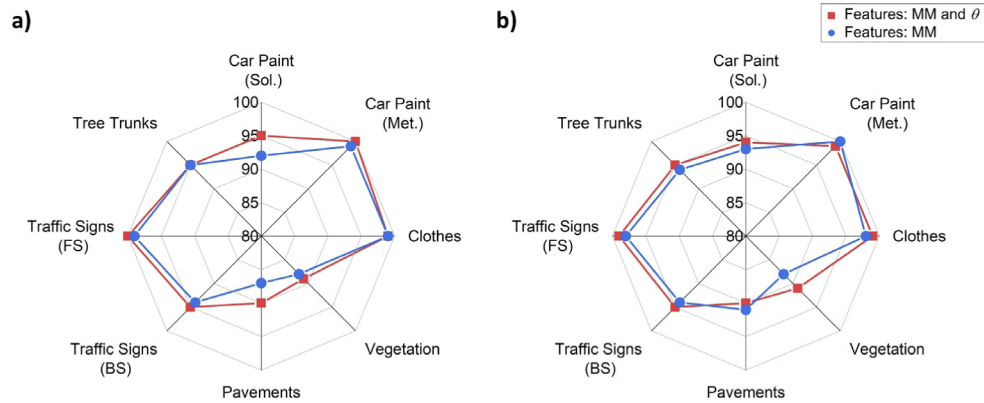

Fig. 9. Performance of the a) SVM and b) ANN classifiers for a different set of features: the Mueller matrix and the rotation angle θ (red squares), and only the Mueller matrix (blue circles).

It can be observed that pavements and vegetation have a higher misclassification rate than the remaining classes for both models. However, in the case of autonomous vehicles, this misclassification error is not critical since these classes are related to static landscape elements, contrary to car paints or clothes, and their discrimination, in an autonomous driving context, probably is not as relevant as distinguishing vehicles, traffic signs or pedestrians. The misclassification rate between pavements, vegetation classes and tree trunks is comparable for both trained models and is related to the similar polarimetric response of these three urban object classes.

4. Conclusions

We presented results that highlight the potential of Mueller matrix measurements to classify urban objects. We demonstrated how the determination of the MM of several real-world urban objects combined with classification models (in particular, SVM and ANN) can be used for vehicles, pedestrians and traffic signs (among others) recognition. With this aim, a quantitative analysis of the MM of several real urban objects has been provided, such as traffic signs, automotive (metallic and solid) paints, vegetation, tree trunks, clothes of pedestrians and pavements. For that, a complete Mueller matrix imaging polarimeter working at 1550 nm was developed. The wavelength used for this study was chosen by considering eye safety, lower solar background and atmospheric attenuation reduction. Moreover, the samples were illuminated at different angles in order to study its importance in urban objects classification. The measured experimental Mueller matrices show differences between the considered groups of urban objects.

The Boruta algorithm was used to identify the most relevant features of our database. Among the different studied inputs, it showed that the knowledge of the angle of illumination of the samples is less important than the knowledge of all (normalized) MM elements, while m_{33}, m_{10} and m_{01} are the features with the highest importance. Afterwards, we developed two different types of classifiers: the ANN and SVM. Both models provided a good performance for classifying new samples into the predefined object classes, with an overall accuracy higher than 95%, proving that polarimetric measurements combined with machine learning can be successfully used for road scene recognition and classification applications, even for materials with similar polarimetric properties (such as clothes, vegetation, tree trunks and pavements). Moreover, we also studied the ANN and SVM classifiers when considering only the elements of

the MM as features with only a slight decrease in the overall accuracy. These results suggest that a remote sensor intended to measure the MM of urban objects and afterwards to classify this information by using a pre-trained ANN or SVM would be an additional valuable method for outdoor environments recognition, even if the rotation angle of the objects is unknown. To reduce misclassification errors, we propose to study the benefits of combining polarimetry with image processing techniques. This method could be of interest to improve the performance in the classification of real urban objects.

The results of this study can be of interest for autonomous driving, highlighting the use of machine learning models for identifying different categories of real-world objects based on Mueller matrix measurements. Future work will focus on implementing this technique in sensors used by autonomous vehicles, such as LiDAR sensors. In addition, the possibility of determination of the angle of incidence from the set of MMs will be investigated.

Disclosures. The authors declare that there are no conflicts of interest related to this article.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

Supplemental document. See Supplement 1 for supporting content.

References

1. F. Duarte and C. Ratti, "The Impact of Autonomous Vehicles on Cities: A Review," J. Urban Technol. **25**(4), 3–18 (2018).
2. A. O. Salonen and N. Haavisto, "Towards Autonomous Transportation. Passengers' Experiences, Perceptions and Feelings in a Driverless Shuttle Bus in Finland," Sustainability **11**(3), 588 (2019).
3. C. Urmson and W. Whittaker, "Self-driving cars and the Urban challenge," IEEE Intell. Syst. **23**(2), 66–68 (2008).
4. T. J. Crayton and B. M. Meier, "Autonomous vehicles: Developing a public health research agenda to frame the future of transportation policy," J. Transp. Heal. **6**, 245–252 (2017).
5. S. Pettigrew, "Why public health should embrace the autonomous car," Aust. N. Z. J. Public Health **41**(1), 5–7 (2017).
6. M. Alawadhi, J. Almazrouie, M. Kamil, and K. A. Khalil, "A systematic literature review of the factors influencing the adoption of autonomous driving," Int. J. Syst. Assur. Eng. Manag. **11**(6), 1065–1082 (2020).
7. D. J. Yeong, G. Velasco-hernandez, J. Barry, and J. Walsh, "Sensor and Sensor Fusion Technology in Autonomous Vehicles: A Review," Sensors **21**(6), 2140 (2021).
8. S. Campbell, N. O'Mahony, L. Krpalcova, D. Riordan, J. Walsh, A. Murphy, and C. Ryan, "Sensor Technology in Autonomous Vehicles : A review," *2018 29th Irish Signals Syst. Conf.* 1–4 (2018).
9. J. Kocic, N. Jovicic, and V. Drndarevic, "Sensors and Sensor Fusion in Autonomous Vehicles," 2018 26th Telecommun. Forum 420–425 (2018).
10. S. Kuutti, R. Bowden, Y. Jin, P. Barber, and S. Fallah, "A Survey of Deep Learning Applications to Autonomous Vehicle Control," IEEE Trans. Intell. Transport. Syst. **22**(2), 712–733 (2021).
11. M. Masmoudi, H. Ghazzai, M. Frikha, and Y. Massoud, "Object detection learning techniques for autonomous vehicle applications," *Proc IEEE Int. Conf. Veh. Electron. Saf.*, 1–5 (2019).
12. H. J. Vishnukumar, B. Butting, C. Muller, and E. Sax, "Machine learning and deep neural network - Artificial intelligence core for lab and real-world test and validation for ADAS and autonomous vehicles: AI for efficient and quality test and validation," *Proc. IEEE Intell. Syst. Conf.*, 714–721 (2017).
13. D. H. Goldstein, *Polarized Light* (CRC Press, 2014).
14. F. Snik, J. Craven-Jones, M. Escuti, S. Fineschi, D. Harrington, A. De Martino, D. Mawet, J. Riedi, and J. S. Tyo, "An Overview of Polarimetric Sensing Techniques and Technology with Applications to Different Research Fields," Proc. SPIE **9099**, 909901–90990B20 (2014).
15. M. Anastasiadou, A. De Martino, D. Clement, F. Liège, B. Laude-Boulesteix, N. Quang, J. Dreyfuss, B. Huynh, A. Nazac, L. Schwartz, and H. Cohen, "Polarimetric imaging for the diagnosis of cervical cancer," phys. stat. sol. (c) **5**(5), 1423–1426 (2008).

16. A. Pierangelo, A. Benali, M. R. Antonelli, T. Novikova, P. Validire, B. Gayet, A. De Martino, J. H. Wei, D. Xing, J. J. Lu, H. M. Gu, G. Y. Wu, and Y. Jin, "Ex-vivo characterization of human colon cancer by Mueller polarimetric imaging," Opt. Express **19**(2), 1582–1593 (2011).

17. A. Van Eeckhout, A. Lizana, E. Garcia-Caurel, J. J. Gil, A. Sansa, C. Rodríguez, I. Estévez, E. González, J. C. Escalera, I. Moreno, and J. Campos, "Polarimetric imaging of biological tissues based on the indices of polarimetric purity," J. Biophotonics **11**(4), e201700189 (2018).

18. L. M. S. Aas, P. G. Ellingsen, and M. Kildemo, "Near infra-red Mueller matrix imaging system and application to retardance imaging of strain," Thin Solid Films **519**(9), 2737–2741 (2011).

19. N. Hong, R. A. Synowicki, and J. N. Hilfiker, "Mueller matrix characterization of flexible plastic substrates," Appl. Surf. Sci. **421**, 518–528 (2017).

20. A. Van Eeckhout, E. Garcia-Caurel, T. Garnatje, M. Durfort, J. C. Escalera, J. Vidal, J. J. Gil, J. Campos, and A. Lizana, "Depolarizing metrics for plant samples imaging," PLoS One **14**(3), e0213909 (2019).

21. J. Hough, "Polarimetry: A powerful diagnostic tool in astronomy," Astron. Geophys. **47**(3), 3.31–3.35 (2006).

22. D. V. Vorobiev, Z. Ninkov, and N. Brock, "Astronomical Polarimetry with the RIT Polarization Imaging Camera," PASP **130**(988), 064501 (2018).

23. D. Vorobiev, Z. Ninkov, L. Bernard, and N. Brock, "Imaging Polarimetry of the 2017 Solar Eclipse with the RIT Polarization Imaging Camera," PASP **132**(1008), 024202 (2020).

24. J. S. Tyo, D. L. Goldstein, D. B. Chenault, and J. A. Shaw, "Review of passive imaging polarimetry for remote sensing applications," Appl. Opt. **45**(22), 5453–5469 (2006).

25. O. Dubovik, Z. Li, M. I. Mishchenko, D. Tanré, Y. Karol, B. Bojkov, B. Cairns, D. J. Diner, W. R. Espinosa, P. Goloub, X. Gu, O. Hasekamp, J. Hong, W. Hou, K. D. Knobelspiesse, J. Landgraf, L. Li, P. Litvinov, Y. Liu, A. Lopatin, T. Marbach, H. Maring, V. Martins, Y. Meijer, G. Milinevsky, S. Mukai, F. Parol, Y. Qiao, L. Remer, J. Rietjens, I. Sano, P. Stammes, S. Stamnes, X. Sun, P. Tabary, L. D. Travis, F. Waquet, F. Xu, C. Yan, and D. Yin, "Polarimetric remote sensing of atmospheric aerosols: Instruments, methodologies, results, and perspectives," J. Quant. Spectrosc. Radiat. Transf. **224**, 474–511 (2019).

26. D. A. LeMaster, A. H. Mahamat, B. M. Ratliff, A. S. Alenin, J. S. Tyo, and B. M. Koch, "SWIR active polarization imaging for material identification," Proc. SPIE **8873**, 887301–88730O8 (2013).

27. I. J. Vaughn, B. G. Hoover, and J. S. Tyo, "Classification using active polarimetry," Proc. SPIE **8364**, 836401–836401-8 (2012).

28. D. G. Jones, D. H. Goldstein, and J. C. Spaulding, "Reflective and polarimetric characteristics of urban materials," Proc. SPIE **6240**, 624001–62400A10 (2006).

29. M. Kupinski and L. Li, "Evaluating the utility of mueller matrix imaging for diffuse material classification," J. Imaging Sci. Technol. **64**(6), 60409 (2020).

30. B. J. DeBoo, J. M. Sasian, and R. A. Chipman, "Depolarization of diffusely reflecting man-made objects," Appl. Opt. **44**(26), 5434–5445 (2005).

31. R. Blin, S. Ainouz, S. Canu, and F. Meriaudeau, "Road scenes analysis in adverse weather conditions by polarization-encoded images and adapted deep learning," *Proc. IEEE Intell. Transp. Syst. Conf.* 27–32 (2019).

32. R. Blin, S. Ainouz, S. Canu, and F. Meriaudeau, "A new multimodal RGB and polarimetric image dataset for road scenes analysis," *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work.* 867–876 (2020).

33. K. Omer and M. Kupinski, "Mid-fusion of road scene polarization images on pretrained RGB neural networks," J. Opt. Soc. Am. A **38**(4), 515–525 (2021).

34. B. Javidi, G. Krishnan, K. Usmani, and T. O'Connor, "Deep learning polarimetric three-dimensional integral imaging object recognition in adverse environmental conditions," Opt. Express **29**(8), 12215–12228 (2021).

35. R. Blin, S. Ainouz, S. Canu, and F. Meriaudeau, "Multimodal Polarimetric And Color Fusion For Road Scene Analysis In Adverse Weather Conditions," *Proc. IEEE Int. Conf. Image Process.* 3338–3342 (2021).

36. J. P. Brown, R. G. Roberts, D. C. Card, C. L. Saludez, and C. K. Keyser, "Hybrid passive polarimetric imager and lidar combination for material classification," Opt. Eng. **59**(07), 1 (2020).

37. E. Korevaar, H. Willebrand, J. Schuster, and S. Bloom, "Understanding the performance of free-space optics," J. Opt. Netw. **2**(6), 178–200 (2003).

38. R. Henderson and K. Schulmeister, *Laser Safety*, 1st ed. (CRC Press, 2003).

39. D. A. Rockwell and G. S. Mecherle, "Wavelength selection for optical wireless communications systems," Proc. SPIE **4530**, 27–35 (2001).

40. A. Peinado, A. Lizana, and J. Campos, "Optimization and tolerance analysis of a polarimeter with ferroelectric liquid crystals," Appl. Opt. **52**(23), 5748–5757 (2013).

41. S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*, 1st ed. (Cambridge University Press, 2014).

42. H. K. Lam, U. Ekong, H. Liu, B. Xiao, H. Araujo, S. H. Ling, and K. Y. Chan, "A study of neural-network-based classifiers for material classification," Neurocomputing **144**, 367–377 (2014).

43. C. J. C. Burges, "A Tutorial on Support Vector Machines for Pattern Recognition," Data Min. Knowl. Discov. **2**(2), 121–167 (1998).

44. C. M. Bishop, *Pattern Recognition and Machine Learning*, 1st ed. (Springer New York, 2006).

45. C. C. Chang and C. J. Lin, "LIBSVM: A library for support vector machines," ACM Trans. Intell. Syst. Technol. **2**(3), 1–27 (2011).

46. J. J. Gil, "Characteristic properties of Mueller matrices," J. Opt. Soc. Am. A **17**(2), 328–334 (2000).
47. S. R. Cloude, "Conditions For The Physical Realisability Of Matrix Operators In Polarimetry," Proc. SPIE **1166**, 177–185 (1990).
48. M. B. Kursa, A. Jankowski, and W. R. Rudnicki, "Boruta – A System for Feature Selection," Fundam. Informaticae **101**(4), 271–285 (2010).
49. M. B. Kursa and W. R. Rudnicki, "Feature Selection with the Boruta Package," J. Stat. Softw. **36**(11), 1–13 (2010).
50. N. Umov, "Chromatische depolarisation durch lichtzerstreuung," Phys. Z **6**, 674–676 (1905).
51. M. K. Kupinski, C. L. Bradley, D. J. Diner, F. Xu, and R. A. Chipman, "Angle of linear polarization images of outdoor scenes," Opt. Eng. **58**(08), 1 (2019).