



OPEN Physics-assisted machine learning for THz time-domain spectroscopy: sensing leaf wetness

Milan Koumans¹, Daan Meulendijks¹, Haiko Middeljans¹, Djero Peeters¹, Jacob C. Douma² & Dook van Mechelen¹✉

Signal processing techniques are of vital importance to bring THz spectroscopy to a maturity level to reach practical applications. In this work, we illustrate the use of machine learning techniques for THz time-domain spectroscopy assisted by domain knowledge based on light–matter interactions. We aim at the potential agriculture application to determine the amount of free water on plant leaves, so-called leaf wetness. This quantity is important for understanding and predicting plant diseases that need leaf wetness for disease development. The overall transmission of 12,000 distinct water droplet patterns on a plastized leaf was experimentally acquired using THz time-domain spectroscopy. We report on key insights of applying decision trees and convolutional neural networks to the data using physics-motivated choices. Eventually, we discuss the generalizability of these models to determine leaf wetness after testing them on cases with increasing deviations from the training set.

The quest to implement societally and industrially relevant applications of THz technology is impeded by aspects such as costs and performance as compared to alternatives. Despite the wide range of exploratory studies of THz technology since the late 1980s, and the numerous suggested applications¹, very few products exist that use THz-based techniques. Often demonstrators underperform compared to a cheaper alternative that already exists or that is conveniently adapted to a new application. However, since THz technology is rather new on the market, costs will remain high at least for a while. The only promising way out is to find a novel application for the technology, which it can almost uniquely serve, and for which the business case is strong enough to support the high costs². In this case, its overall performance needs to justify its usage. On the hardware side, technological maturity has strongly improved over the last decades, although the progress at an integrated level is lagging³. On the software side, despite crucial advancements such as model-based analysis that is now widely employed⁴, the limitations are generality, robustness, and speed, which are essential for realistic application cases.

The agriculture sector is an area where a multitude of sensing technologies are employed to aid management decisions. Certainly due to this reason, there is a vivid interest in searching for applications in this field where THz spectroscopy can make a difference⁵. Globally, pests and pathogens are a big threat to crop production, with yield losses reported in the range of 9–21%⁶. For some pathogens, such as water molds and some fungi, the presence or absence of free water on the surface of leaves, so-called leaf wetness, is key for infection and/or sporulation and is therefore an important parameter in disease epidemiology. A famous example is *Phytophthora infestans*, the causal agent of late blight, responsible for the Irish Potato Famine in the mid-19th century. The development of *P. infestans* depends on the presence of leaf wetness and the surrounding temperature^{7,8}. In ideal circumstances, *P. infestans* can decimate a potato crop in less than 10 days. Control of late blight, as well as that of other pests and pathogens, is nowadays mostly done using crop protection products⁹. There is, however, a strong push from policymakers to reduce this. Early detection and improved predictions of when and where diseases may be expected can help targeted (preventive) measures. Instead of directly detecting the molds, which in the field are difficult to observe, predicting leaf wetness during the growth season is an important input for decision support systems to advise on spraying crop protection products.

Terahertz spectroscopy is particularly suited to accurately sense little amounts of water and is proven to probe leaf properties^{10–14}. The current state-of-the-art uses electrical measurements to infer upon leaf wetness. They are, however, not directly performed on a plant leaf and lack interpretability¹⁵. Other technologies that can sense water, such as visible imaging, will have difficulties estimating very small amounts of water due to a lack of contrast. The infrared range, on the other hand, contains several strong water absorptions bands, but

¹Department of Electrical Engineering, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands. ²Centre for Crop System Analysis, Wageningen University, 6700 AK Wageningen, The Netherlands. ✉email: j.l.m.v.mechelen@tue.nl

its frequency domain operation will make it difficult to distinguish between surface water and water contained inside a leaf, and its shorter wavelength also causes more local sensing.

Model-based signal processing, e.g., based on the transfer matrix method, can be powerful in well-behaved conditions, such as paint layer inspection in an automotive paint shop and wafer metrology in a cleanroom¹⁶, it will not describe well the complex situation at hand of an anisotropic and curved plant leaf in a crop canopy exposed to weather conditions. Signal processing using a data-based learning method could be the solution in this case. Many other studies have used machine learning on THz spectroscopic data^{17–20}. However, most of these studies suffer from a lack of transparency of the used method such that the quality of the result is unclear in aspects of generality and reproducibility.

Here, we report on the application of machine learning models to THz time-domain data in order to determine leaf wetness, defined as the weight of a pattern of water droplets at the surface of a plant leaf. We explore two distinct methods, decision trees and convolutional neural networks, and design the models based on our physics-based interpretation of the data. Both methods accurately predict leaf wetness, independent of the droplet pattern. We obtain an accuracy of about 4% when training and testing on an acquired data set of about 10,600 distinct droplet patterns. In order to explore the performance of our method on closely related data, we test their generality by inferring onto a second acquired data set of about 1500 droplet patterns where the leaf material was flipped in the beam path. The majorly reduced predictability of 40–50% indicates the sensitivity of the models to learning less relevant features. As a solution, we suggest the use of physics-informed algorithms, where a model-based constraint may increase the correlation with the target variable.

The paper is structured as follows: “Results” section first discusses the acquired data sets with a physics-based interpretation, and subsequently sets forth a didactic approach of applying decision trees and convolutional neural networks to these time-domain data and give particular emphasis to feature engineering. The generality of the said models is evaluated in “Discussion” section, where we have designed three validation cases with increasing deviations of the test set as compared the training set. Although from a light-matter-perspective these changes are minor, for the used machine learning algorithms they turn out to be significant. We eventually list challenges that the proposed sensing method will face prior to integration as a demonstrator in a crop canopy.

Results

The transmitted electric field $E(t)$ of water droplet patterns on plasticized plant leaves has been experimentally recorded in transmission geometry at THz frequencies (0.1–3 THz) in the time-domain for about 12,000 distinct patterns as described in “Methods” (Fig. 1a,b). We like to highlight the size of this experimental data set, which is minimally required for meaningful training and testing using machine learning algorithms. The data set is divided into set $\mathcal{E}^+(t)$ for which the droplet patterns are deposited on the top side of the leaf sample, and set $\mathcal{E}^-(t)$ for which the patterns are on the bottom side of the leaf sample. For the first data set, $\mathcal{E}^+(t)$, Fig. 1c shows the distribution of the total sample weight G , being the sum of the weight of the plasticized leaf and the leaf wetness weight g , whereas Fig. 1d displays the distribution of only g . Both gravimetric weights G and g are determined with a scale (see “Methods”). Although a more homogeneous distribution would have been preferred for the latter, the deposition technique and vertical mount of the leaf sample (see “Methods”) favor small droplets above large ones that run down above a given level of g . Fig. 2a shows a given measurement series $E_r(t) \in \mathcal{E}^+(t)$ for selected values of g , which visually drastically changes with g . Besides, each $E(t) \in E_r(t)$ is also related to a given leaf wetness pattern with its droplet size distribution. Therefore, spectra $E(t)$ with identical value g may mutually differ. The standard deviation $\sigma(t)$ of $E_r(t)$ demonstrates that the influence of water on $E(t)$ is predominantly present in the range (4, 5) ps, related to a single passage of THz radiation through the droplet pattern and leaf sample (Fig. 2b). However, maybe of more importance for sensing water is the range of $E(t)$ where the radiation internally reflects within the water droplets. From simulating $E(t)$ using a transfer matrix model, we find that the first internal reflection inside droplets shows up in $E(t)$ as a shoulder right after the largest positive peak, that is, around 5.5 ps. This can be better visualized through the quantity $\xi = E(t)/\max(E(t)) - E_0(t)/\max(E_0(t))$ with E_0 referring to a leaf sample with $g = 0$ (see Fig. 2c). When plotting ξ for the entire data set $\mathcal{E}^+(t)$ as a function of g (Fig. 2d), this effect shows up a fork for low g around 5.5 ps. The ray trajectory in the droplet-plastic-leaf-plastic system is, however, sufficiently complex that more subtle features in $E(t)$ cannot be uniquely attributed to a specific radiation path. To obtain an accurate prediction of the leaf wetness g , we employ two different data-driven methods, decision tree regression and convolutional neural networks. For each algorithm, we determine the mean absolute error and median percentage difference, defined as the median of $|g_p - g|/(g + \epsilon)$ with the predicted leaf wetness weight g_p and small ϵ for stability when $g \approx 0$.

Decision trees

The above-described problem to predict g_p from the experimental input matrix $\mathcal{E}^+(t)$ is equivalent to determining a function f for which holds that $f : \mathcal{E}^+(t) \rightarrow g$. We have ascertained that the relation between distinct features of $\mathcal{E}^+(t)$, as discussed below, and the target variable g is not linear. To evaluate this problem, we choose a decision tree method which is proven to work with non-linear mappings f for mid-sized data sets such as that of $\mathcal{E}^+(t)$. A decision tree is a supervised machine learning method that is shaped as a flowchart in which iterative decisions lead to a piece-wise approximation of the target variable (see Fig. 3a for a single class and two features x_1 and x_2). We refer to “Methods” for further details. As common for machine learning algorithms, decision trees fit on features that are reminiscent of the data. Many reports employing machine learning methods on THz time-domain data either first convert to the frequency domain, thereby losing crucial spatial information, or simply use all data of the time trace $E(t)$ ^{18,19,21}. However, points within $E(t)$ are time-correlated and when directly used as features, the correlation may hamper determination of the independent effect of each feature on the target variable. For this reason, and in order to make the computation resource efficient, feature engineering was used to extract the

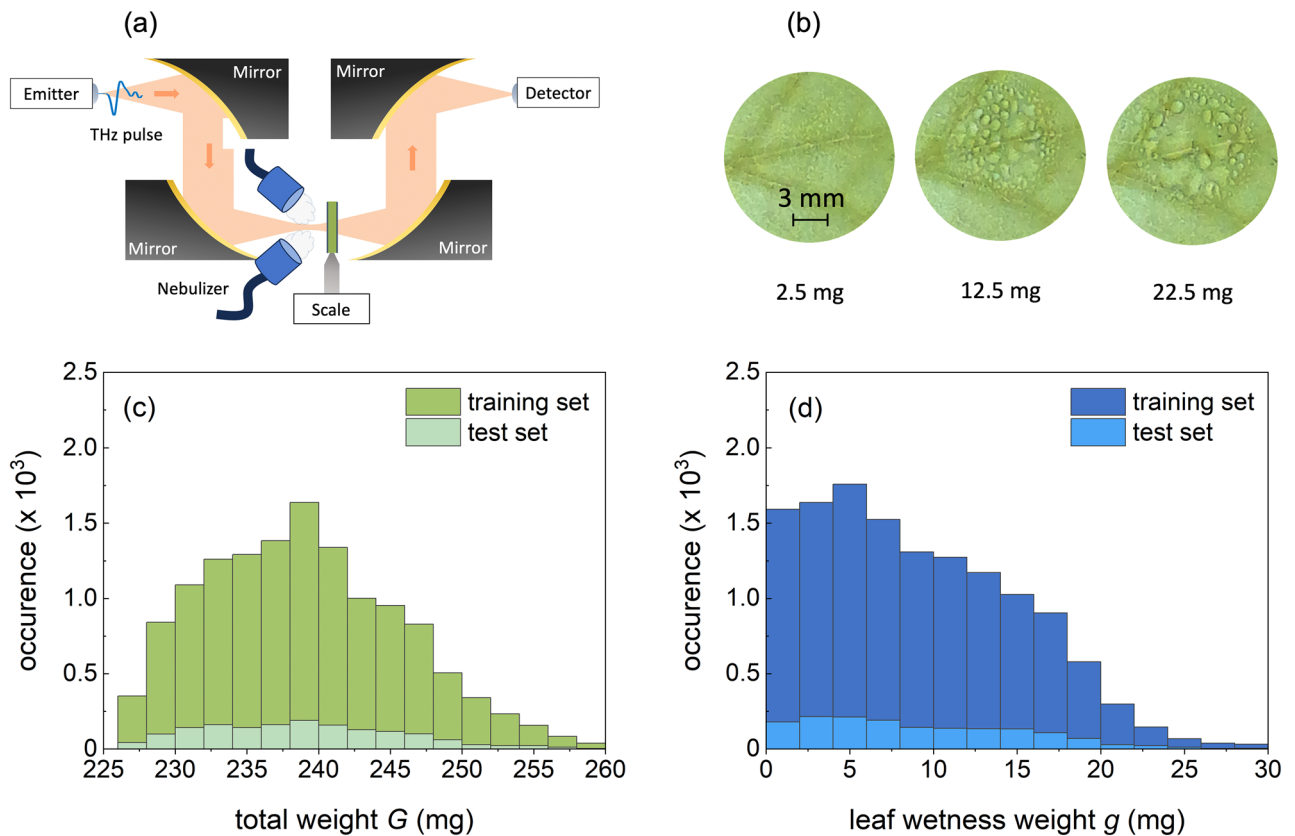


Figure 1. (a) Experimental setup showing a plasticized leaf on a scale in a THz time-domain transmission configuration. (b) Selected images from a measurement series showing the formation of a droplet pattern with leaf wetness weight $g = 2.5, 12.5$ and 22.5 mg, respectively. The THz beam size is comparable to the size of the droplet pattern. (c) Distribution of the total weight G for data set $\mathcal{E}(t)$. (d) Distribution of the leaf wetness weight g for data set $\mathcal{E}(t)$.

most reminiscent features from $E(t)$, which are not temporally correlated. We hereto fit a polynomial function of degree n to relevant temporal ranges of $E(t)$ (Fig. 3b), as further detailed in “Methods”. n is chosen as small as possible, but large enough to capture reminiscent features of $E(t)$ to best predict g . Although visually the fit may not seem optimal, for the algorithm it is. Besides the coefficients of the polynomial terms, also the beginning of the time window t_{start} is added as a feature as well as the absolute air humidity a , which has a strong influence on $E(t)$. Optimization for the displayed range in Fig. 3b leads to $n = 11$ as described in “Methods”. The feature vector thus reads $[t^0, t^1, \dots, t^{11}, t_{\text{start}}, a]$, where t^0 is the bias term and will be further indicated like that. Among these features there can be, however, ones that are mutually correlated and ones that are only a little related to the target variable. We use recurrent feature elimination to reduce the dimensionality of the feature matrix and enhance the performance of the model, as further detailed in “Methods”. Eventually, hyperparameters responsible for the regularization of the final ensemble model are determined using 5-fold cross-validation on a training set consisting of 85 % of data set $\mathcal{E}(t)$.

The predictive performance of a bagged decision tree as described here and further detailed in the methods was evaluated on an unseen test set consisting of 15 % of data set $\mathcal{E}(t)$ and is shown in Fig. 5a. The mean absolute error on g_p is $0.35_{-0.1}^{+0.17}$ mg and the median percentage difference is $3.4_{-1}^{+1.5}$ %. The mean inference time is 139 ms ($\sigma = 22$ ms) per sample using the hardware as mentioned in “Methods”. The indicated error bars are motivated in “Discussion”.

Convolutional neural network

An artificial neural network is a convenient tool to autonomously discover intricate patterns and representations of signals. The simplest neural network predicts an output variable y_i given input data X_{train} and known target variable y_{train} as in linear regression by learning a mapping function $f: X_{\text{train}} \rightarrow y_{\text{train}}$. The architecture of such a neural network consists of an input layer with input data $x_i \in X_{\text{train}}$ and an output layer with output data y_i . The real strength of a neural network is, however, to find patterns and representations in high-dimensional input data. Hereto, cross-correlations between the inputs are learned using so-called perceptrons (artificial neurons) that output a latent variable h depending on input x_i . Subsequently, an activation function, for instance $\varphi(h) = \max(0, h)$, is applied to h . This has the effect that only neuron output of sufficient importance is fed deeper into the network and correlated to the output of neurons in the subsequent layer. In this way, the network can learn non-linear relationships beyond the simple perceptron model. In case the input can be represented as an image, patterns are more effectively learned by using a so-called convolutional neural network (CNN). In such

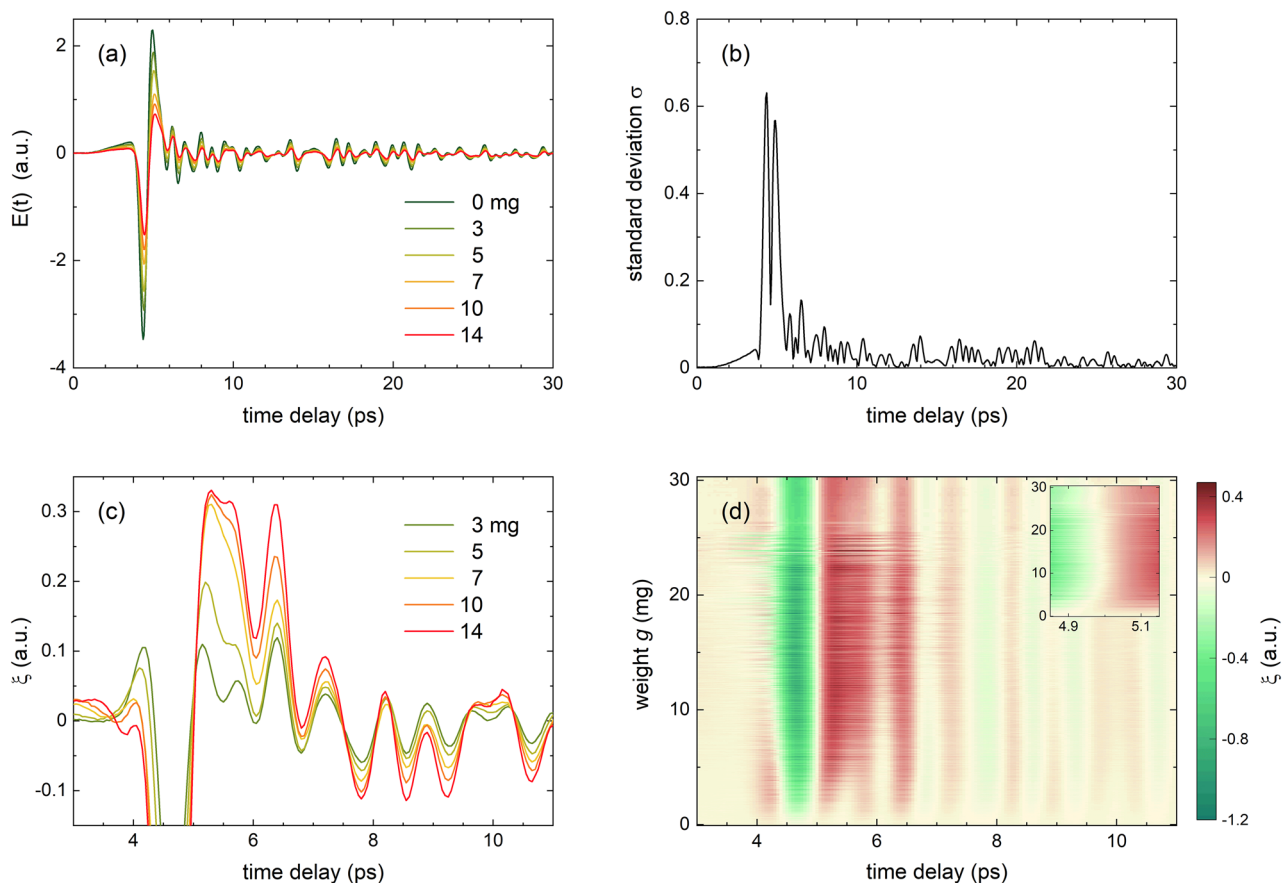


Figure 2. (a) Experimentally determined $E(t) \in E_r(t)$ for selected values of g at ambient conditions (see “Methods”). (b) Standard deviation $\sigma(t)$ of $E_r(t)$. (c) $\xi = E(t)/\max(E(t)) - E_0(t)/\max(E_0(t))$ with $E_0(t)$ the transmission for $g = 0$. This quantity provides detailed insight into temporal ranges of large variation. (d) ξ vs. g . The inset highlights the temporal shift around 5 ps as a function of g .

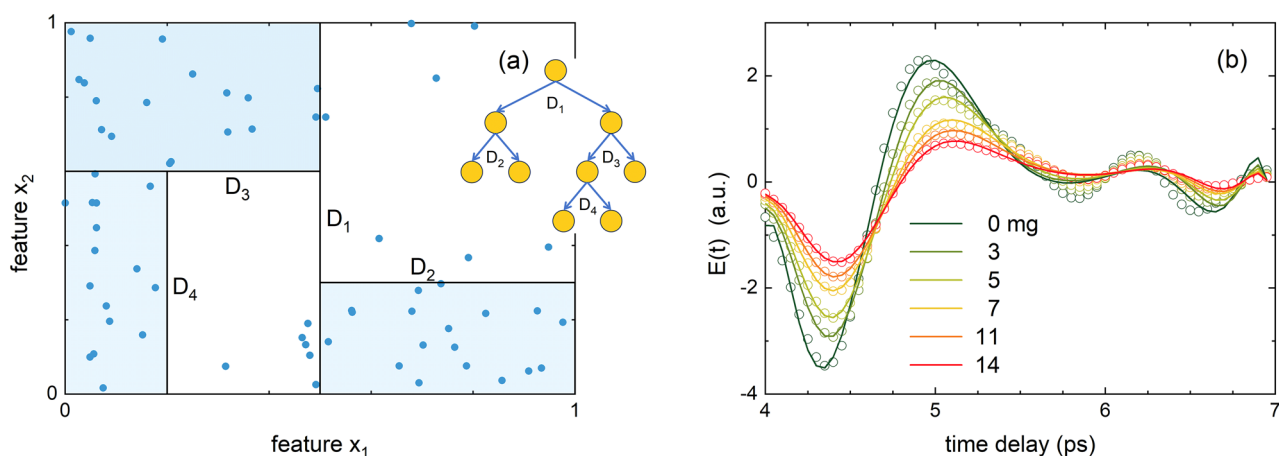


Figure 3. (a) Schematic view of the mechanism of a decision tree. The dots represent data containing features x_1 and x_2 , and the black lines show the division of the parameter space according to split criteria D_s with $s = 1, 4$. The predicted value of each section is the average of the benchmark values in the corresponding division. (b) Experimentally determined $E(t)$ for selected values of g (circles) together with polynomial fits with $n = 11$ (solid lines).

a network a small kernel matrix is scanned over the image matrix to learn local relations between data points. For pictures, these kernel matrices can represent lines and circles, but also more complicated patterns, which are learned by the network. In addition, the method also works for correlated 1-dimensional signals.

The aim here is to construct a CNN, train it on the experimentally acquired time-domain data $\mathcal{E}(t)$, and compare its performance to predict g to prediction using decision trees as discussed in the previous section. For the few studies where THz time-domain data is used for CNNs, rarely the full functional $E(t)$ is used despite its spectral richness and possibility to effectively augment the number of samples as mentioned in “Discussion”^{20,22,23}. A typical CNN consists of many layers which can be grouped into a feature extraction part and a regression part. The feature extraction part consists of multiple convolutional layers ℓ that utilize kernel operations to convolve over the input vector, which enables the detection of pertinent features in signal $E(t)$. Early layers reveal the local context of signal $E(t)$, whereas deeper layers combine activations of different temporal regions to extract the global context of signal $E(t)$. The patterns of the kernels are determined through backpropagation as described in “Methods”. For each layer, the number of distinct kernels k , also called filters, has been empirically chosen as a function of the size of $E(t)$, the complexity of detected patterns in $E(t)$, and the condition to keep the network size as small as possible. Hence, to capture basic patterns, layer $\ell = 1$ of our CNN has four filters, each of dimension (1×3) applied to each $E(t)$ of dimension (1×760) (see Fig. 4m). This results in a feature map of dimension $(1 \times 760 \times 4)$. Fig. 4a,g displays the activation of $E(t)$ for $g = 4$ and 15 mg, respectively, that is, the average of the four filters projected onto t . Similar to model-based signal processing of $E(t)$ ²⁴, CNN shows the largest activation in the range around the absolute minimum of $E(t)$ and hence demonstrates the local character of this first layer. Batch normalization has been applied after each convolutional layer, which normalizes the features during forward propagation. This ensures activation throughout a deep network and results in improved convergence while it simultaneously works as a regularizer²⁵. A so-called max-pooling layer is inserted behind each normalization layer to reduce the network size by downsampling the resulting feature maps. Conceptually, this layer structure increases the receptive field of each neuron. ℓ should be chosen such that the receptive field of neurons in the last layer covers the complete input signal $E(t)$. In our case, we empirically evaluated that $\ell = 6$. We can verify the feature extraction performance of this architecture by inspecting the activations for two cases of $g = 4$ and 15 mg (see Fig. 4a–l). With increasing ℓ , the activation loses its local character and spreads throughout the entire temporal range. The concomitant widening and increasing height of the block functions is the effect of the max-pooling operations, where every iteration halves the time window and eliminates the smallest values. The activation shows that for earlier layers, the range (4, 5) ps is of main importance for the network. This is in agreement with the earlier observation that the largest amplitude of $E(t)$ varies most with g (cf. Fig. 2a), and as such is the most basic pattern of $E(t)$. Interestingly, for $\ell = 3, 4$ and 5, the activations show increased values also in the range (5, 7) ps. As mentioned before, for $g < 7$ mg first internal reflections within the droplets occur around 5.5 ps. In addition, from modeling $E(t)$ using the transfer matrix method, we find that the first internal reflection within the leaf material, although strongly damped, occurs around 7 ps. It thus turns out that these regions, which have an increased importance from a light–matter perspective, are likewise important for a CNN.

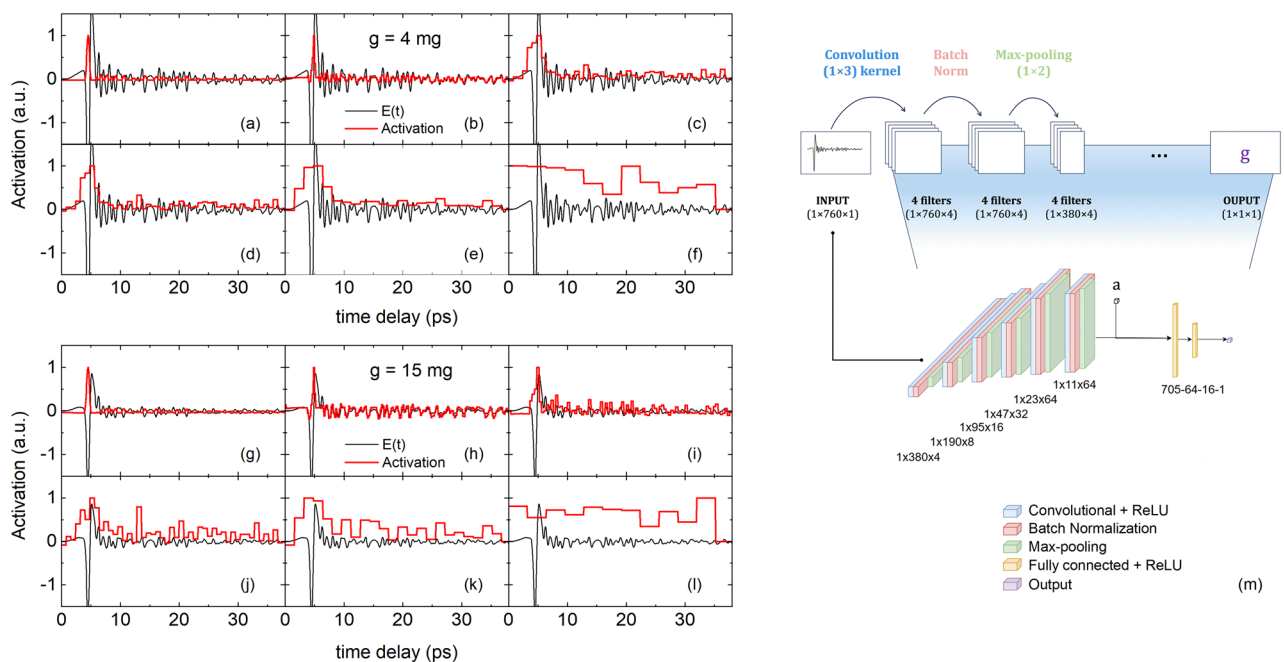


Figure 4. Activation of $E(t)$ per convolutional layer ℓ for (a–f) $g = 4$ mg and (g–l) $g = 15$ mg averaged over all feature maps (panels (a,g) correspond to $\ell = 1$, and so forth). (m) Architecture of the convolutional neural network, indicating the output dimensions for each layer. The dimensions after each max-pooling layer are given as $1 \times d_\ell \times k_\ell$ where d_ℓ is the length of layer ℓ and k_ℓ the number of filters of layer ℓ .

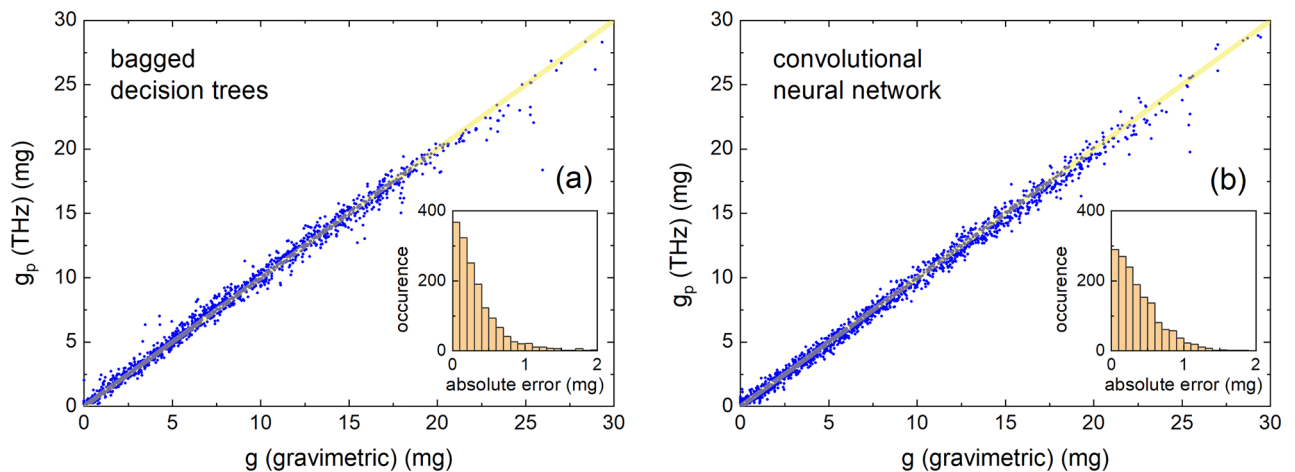


Figure 5. Predicted leaf wetness weight g_p using (a) bagged decision trees and (b) a convolutional neural network, both versus the benchmark leaf wetness weight g . The insets display the distribution of the absolute error of g_p for both models.

Moreover, the increased activity in other regions makes our CNN sensitive to details of $E(t)$ that can contain aspects that are difficult to incorporate into a physical model.

After feature extraction, regression is performed by a fully connected artificial neural network. Hereto, all previous activation layers are flattened to a 1D vector consisting of a concatenation of the 641×11 activations as well as the absolute humidity a . The fully connected network consisting of two layers subsequently learns the nonlinear relation between determined input features and the target variable by minimizing its loss function given by the squared error $|\mathbf{g} - \mathbf{g}_p|^2$, with $|\mathbf{g}| = |\mathbf{g}_p| = |\mathcal{E}(t)|$, using gradient descent optimization.

The CNN is trained on $\mathcal{E}(t)$ for 300 epochs with a batch size of 128 using a train-test ratio of 85:15, where an epoch is a single iteration for which the complete training set has been (forward and backward) propagated through the network. During training, a validation set of 10% of the training set is used to validate the performance after each epoch. Figure 5b graphically shows the performance of the model, having a mean absolute error of $0.38^{+0.17}_{-0.1}$ mg and a median percentage difference of $4.1^{+1.5}_{-1}\%$. The mean inference time is 70 ms ($\sigma = 9$ ms) per sample using the hardware as mentioned in “Methods”.

Discussion

The presented results demonstrate that using two very different data-driven approaches, leaf wetness can be determined using THz spectroscopy within a confidence interval of about 3–4%, independent of the method. This is well beyond the performance of model-based analysis to determine leaf wetness using THz spectroscopy²⁴. Upon comparison to methodologies reported for the closely related property of leaf water content^{12,13,26}, we estimate that our study, with its large data set and comprehensible approach, gives a realistic view of the chances of performing quality control on plants.

Although most studies on machine learning conclude at this point, we are here interested in the generalizability of the models. Therefore, we consider the following further test cases. For case I, we use 37 out of the 39 measurement series of dataset $\mathcal{E}(t)$, train the learning methods according to the same procedure as described, and test on the two remaining series E_{r_1, r_2} . Both methods underperform as compared to the results shown in Fig. 5, with decision trees having a median percentage difference of 8.8% and CNN 6.7% (see Fig. 6a,c). The question thus poses about the origin of this decreased performance. The shape of $E(t)$ is both determined by g and the droplet patterns, as mentioned before. However, the temporal range (4, 5) ps with the largest amplitudes is of leading importance for predicting g , and $\min_t \xi(t) \approx 5$ ps manifests a monotonous variation with g (see the inset of Fig. 2d), independent of the droplet patterns, both suggesting that g predominantly determines $E(t)$. Yet, the values of g of the test set are a subset of the range of g on which the models are trained. The effect of droplet patterns on $E(t)$, on the other hand, is more subtle and beyond the range of the largest amplitudes, as discussed before. We conjecture that the test set contains droplet patterns, rather than different values of g , that are unseen to both methods. The relative underperformance of decision trees is then probably directly linked to its mechanism for which a slight variation in $E(t)$, here due to a different pattern, alters the polynomial fits and thus features on which it has been trained. CNN on the other hand directly learns regions in $E(t)$ which are of relevance for learning g for all trained droplet patterns. Case II considers $\mathcal{E}(t)$ as well as an additional 1501 time traces $E(t)$, referred to as $\mathcal{E}'(t) \notin \mathcal{E}(t)$, for which the droplet pattern is created on the bottom side of the leaf sample. Instead of shining on the smooth and reflective top surface as for $\mathcal{E}(t)$, for $\mathcal{E}'(t)$ the beam now enters the leaf through the dull and rough bottom surface, thereby significantly changing $E(t)$ (see inset of Fig. 6). Please note that the droplet pattern is always on the emitter side. The methods are trained on a subset of $\mathcal{E}(t) \cup \mathcal{E}'(t)$ and tested on an unseen subset of the same data set. Figure 6b,d shows a similar performance in predicting g_p as when considering only $\mathcal{E}(t)$ (cf. Fig. 5), also indicated by the median percentage difference of 4.4% for decision trees and 3.8% for CNN. This indicates that the models can extract relevant features within $E(t)$ no matter the underlying leaf material. In case III, we train the models on $\mathcal{E}(t)$ and test them on the unseen

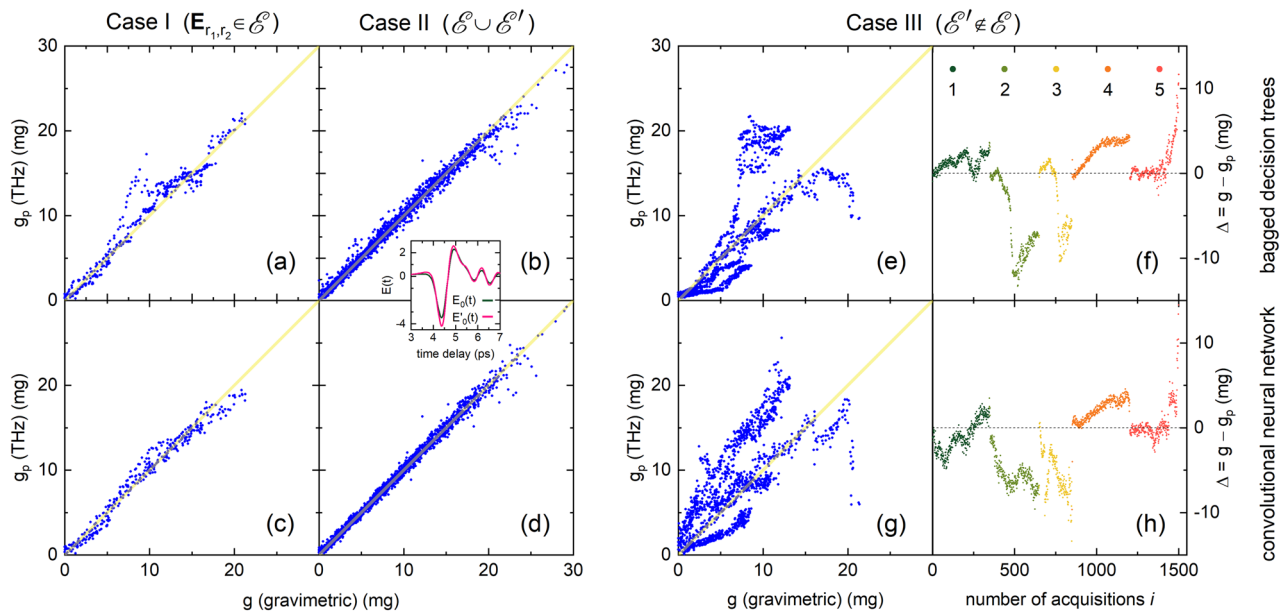


Figure 6. g_p vs. g using (a,b,e) bagged decision trees and (c,d,g) CNN, for three test cases. For each case, the test data set is mentioned in the figure title. Case I tests generalizability towards unseen droplet patterns using two unseen measurement series $E_r(t) \in \mathcal{E}(t)$. Case II tests generalizability towards droplet patterns on top and below a leaf sample using the enlarged data set $\mathcal{E}(t) \cup \mathcal{E}'(t)$. Case III tests generalizability towards droplet patterns on an unseen leaf surface by training on $\mathcal{E}(t)$ and testing on $\mathcal{E}'(t)$. The inset displays $E(t)$ and $E'(t)$, both at $g = 0$, indicating the difference in the transmitted electric field between irradiating a leaf from the top side and the bottom side, respectively. $\Delta = g - g_p$ for the 5 data sets (indicated by the color scheme) of $\mathcal{E}'(t)$ vs. the number of acquisitions i for (f) bagged decision trees and (h) CNN.

data set $\mathcal{E}'(t)$. This would correspond to the practical situation where the method has learned based on droplet patterns on leaves which physically are not the same as those on which the model is inferred. The performance of the methods is rather lousy (Fig. 6e,g), as indicated by the median percentage difference of 39% for decision trees and 50% for CNN. Figure 6f,h displays the absolute deviation $\Delta = g - g_p$ grouped according to the five measurement series of $\mathcal{E}'(t)$ as a function of the number of acquisition i . Each series starts at low i with $g = 0$ after which both i and g increase concomitantly. It turns out that Δ is small at low g and deviates with increasing g . Signatures of droplet patterns with small droplets are thus overall well recognized by the models, despite the unseen patterns of $E(t)$ due to the flipped leaf. We conjecture that this sensitivity may be related to the presence of reminiscent features of $E(t)$ for $g < 7$ mg, due to internal reflections inside the droplets and leaf material (see “Results”), on which the network can train. These three test cases indicate that the learning model needs input training data that are quite close to the data set from which will be inferred. Although it is promising that g is well predicted from data corresponding to droplet patterns that have not been seen before, unseen variations of the leaf material cause changes in $E(t)$ to which the methods are not robust. This could be resolved by training the models on data that besides a large variation of g and droplet patterns also include a large diversity of leaves.

The confidence and related error bars of each machine learning model also depend on specific parameters that are chosen within the architecture and implementation. For decision trees, for instance, we made use of domain knowledge to select four temporal regions within $E(t)$. However, a different choice of the temporal intervals will alter the optimal polynomial order n , modify the feature vector, and thus the performance of the method. For the convolutional neural network, the specific network architecture expressed by variables such as the number of layers, filters, and the number of epochs significantly varies the performance of the network. For the current data set, the quoted performance due to the described variations has an estimated error bar of the order of around 1.5%. However, also the data set size is an important parameter. Whereas for decision trees 9018 time traces with 11 features each gives a significant data set of 10^5 features to train on, for CNNs the final feature size is more difficult to estimate. Although all temporal points of $\mathcal{E}(t)$ with size (9018, 760) are used as input, feature extraction leads to much fewer features than the number of elements of this matrix. By visual inspection of the activations (see Fig. 4a–l), we estimate that each $E(t)$ provides some 10^2 patterns reminiscent of g , summed over all feature maps. On the other hand, for each $E(t)$ many patterns together are needed for an accurate prediction of g , which reduces the feature space from the number estimated before. To clarify this point we draw the analogy with having a picture of a cat, where a cat is defined by its eyes, ears, tail, fur, etc. Having only one of these features will not lead to an accurate prediction of the picture showing a cat. We therefore estimate that the total data size on which the CNN trains should be also of the order of 10^5 , which is on the low side of what is common for CNNs.

The performance of the studied method is also determined by the experimental setup. The nebulized water does not always entirely end up in the area covered by the THz beam, thereby contributing to the measured gravimetric weight that is not seen by the THz beam. We expect that these errors are the cause of the horizontal lines in Fig. 2d. Additionally, the accuracy of the gravimetric measurements of g is rather low, mostly due to

natural air convection and air streams caused by the nebulizers. However, the leaf cannot be placed in an enclosed box, as then the relative humidity gets spatial and temporal fluctuations with drastic consequences on $E(t)$. We estimate the absolute error on g between 0.1 and 0.2 mg.

Although these results indicate that our data-driven approaches provide a performant model to determine leaf wetness in a lab setting, the real baptism of fire for the methodology should be a test in the application environment. A first change would be to determine leaf wetness of water nebulized onto a leaf, instead of onto a plastic sheet attached to the leaf. As the analyses are merely sensitive to the influence of the droplet pattern to $E(t)$, we estimate that as long as droplets are formed on the leaf, the performance of the used algorithms is largely comparable. A leaf wetness sensor for the agriculture sector would need to continuously sense leaf wetness during the growth season on a representative size of the crop canopy. A reflection geometry would be preferred over a transmission one, and instead of a single leaf, many leaves will need to be probed simultaneously. In addition, leaves will not be clean and flat but occur with a large distribution of appearances. Moreover, the air can be highly humid or contain rain, and dust, and the beam path may be (temporally) disturbed by passing insects. We estimate that despite the reported performance of the algorithms presented in this work, this real setting, which may vary depending on the canopy, will be challenging for the described methods. A solution may be found in adding model-based constraints to data-driven learning algorithms, which may increase correlation to the target variable rather than to less relevant variations within the data, as explored in physics-informed neural networks²⁷.

Conclusions

We have studied the application of conventional machine learning methods to THz time-domain data to determine leaf wetness. Hereto, we experimentally acquired a large data set of 12,000 distinct time domain traces $E(t)$ corresponding to a distribution of droplet patterns on plant leaves. Using domain knowledge related to the light-matter interaction, we designed and trained a decision tree model and a convolutional neural network and gave special attention to feature extraction. Both models predict leaf wetness with an accuracy of about 4%. The generalizability of the methods was evaluated on unseen datasets with increasing deviations from the training set. We observe a similarity between features important for CNN and those having a physical interpretation. In conclusion, conventional machine learning models can be of additional value when compared to model-based signal processing, especially in case the sample configuration is complex and well beyond a multilayer structure, although the sensitivity upon input data is an obstacle for practical application.

Methods

Samples

The plant material used in this work is *Alliaria petiolata*, also known as *garlic mustard*. The plant is widely spread in the Netherlands, in the wild, and in private gardens. The plant has been grown in the open ground, and leaves are harvested in May to June. Of the selected leaf, a circle of 30 mm diameter is cut and immediately embedded in between two plastic sheets each 0.08 mm thick to prevent it from fast drying out. The plastic sheets have not been sealed at their edges. During the measurement campaign, the leaf naturally dried out on the order of 18% (that is, about 11 mg) in 5 days. This range of relative leaf water content corresponds to variations due to well-watered vs. severe-drought growth conditions.

Experimental setup

The optical properties of a moistened plant leaf have been obtained by performing THz time-domain transmission spectroscopy (Toptica Teraflash Pro). The sample has been put out-of-focus such that the THz beam covers the surface area on which water is nebulized. Water droplet patterns on the plasticized leaf sample have been created by using two nebulizers (Medisana, Inhalator IN 500). The nebulizers are equipped with specially designed nozzles and are symmetrically positioned in the plane of incidence and facing the sample (see Fig. 1a). The nebulizers are filled with distilled water and have been switched on at random to create a wide range of different droplet patterns. The gravimetric leaf wetness weight g of the moist sample is recorded simultaneously with the THz data. Hereto, the plasticized leaf sample was mounted on a dedicated holder attached to a precision scale (Sartorius WZA224-L). In addition, an RGB picture is taken from each droplet pattern using an optical camera (Logitech Brio 4K). Simultaneously, also the absolute air humidity a has been recorded. During the entire measurement campaign, which lasted 7 days, the temperature was 25.0 ± 0.3 °C, and the relative humidity was $42 \pm 6\%$.

Training and inference of the described models are performed using python libraries Sklearn (for Decision trees), TensorFlow (for CNN) on a CPU (Intel i7-7700HQ @ 2.8 GHz), using 4/4 cores and multi-threading enabled.

Data acquisition

Using the experimental setup described above, the transmitted electric field $E(t)$ has experimentally been determined of a plasticized leaf sample that has been moistened. A measurement sequence is divided into separate time slots: 0.5 s for nebulization, 1.0 s waiting time where the nebulization cloud deposits itself onto the sample and diffuses away from the beam path, and 1.5 s for acquiring 50 averages of $E(t)$, recording g and a , and obtaining an RGB picture of the droplet pattern. A measurement series starts with a dry plasticized leaf sample, i.e. $g = 0$, and ends when a first droplet runs down the sample. For a subsequent measurement series, the sample is first wiped dry, and remoistened.

The THz beam shines onto the upper side of the leaf for data contained in $\mathcal{E}(t)$, whereas it shines onto the bottom side of the leaf for data within $\mathcal{E}'(t)$. In both cases, the droplet pattern is directed towards the emitter side. $\mathcal{E}(t)$ contains a total of 10,609 different droplet patterns organized in $|r| = 39$ measurement series $E_r(t)$,

which all together are acquired within 5 days. $\mathcal{E}'(t)$ contains 1501 patterns recorded in $|s| = 5$ measurement series $E'_s(t)$ acquired within 2 days. Each $E(t)$ contains 760 data points and spans a delay of 38 ps.

Data processing

Decision trees

Decisions within regression trees use quantitative split criteria such as the absolute error (L1-norm) or squared error (L2-norm) to determine the best split at a decision point D (see Fig. 3a). A single decision tree is deterministic in its prediction, making it prone to overfitting. To overcome this shortcoming, an ensemble of simplified trees is used. Each simplified tree independently predicts the target value, after which the predictions are averaged. This reduces the variance of the model while maintaining the predictive power of the decision trees. Simplifying the trees by limiting the seen input data is called bagging and is a well-proven method that is capable of learning nonlinear relations between input features and the target variable. Due to the simplistic nature of decision trees, inference is interpretative and most importantly, ensemble modeling leads to improved robustness compared to single estimators.

Feature extraction is performed by fitting a polynomial function of order n to a selected temporal range m of $E(t)$. As $E(t)$ has a high temporal variation, and to keep n low for optimal performance of the method, we choose several ranges to capture all relevant patterns of $E(t)$. The standard deviation σ of $E_r(t)$ is enhanced in the range (4, 22) ps (see Fig. 2b) which we split into $m = 4$ equal temporal regions (4, 7); (7, 10); (15.5, 18.5) and (18.5, 21.5) ps. As $E(t)$ for $g = 0$, $E_0(t)$, manifests a slight temporal shift due to factors such as leaf water content, air humidity, and drift, the onset of all measured traces $E(t)$ occurs at a slightly different value of t . To correct for this, the onset of region $m = 1$ is determined by the slope d/dt of $E(t)/\max(E(t))$. Hereto, $E(t)$ has been interpolated to enhance the temporal precision. Determination of the order n of each range has been done using a grid search algorithm, which is a brute-force search of a predefined set of options. Hereto, we calculate the L2 loss of the validation set after training on $n \in [0, 20]$. For the four different temporal regions of interest we obtain optimal polynomial fits with order $n_1 = 11$, $n_2 = 2$, $n_3 = 4$, and $n_4 = 8$, respectively. Subsequently, we eliminate recursive features from the input features using permutation importance as a ranking criterion²⁸. By shuffling a feature column in the input matrix, the correlation between the feature and target value is evaluated. As such, features with little correlation to the target value are iteratively eliminated, as well as features that are correlated to others that already show dependence on the target value. This reduces the number of input features and in turn, increases the predictive power of the validation set. In this way, the feature vectors obtained by applying the decision tree algorithm simultaneously to all regions m read $[t^2, t^4, t_{\text{start}}, a]$, $[\text{bias}, t^2]$, $[t^1, t^2]$, $[\text{bias}, t^2, t^6]$, for $m = 1, 4$, respectively. One can see that although both a and t_{start} are present within the features for $m = 1$, they are eliminated from the features for $m > 1$ due to mutual correlations.

To train a generalizable model, we use an ensemble of randomized decision trees, opposite to for instance a single tree which is prone to overfitting. The randomization is effectuated by the well-established principle of bootstrap aggregation (or bagging), where different subsets D_i are sampled from training set D with replacement, meaning that the same sample can be sampled multiple times. We further optimize the trees by tuning hyperparameters that control regularization, including the number of samples per tree N_{SAMPLES} , the number of features per tree N_{FEATURES} , and the maximum depth of branches of a tree MAX_DEPTH . After optimization of these hyperparameters using the mentioned grid search algorithm, we obtain $N_{\text{SAMPLES}} = 2000$, and MAX_FEATURES turns out to be equal to N_{FEATURES} such that the decision tree becomes a bagged decision tree. The minimum number of samples after a split is set to 5 to smooth the predicted value of a tree²⁹.

Convolutional neural network

The architecture of the CNN is shown in Fig. 4m. The input data is inserted in a concatenation of six groups each containing a convolutional layer. This structure is motivated by a receptive field which increases with each deeper layer. Since deeper layers learn more complex patterns, the number of filters is increased by a factor of two for each layer. Moreover, the convolutional layers are (after batch normalization, see below) succeeded by a max-pooling layer that decreases the length by a factor of two. Given the kernel dimension (1×3) and that of the input vector $E(t)$ (1×760) , we have chosen for zero padding the extremes of $E(t)$ with a single entry to have the dimension unaltered after convolution. After this feature extraction part, the output is flattened and feature \mathbf{a} is appended before inserting it in a two-layer fully connected network. The entire network has 72,385 trainable parameters: 26,144 are responsible for feature detection and 46,241 for regression.

Input data needs to be normalized before feeding it into the network to ensure proper convergence during gradient descent, as steepest gradient descent algorithms do not possess the property of scale invariance. Normalization of the absolute humidity is performed as $a'_i = (a_i - \mu_a)/\sigma_a$ where μ_a is the mean of a with respect to the training population and σ_a its standard deviation. For all $E(t)$, a uniform division factor of 4 is applied to ensure all data points are within the range $(-1, 1)$. For activation, we make use of the Rectified Linear Unit (ReLU) function, defined as $\varphi(h) = \max(0, h)$.

Convolutional networks require supervised learning to learn the intrinsic patterns to predict the target value. Supervised deep learning models use backpropagation to autonomously shape function f to find relation $f: x \rightarrow y$, where x is the input and y the target. We have used the Adam optimizer to find the optimal model parameters. Adam is a simple and computationally efficient algorithm that combines the advantages of AdaGrad and RMSProp, resulting in an optimizer that is robust and well-suited to a wide range of non-convex optimization problems in the field of machine learning³⁰.

Ethical approval

The authors state permission to harvest leaves of *Alliaria petiolata* used in this study. All methods were carried out in accordance with relevant guidelines and regulations.

Data availability

The datasets and algorithms used during the current study are available from the corresponding author on reasonable request.

Received: 17 October 2023; Accepted: 14 March 2024

Published online: 25 March 2024

References

- Koch, M., Mittleman, D. M., Ornik, J. & Castro-Camus, E. Terahertz time-domain spectroscopy. *Nat. Rev. Methods Primers* **3**, 48. <https://doi.org/10.1038/s43586-023-00232-z> (2023).
- van Mechelen, D. An industrial THz killer application?. *Opt. Photon. News* **26**, 16–18 (2015).
- Leitenstorfer, A. *et al.* The 2023 terahertz science and technology roadmap. *J. Phys. D Appl. Phys.* **56**, 223001. <https://doi.org/10.1088/1361-6463/acbe4c> (2023).
- van Mechelen, J. L. M., Kuzmenko, A. B. & Merbold, H. Stratified dispersive model for material characterization using terahertz time-domain spectroscopy. *Opt. Lett.* **39**, 3853–3856 (2014).
- Rawson, A. & Sunil, C. K. Recent advances in terahertz time-domain spectroscopy and imaging techniques for automation in agriculture and food sector. *Food Anal. Methods* **15**, 498–526. <https://doi.org/10.1007/s12161-021-02132-y> (2022).
- Oerke, E.-C. Crop losses to pests. *J. Agric. Sci.* **144**, 31–43. <https://doi.org/10.1017/S0021859605005708> (2006).
- Bregaglio, S., Donatelli, M., Confalonieri, R., Acutis, M. & Orlandini, S. Multi metric evaluation of leaf wetness models for large-area application of plant disease models. *Agric. For. Meteorol.* **151**, 1163–1172. <https://doi.org/10.1016/j.agrformet.2011.04.003> (2011).
- Huber, L. & Gillespie, T. J. Modeling leaf wetness in relation to plant disease epidemiology. *Annu. Rev. Phytopathol.* **30**, 553–577. <https://doi.org/10.1146/annurev.py.30.090192.003005> (1992).
- Goffart, J.-P. *et al.* Potato production in northwestern Europe (Germany, France, the Netherlands, United Kingdom, Belgium): Characteristics, issues, challenges and opportunities. *Potato Res.* **65**, 503–547. <https://doi.org/10.1007/s11540-021-09535-8> (2022).
- Mittleman, D., Jacobsen, R. & Nuss, M. T-ray imaging. *IEEE J. Sel. Top. Quantum Electron.* **2**, 679–692. <https://doi.org/10.1109/2944.571768> (1996).
- Gente, R. *et al.* Determination of leaf water content from terahertz time-domain spectroscopic data. *J. Infrared Millimeter Terahertz Waves* **34**, 316–323. <https://doi.org/10.1007/s10762-013-9972-8> (2013).
- Gente, R. & Koch, M. Monitoring leaf water content with THz and sub-THz waves. *Plant Methods* **11**, 15. <https://doi.org/10.1186/s13007-015-0057-7> (2015).
- Li, R., Lu, Y., Peters, J. M. R., Choat, B. & Lee, A. J. Non-invasive measurement of leaf water content and pressure-volume curves using terahertz radiation. *Sci. Rep.* **10**, 21028. <https://doi.org/10.1038/s41598-020-78154-z> (2020).
- Singh, A. K., Pérez-López, A. V., Simpson, J. & Castro-Camus, E. Three-dimensional water mapping of succulent *Agave victoriae-reginae* leaves by terahertz imaging. *Sci. Rep.* **10**, 1404. <https://doi.org/10.1038/s41598-020-58277-z> (2020).
- Rowlandson, T. *et al.* Reconsidering leaf wetness duration determination for plant disease management. *Plant Dis.* **99**, 310–319. <https://doi.org/10.1094/PDIS-05-14-0529-FE> (2015).
- van Mechelen, J. L. M., Frank, A. & Maas, D. J. H. C. Thickness sensor for drying paints using THz spectroscopy. *Opt. Express* **29**, 7514. <https://doi.org/10.1364/OE.418809> (2021).
- Park, H. & Son, J.-H. Machine learning techniques for THz imaging and time-domain spectroscopy. *Sensors* **21**, 1186. <https://doi.org/10.3390/s21041186> (2021).
- Wang, Y. *et al.* Terahertz spectroscopic diagnosis of early blast-induced traumatic brain injury in rats. *Biomed. Opt. Express* **11**, 4085. <https://doi.org/10.1364/BOE.395432> (2020).
- Cao, C., Zhang, Z., Zhao, X. & Zhang, T. Terahertz spectroscopy and machine learning algorithm for non-destructive evaluation of protein conformation. *Opt. Quant. Electron.* **52**, 225. <https://doi.org/10.1007/s11082-020-02345-1> (2020).
- Wang, Q. *et al.* Automatic defect prediction in glass fiber reinforced polymer based on THz-TDS signal analysis with neural networks. *Infrared Phys. Technol.* **115**, 103673. <https://doi.org/10.1016/j.infrared.2021.103673> (2021).
- Li, R. *et al.* Nondestructive evaluation of thermal barrier coatings thickness using terahertz time-domain spectroscopy combined with hybrid machine learning approaches. *MDPI* **12**, 1875. <https://doi.org/10.3390/coatings12121875> (2022).
- Mao, Q. *et al.* Convolutional neural network model based on terahertz imaging for integrated circuit defect detections. *Opt. Express* **28**, 5000. <https://doi.org/10.1364/OE.384146> (2020).
- Wang, C. *et al.* Convolutional neural network-based terahertz spectral classification of liquid contraband for security inspection. *IEEE Sens. J.* **21**, 18955–18963. <https://doi.org/10.1109/JSEN.2021.3086478> (2021).
- Koumans, M., Perez-Casanova, A. & Van Mechelen, J. L. M. Sensing moisture patterns using terahertz spectroscopy. In *2022 47th International Conference on Infrared, Millimeter and Terahertz Waves (IRMMW-THz)*, 1–2. <https://doi.org/10.1109/IRMMW-THz50927.2022.9895781> (IEEE, Delft, Netherlands, 2022).
- Ioffe, S. & Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift (2015). [arXiv:1502.03167](https://arxiv.org/abs/1502.03167) [cs].
- Jördens, C., Scheller, M., Breitenstein, B., Selmar, D. & Koch, M. Evaluation of leaf water status by means of permittivity at terahertz frequencies. *J. Biol. Phys.* **35**, 255–264. <https://doi.org/10.1007/s10867-009-9161-0> (2009).
- Karniadakis, G. E. *et al.* Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440. <https://doi.org/10.1038/s42254-021-00314-5> (2021).
- Gregorutti, B., Michel, B. & Saint-Pierre, P. Correlation and variable importance in random forests. *Stat. Comput.* **27**, 659–678. <https://doi.org/10.1007/s11222-016-9646-1> (2017).
- Geurts, P., Ernst, D. & Wehenkel, L. Extremely randomized trees. *Mach. Learn.* **63**, 3–42. <https://doi.org/10.1007/s10994-006-6226-1> (2006).
- Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization (2017). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) [cs].

Acknowledgements

D.v.M. and J.C.D. are grateful to fruitful discussions with Niels Anten (Wageningen University, The Netherlands), and for the facilitating role of research program Synergia-SYstem change for New Ecology-based and Resource efficient Growth with high tech In Agriculture, financed by NWO, the Dutch Research Council (project number

17626), industrial and scientific/research partners, as well as to Rik Vullings (Eindhoven University of Technology) for proofreading the manuscript.

Author contributions

D.v.M. designed and supervised the project, J.C.D. conceived the problem statement, M.K., D.M., and H.M. made the experimental setup and conducted the experiments, H.M. identified the plant material, M.K. analysed the results. D.v.M., M.K. and J.C.D. wrote the manuscript. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to D.M.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024