Making deep neural networks right for the right scientific reasons by interacting with their explanations

Patrick Schramowski, Wolfgang Stammer, Stefano Teso, Anna Brugger,
Franziska Herbert, Xiaoting Shao, Hans-Georg Luigs,
Anne-Katrin Mahlein & Kristian Kersting

Abstract

Deep neural networks have shown excellent performances in many real-world applications. Unfortunately, they may show “Clever Hans”-like behavior—making use of confounding factors within datasets—to achieve high performance. In this work, we introduce the novel learning setting of “explanatory interactive learning” (XIL) and illustrate its benefits on a plant phenotyping research task. XIL adds the scientist into the training loop, such that she interactively revises the original model via providing feedback on its explanations. Our experimental results demonstrate that XIL can help to avoid Clever Hans moments in machine learning and encourages (or discourages, if appropriate) trust into the underlying model.¹

Imagine a plant phenotyping team attempting to characterize crop resistance to plant pathogens. The plant physiologist records a large amount of hyperspectral imaging data. Impressed by the results of deep learning in other scientific areas, she wants to establish similar results for phenotyping. Consequently, she asks a machine learning expert to apply deep learning to analyze the data. Luckily, the resulting predictive accuracy is very high. The plant physiologist, however, remains skeptical. The results are “too good, to be true”. Checking the decision process of the deep model using explainable artificial intelligence (AI), the machine learning expert is flabbergasted to find that the learned deep model uses clues within the data that do not relate to the biological problem at hand, so-called confounding factors. The physiologist loses trust in AI and turns away from it, proclaiming it to be useless.

This example encapsulates a critical issue of current explainable AI [1, 2]. Indeed, the seminal paper of Lapuschkin *et al.* [3] helps in “unmasking Clever Hans predictors and assessing what machines really learn”, however, rather than proclaiming, as the plant physiologist might, that the machines have learned the right predictions for the wrong reasons and can therefore not be trusted, we here showcase that interactions between the learning system and the human user can correct the model, towards making the right predictions for the right reasons [4]. This may also increase the trust in machine learning models. Actually, trust lies at the foundation of major theories of interpersonal relationships in psychology [5, 6], and we argue that interaction and understandability are central to trust in learning machines. Surprisingly, the link between interacting, explaining, and building trust has been largely ignored by the machine learning literature. Existing approaches focus on passive learning only and do not consider the interaction between the user and the learner [7, 8, 9], whereas, interactive learning frameworks such as active [10] and coactive learning [11] do not consider the issue of trust. In active learning, for instance, the model presents unlabeled instances to a user, and in exchange, obtains their label. This is completely opaque—the user is oblivious to the models beliefs and reasons for predictions and how they change in time, and cannot see the consequences of her instructions. In coactive learning, the user sees and corrects the systems prediction, if necessary, but the predictions are not explained to her. So, why should users trust models learned interactively?

Furthermore, although an increasing amount of research investigates methods for explaining machine learning models, even here, the notion of interaction has been largely ignored. Reconsider the study by Lapuschkin *et al.* [3]. They showed that one can find “Clever Hans”-like behavior in popular computer vision models basing their decisions on confounding factors. Based on these findings, the authors recommended a word of caution towards the interest in such models, but they did not offer a solution for correcting their behavior. Particularly in real-world applications, where monitoring for every possible confounding factor or acquiring a new dataset due to existing confounders is time and resource consuming, it is inevitable to move beyond revealing the (wrong) reasons by making a step towards correcting the reasons underlying a models decisions.

Doing so is exactly the main technical contribution of the present study. We introduce the novel learning setting of “explanatory interactive learning” (XIL) and illustrate its benefits in an important scientific endeavor, namely, plant phenotyping. Starting with a learning system that does not deliver biologically plausible explanations for a relevant, real-world task in plant phenotyping, we add the scientist into the training loop, who interactively revises the original model by interacting via its explanations so that it produces trustworthy decisions without a major drop in performance.

¹Published as Nature Machine Intelligence volume 2, pages 476486(2020). DOI: <https://doi.org/10.1038/s42256-020-0212-3>. Final manuscript available at: <https://www.nature.com/articles/s42256-020-0212-3>.

Fig. 1: **Explanatory Interactive Learning (XIL)**. Human users revise learning machines towards trustworthy decision strategies. **(a-left)** Data samples, expert-classifications (checks and Xs with colors indicating the class) and explanations (overlaid with an edge filtered original image for better interpretability) that an expert expects of an ML model. Yellow corresponds to relevant regions, blue to irrelevant regions for a classification. Not even an expert can be certain about potential samples from an early disease stage and what a valid explanation should be. **(a-middle)** Illustration of hyperspectral data consisting of spatial and spectral dimensions. The planes on the top and left sides of the cube correspond to slices taken from the center of the cube but placed on the edges for visualization. **(a-right)** The characteristic reflectance of healthy tissue vs. disease spots. The vertical red, green and blue lines depict the three wavelengths of the RGB dataset. **(b,c)** Classifications of a deep neural network **(b)** and its explanations **(c)**. The learned model clearly uses confounding factors, identified as the embedding agar solution. **(d)** The human user provides feedback on the reasons and the machine can continue learning. The human-revised deep network yields classifications matching biologically plausible strategies. (All shown RGB images correspond to real RGB images, while the edge overlays resulted from pseudo-RGB images generated from the original hyperspectral dataset, cf. Methods RGB/HS classification.)

Specifically, XIL takes the form illustrated in Fig. 1. In each step, the learner explains its interactive query to the domain expert, and she responds by correcting the explanations, if necessary, to provide feedback. This allows the user not only to check whether the model is right or wrong on the chosen instance but also if the answer is right (or wrong) for the wrong reasons, e.g., when there are ambiguities in the data such as confounders [4]. By witnessing the evolution of the explanations, similar to a teacher supervising the progress of a student, the human user can see whether the model eventually “gets it”. The user may even correct the explanation presented to guide the learner. This correction step is crucial for more directly affecting the learner’s beliefs and is integral to modulating trust [6, 12].

Specifically, we make the following contributions: (i) Introduction of XIL with counterexamples (CE) to revise “Clever Hans” behavior in a model-agnostic fashion. (ii) Adaption of the “right for the right reasons” (RRR) loss to latent layers of deep neural networks. (iii) Showcasing XIL on the computer vision benchmark datasets PASCAL VOC 2007 [14] and MSCOCO 2014 [15]. (iv) Evaluation of XIL on a highly relevant dataset for plant phenotyping, demonstrating its potential as an enabler of scientific discovery. (v) Gathering of the plant phenotyping dataset and the creation of a version with confounders. (vi) A user study on trust development within XIL [16].

A preliminary version of this manuscript has been published as a conference paper [17]. The present paper significantly extends the conference version by (ii-v). Moreover, the ad-hoc XIL user study presented in [17] was re-designed entirely, newly conducted, and now includes a thorough statistical analysis (vi). To encourage further research, we provide the created plant phenotyping dataset.

We proceed as follows. We start by formally introducing Explanatory Interactive Machine Learning (XIL) and instantiate it in the CAIPI method [17] as well as the RRR method [4]. After introducing XIL, we discuss quantitative results on test datasets before providing details on how domain experts can revise learning machines and, in turn, enable the machines to correct their abilities to solve the scientific real-world task of plant disease prediction. Finally, we demonstrate the importance of explaining decisions for building trustful machines via a user study. Thus, our contributions address a central part of building trustworthy AI methods by providing an end-to-end, interactive method to evaluate and revise machine learning models. This provides an important add-on to Rudin’s [18] message “*Stop explaining black-box machine learning models for high stakes decisions and use interpretable models instead*”, namely “*Start interacting with explanations of machine learning models to avoid ‘Clever Hans’-like behavior.*”

Explanatory Interactive Machine Learning (XIL)

In XIL, a learner can interactively query the user (or some other information source) to obtain the desired outputs of the data points. The interaction takes the following form. At each step, the learner considers a data point (labeled or unlabeled), predicts a label, and provides explanations of its prediction. The user responds by correcting the learner if necessary, providing a slightly improved—but not necessarily optimal—feedback to the learner.

Let us now instantiate this schema to *explanatory active learning*—combining active learning with local explainers (cf. Methods). Indeed, other interactive learning can be made explanatory too, including coactive learning [11], active imitation learning [19], and mixed-initiative interactive learning [20], but this is beyond the scope of this paper.

Explanatory Active Learning. In Explanatory Active Learning, we require black-box access to an active learner and an explainer. We assume that the active learner provides a procedure $\text{SELECTQUERY}(f, \mathcal{U})$ for selecting an informative instance $x \in \mathcal{U}$ based on the current model f , and a procedure $\text{FIT}(\mathcal{L})$ for fitting a new model (or update the current model) on the examples in $\mathcal{L}$. The explainer is assumed to provide a procedure $\text{EXPLAIN}(f, x, \hat{y})$ for explaining a particular prediction $\hat{y} = f(x)$. The framework is intended to work for any reasonable learner and explainer.

When using LIME for computing an interpretable model locally around the queries to visualize explanations for

Algorithm 1 XIL takes as input a set of labeled examples $\mathcal{L}$, a set of unlabeled instances $\mathcal{U}$, and iteration budget T .

```

1:  $f \leftarrow \text{FIT}(\mathcal{L})$ 
2: repeat
3:    $x \leftarrow \text{SELECTQUERY}(f, \mathcal{U})$ 
4:    $\hat{y} \leftarrow f(x)$ 
5:    $\hat{z} \leftarrow \text{EXPLAIN}(f, x, \hat{y})$ 
6:   Present  $x$ ,  $\hat{y}$ , and  $\hat{z}$  to the user
7:   Obtain  $\bar{y}$  and explanation correction  $\mathcal{C}$ 
8:   if CAIPI:
9:      $\{(\bar{x}_i, \bar{y})\}_{i=1}^c \leftarrow \text{TOCOUNTEREXAMPLES}(\mathcal{C})$ 
10:     $\mathcal{L} \leftarrow \mathcal{L} \cup \{(x, \bar{y})\} \cup \{(\bar{x}_i, \bar{y})\}_{i=1}^c$ 
11:   else if RRR:
12:      $\{(x, \bar{y}, A)\} \leftarrow \text{TOBINARYCORRECTIONMASK}(\mathcal{C})$ 
13:      $\mathcal{L} \leftarrow \mathcal{L} \cup \{(x, \bar{y}, A)\}$ 
14:    $\mathcal{U} \leftarrow \mathcal{U} \setminus \{x\}$ 
15:    $f \leftarrow \text{FIT}(\mathcal{L})$ 
16: until budget  $T$  is exhausted or  $f$  is good enough
17: return  $f$ 

```

current predictions, this results in CAIPI as summarized in Alg. 1. At each iteration $t = 1, \dots, T$ an instance $x \in \mathcal{U}$ is chosen using the query selection strategy implemented by the **SELECTQUERY** procedure. Then its label $\hat{y}$ is predicted using the current model f , and **EXPLAIN** is used to produce an explanation $\hat{z}$ of the prediction. The triple $(x, \hat{y}, \hat{z})$ is presented to the user as a (visual) artifact. The user checks the prediction and the explanation for correctness and provides the required feedback. Upon receiving the feedback, the system updates $\mathcal{U}$ and $\mathcal{L}$ accordingly and re-fits the model. The loop terminates when the iteration budget T is reached or the model is good enough.

During interactions between the system and the user, three cases can occur: **(1) Right for the right reasons:** The prediction and the explanation are both correct. No feedback is requested. **(2) Wrong for the wrong reasons:** The prediction is wrong. As in active learning, we ask the user to provide the correct label. While the explanation may provide some signal as to why the prediction was wrong, we currently do not require the user to act on it—this is an interesting avenue for future work—but treat the explanation to be simply wrong. **(3) Right for the wrong reasons:** The prediction is correct, but the explanation is wrong—the main target of XIL.

Model-agnostic XIL using counterexamples (CE). The “right for the wrong reasons” case is novel in active learning, and we propose *explanation corrections* to deal with it. They can assume different meanings depending on whether the focus is on component relevance, polarity, or relative importance (ranking), among others. In our experiments we ask the annotator to indicate the components that have been wrongly identified by the explanation as relevant, that is,

$$\mathcal{C} = \{j : |w_j| > 0 \wedge \text{the user believes the } j\text{th component to be irrelevant}\}.$$

Given the correction $\mathcal{C}$, we are faced with the problem of explaining it back to the learner. We propose a simple strategy to achieve this. This strategy is embodied by **TOCOUNTEREXAMPLES**. It converts $\mathcal{C}$ to a set of *counterexamples* that teach the learner not to depend on the irrelevant components. In particular, for every $j \in \mathcal{C}$ we generate c examples $(\bar{x}_1, \bar{y}_1), \dots, (\bar{x}_c, \bar{y}_c)$, where c is an application-specific constant. Here, the labels $\bar{y}_i$ are identical to the prediction $\hat{y}$. The instances $\bar{x}_i$, $i = 1, \dots, c$ are also identical to the query x , except that the j th component (i.e. $\psi_j(x)$) has been either randomized, changed to an alternative value, or substituted with the value of the j th component appearing in other training examples of the same class. This counterexample strategy (CE) produces $c \cdot |\mathcal{C}|$ counterexamples, which are added to $\mathcal{L}$, as summarized in Alg. 1. Importantly, this method is model-agnostic and can also be used when applying a non-differentiable model.

XIL using gradients. If the model is differentiable, the learner can also be regularized to be right for the right reasons using the “Right for the Right Reasons” loss (RRR) introduced by Ross *et al.* [4]. Here one adds a penalty to gradients that lie outside of a binary mask that indicates which features of the input are relevant. We modified the original loss function to:

$$L(\theta, X, y, A) = \underbrace{\sum_{n=1}^N \sum_{k=1}^K -c_k y_{nk} \log(\hat{y}_{nk})}_{\text{Right answers}} + \underbrace{\lambda_1 \sum_{n=1}^N \sum_{d=1}^D \left(A_{nd} \frac{\delta}{\delta h_{nd}} \sum_{k=1}^K c_k \log(\hat{y}_{nk}) \right)^2}_{\text{Right reasons}} + \underbrace{\lambda_2 \sum_i \theta_i^2}_{\text{Weight regularization}}, \quad (1)$$

(a) Fashion-MNIST (Toy) Dataset

(b) Scientific Dataset

(c) HS Scientific Dataset
non-confounded test set

	no corr.	Counterexamples			RRR IG		no. corr.	RRR GRAD-CAM	per-channel average	no corr.	RRR GRAD-CAM
Train	97%	93%	92%	92%	89%	RGB	89%	87%*	non-tissue	81%	87%
Test	48%	82%	85%	85%	85%	HS	99%	95%	full image	50%	82%

Table 1: **Explanatory feedback can boost trust and performance.** Highest performances are bold. (a) Accuracy on the fashion MNIST dataset of an MLP without corrections (no corr.), with our (CE) using varying c (middle), and RRR with input gradient (IG) constraints [4]. (b) The average model balanced accuracy of applying RRR with GRAD-CAM over five cross-validation runs. With “*” we denote situations where decisions made based on the background could not be fully removed. (c) The average model balanced accuracy over five cross-validation runs on a non-confounded test set of the hyperspectral (HS) scientific data. The confounding background features were set to either the per-channel average of the non-tissue regions or the full image of the training samples. The accuracies are reported for HS-CNN.

Fig. 2: **Examples of correcting Clever Hans moments with XIL.** XIL helps avoiding Clever Hans moments on unseen PASCAL VOC images (a). Ignoring user feedback, the model focuses on a source tag present in the lower left corner (b). Training it via interacting with its explanations, it does not consider the source tag to be relevant anymore (c). The visual explanations in (b, c) show relevant regions for the models decision using light and irrelevant ones using dark colors. Note that the images (from Pixabay) in (a) are shown for illustrative purposes; the original PASCAL VOC images to which the explanations in (b) and (c) correspond are not shown due to licensing issues but can be provided upon request for academic purposes.

where θ describes the parameters of the network, X the input, y the ground truth and A the binary mask used in the regularization term that discourages the input gradient from being large in regions marked by A . Instead of regularizing the gradients with respect to X , as originally described in [4], we regularize the gradients of the final convolutional layer h , corresponding to Gradient weighted Class Activation Maps (GRAD-CAM) ([21], *cf.* Methods). Further c is a rescaling weight given to each class of the unbalanced dataset and $\hat{y}$ corresponds to the network prediction. The objective function is split into three terms. The first and the last are the familiar cross-entropy and weight (θ) regularization terms. The second term is the new regularization term. The λ values are used to weight the different regularizations. Ross *et al.* [4] state that the regularization parameter λ_1 should be set such that the “right answers” and “right reasons” terms have similar orders of magnitude.

RRR can easily be incorporated into XIL (see again Alg. 1), and, as demonstrated by Selvaraju *et al.*’s HINT approach [22], Eq. 1 can be extended if the user is confident about how a (visual) explanation should look like. See Methods for an extended discussion of related work.

Demonstrating XIL on Computer Vision datasets. We begin by considering simulated users—as it is common for active learning—to evaluate the contribution of explanation feedback. Indeed, counterexample strategies (e.g., CAIPI) can trivially accommodate more advanced models than the one employed here. We simulate a human annotator that provides correct labels. Explanation corrections are also assumed to be correct and complete (i.e., they identify all false-positive components), for simplicity.

Specifically, we applied our data augmentation strategy to a decoy variant of fashion-MNIST [17], based on [23] (*cf.* Methods). The average test set accuracy of a multilayer-perceptron (with the same hyperparameters as in [4]) is reported in Table 1(a) for three correction strategies: no corrections, our CE, and the input-gradient constraints (RRR). The models’ explanations for CE are computed with LIME. Additionally, for every training image, we added $c = 1, 3, 5$ counterexamples where the decoy pixels are randomized. When no corrections are given, the accuracy on the test set is 48%: the confounders completely fool the network, *cf.* Table 1(b). Providing even a single counterexample increases the accuracy to 82%, i.e., the effect of confounders drops drastically. With more counterexamples, the accuracy of CE is similar to that of RRR. Both methods pose valid improvements, thus showing that explanatory interactive learning (XIL) is an effective mean for correcting “Clever Hans” moments in machine learning and may even improve predictive performance and beliefs.

Furthermore, we conducted experiments on the PASCAL VOC 2007 [14] dataset. We focused on a five-class subset (*cf.* Methods) and revised the model using XIL with the RRR loss. Fig. 2 presents some example images and their explanations with and without user feedback, i.e. default (test accuracy: 78%) and XIL trained (test accuracy: 73%). One can see that the classifier has learned the confounding factor for horse images (the source tag on the bottom left corner) without user feedback. After retraining the classifier using user feedback on the source tag location, we can see that the model no longer focuses on the confounder, demonstrating the benefit and effectiveness of XIL also in

Extended Data Figure 1: **Examples of XIL on MSCOCO 2014 dataset.** The left column (a) presents the original images, the middle column (b) presents the explanations (GRAD-CAM) after training without user feedback (default), the right column (c) presents the explanations after training with user feedback (XIL) using the MSE loss between user and model explanations. Also here, light regions represent relevant regions for the model’s decision, dark regions represent irrelevant regions. As user annotations we use the complete class segmentation to illustrate that XIL can also aid in improving the explanations for non-confounded data. See the Supplementary Information for more details. Due to license issues the presented images are alternatives to the original dataset.

Fig. 3: **Spectral signatures of measured agar plates with sugar beet leaf discs.** Signatures were extracted of agar on which healthy and inoculated sugar beet leaf discs were placed (a), of healthy and inoculated sugar beet leaf discs (b) and *C. beticola* symptoms of sugar beet leaves (c). Signatures were extracted from 100 pixels for each group, and the mean value is presented. The vertical (green, blue, red) lines correspond to the wavelength selected for the pseudo-RGB images.

this setting. Similar benefits can be observed on MSCOCO using HINT-like extensions (Extended Data Figure 3 and Supplementary Information). They may help to more quickly align human and gradient-based network explanations.

High predictive performance for deep plant phenotyping. Next, we showcase the extent, importance, and usability of XIL. To this end, we performed classification and revised corrections of the learned models on a real-world, scientific dataset. This dataset corresponds to RGB and hyperspectral (HS) (*cf.* Methods) images of leaf tissue from inoculated (*Cercospora beticola*) and healthy sugar beet plants. Notably, there is a strong variability in the extent of disease severity over all samples, with some samples clearly showing the characteristic of *Cercospora* Leaf Spot (CLS) (two rightmost samples in Fig. 1) while others do not (second to the left sample in Fig. 1) and for the human eye appear indistinguishable—at least in RGB—from healthy leaves (top sample in Fig. 1). Roughly 50% of inoculated tissue samples showed visible CLS.

We performed classification using convolutional neural networks (CNNs) on the RGB and HS datasets (*cf.* Methods). The task was to classify the leaf samples into one of the two classes: healthy or diseased. The corresponding average balanced accuracies determined over five cross-validation runs are shown in the left column (no corr.) of Table 1(b). They show high accuracies of 88% on the RGB dataset and nearly perfect performance of 99% on the HS dataset. It seems the HS data contains more relevant information to the classification task.

Model might be right for the wrong scientific reasons. The nearly perfect predictive performance is slightly suspicious since plant phenotyping is a rather difficult task. Therefore, we wanted to know the reasons for the predictions and visualized the network’s explanations using GRAD-CAMs. Specifically, we applied a spectral clustering and t-SNE [24] analysis, similar to [3], on the resulting explanations. Fig. 4(a) shows the strategies of the CNN trained on the HS data for data points belonging to the test set only. Supplementary Figure 1(a) shows the strategies of the CNN trained with RGB data. One can identify that the HS-CNN has altogether two prediction strategies, one for each predicted class label (*cf.* Supplementary Figure 4 for more details). In the case of control samples, the HS-CNN focuses on large areas of the tissue, however, for inoculated samples, even if CLS are visible, the network focuses on the nutritional solution (agar) to classify these as inoculated. Moreover, when analyzing the reflectance of the agar across different stages of disease development, we could indeed identify differences between control and inoculated nutrition solution. This can be seen in Fig. 3(a). Given the much smaller data dimensionality of the RGB images compared to the HS data, it seems likely that the RGB-CNN would have more difficulties focusing only on the agar as a classification feature, thus explaining the different classification strategies between HS and RGB-CNNs as well as the reduced classification performance of the RGB-CNN, compared to the HS-CNN.

In any case, both CNNs showed high to very high performances by largely using confounding factors within the dataset. The trained neural networks used strategies, which a biologist would consider as cheating rather than valid problem-solving behavior. The accuracies might not correspond to the true performance when measured in an environment outside of the lab setting, possibly even leading to dangerous consequences if left untackled.

Revising the model to be right for the right reasons. It is too simple to say that we can not trust these models and even question if machines are truly “intelligent”. We now show that with the human in the loop revising the machine, as in the XIL setting, the models can recover from the observed “Clever Hans”-like strategies towards trustful ones.

To this end, we let a plant biologist revise the machine by constraining the machine’s explanations to match her domain knowledge. Since the used models are differentiable, we focused on using RRR rather than using the CE strategy, though both would be valid within the XIL framework.

Fig. 4: **Cluster strategy analysis.** Analysis of the different decision strategies after training CNNs on the HS data with the cross-entropy loss (Default) in (a) and with the RRR loss in (b). The images are visualized in a two-dimensional t-SNE embedding and colored by the spectral clustering assignments.

Fig. 5: **Results of the user study on trust development.** (a) shows the total TiA Score over the three test conditions and (b) show in detail trust development (Q1) in correct rule learning after the three different learning stages of model accuracy (50%, 75%, 100%) for each test condition. Only statistically significant results are highlighted. The centerline of the box plots represents the median of the data, the box the interquartile distance between the first and third quartile, the whiskers the minimum and maximum value, discarding outliers which are plotted individually above the whiskers. The number of asterisks indicate the P values: * $P \leq 0.05$, ** $P \leq 0.01$, *** $P \leq 0.001$.

Specifically, we simulated the interaction between the domain experts and the ML models. After training a model without any interactions, plant physiologists analyzed the provided predictions and corresponding explanations. She decided that it is always a wrong reason to focus on the background, and consequently, her annotations corresponded to binary masks of the whole tissue (*cf.* Methods).

As before, we analyzed the decision strategies of the RRR trained model using t-SNE and spectral clustering. The results are summarized in Fig. 4(b) for the HS-CNN and Supplementary Figure 1(b) for the RGB-CNN.

As one can see, after training the HS-CNN with RRR, the model focuses on image regions lying only on the tissue, regardless of the underlying class. The control samples’ strategies correspond to nearly full activation of the whole tissue, whereas for inoculated samples, the identified relevant image regions are often very specific spots. Particularly, the model now focuses on the CLS, which were previously essentially ignored. Fig. 1(d) shows in more detail several examples of the observed strategies used by the corrected HS-CNN in comparison to the observed “Clever Hans” strategies of the unrevised machine. Although the model’s performance slightly decreased, *cf.* Table 1(b), it is still able to classify samples without visible symptoms. Even exploring different hyper-parameters for RRR, we were not able to force the RGB-CNN to ignore the background entirely, as illustrated in Supplementary Figure 1(b). As shown in Fig. 3(a), the HS-CNN has much more information at hand to focus on the confounding factors in the first place. However, after revision with RRR, it is easier for the HS-CNN to make accurate predictions based on the reflectance of the tissue in comparison to the RGB-CNN (Fig. 3(b)). Particularly, the HS-CNN mainly uses a spectral area for prediction, which is beyond the RGB area. This explains the difficulty of correcting the RGB-CNN.

We now focus on evaluating the default and revised models on a non-confounded test dataset to investigate the generalization improvement of training with XIL. Due to a missing non-confounded test set for the scientific dataset, we performed the simple trick of replacing the confounding features of all test samples with other values (*cf.* Methods). The results are summarized in Table 1(c), reporting the average test accuracy over five cross-validations. One can see that indeed the accuracy of the revised model is higher than that of the default model for both modifications. These results further indicate the generalization improvements due to XIL. Further experiments applying prior knowledge can be found in the Supplementary Information.

Trust development during XIL. After demonstrating that explanations and especially XIL are necessary to reveal and correct so-called “Clever Hans” behavior of ML models, we finally investigate how explanations influence the trust of users in the learning process. To this end, we designed a questionnaire about a machine that learns a simple concept by querying labels (but *not* explanation corrections) to an annotator. The online questionnaire was administered to 106 participants of varying ages and backgrounds.

Specifically, we designed a toy binary classification problem of (3×3) black-and-white images, inspired by the color dataset used in [4]. The subjects were told that an image is positive if the two top corners are white and negative otherwise. They were shown 30 images together with the classification of an AI model and a knowledgeable annotator. The learning of the model was simulated by increasing the model’s classification accuracy from 50% over 70% to 100% after every ten images. Each participant was randomly assigned to perform one of three experimental conditions with varying feedback from the model. In test condition 1 (**TC1**), the participant received feedback for each image in the form of the model’s prediction and the label provided by a knowledgeable annotator. No explanations were shown. Test conditions 2 and 3 (**TC2**, **TC3**) were identical to TC1, meaning that at every stage *the same example, prediction and feedback label* were shown, but now explanations were also provided. The explanations highlighted the two most relevant pixels in form of red dots. In TC2, the explanations converged to the correct rule—they highlight the two top corners—from the 6th image onwards. In TC3, the explanations converged to an incorrect rule—an image was classified as positive if the two top right pixels were white—from the 12th image onward. To assess the participant’s trust in the model’s skills, we used the Trust in Automation Questionnaire (TiA) [25]. After each learning process stage, the subjects were asked to rate (**Q1**) “I trust that the AI has learned the correct rule for classifying such images.”. Lastly, having seen all images, subjects were asked to answer the full TiA.

Fig. 5 summarizes the results, where (a) shows the total TiA score over TC1-TC3 and (b-d) the Q1 results for each test condition over the different stages of the learning process. They confirm previous findings: without explanations, people trust highly accurate machines, but the trust drops when wrong behavior is witnessed [6]. Users expect machines and their explanations to be correct. Indeed, explanations may increase the trust in earlier iterations at lower predictive performances, if they are correct. However, people do not forgive wrong explanations if the predictions are correct. Thus, users really care about the “right for the wrong reasons” case.

Taking all our empirical results together, people care about “Clever Hans”-like moments in machine learning, XIL can eliminate them, and XIL may even improve the models predictive performance.

Conclusion

In recent years, AI methods, especially machine learning with various directions and algorithms [26, 27], have become more and more successful in a wide range of areas like computer vision, natural language processing, and robotics, among others. Consider, for instance, AlphaZero surpassing human-level performance in playing chess and Go. During its self-play training process, AlphaZero discovered a remarkable level of Go knowledge. This included not only fundamental elements of human Go knowledge, but also non-standard strategies beyond the scope of traditional human Go knowledge [28], exemplifying the potential of these methods to discover strategies previously unknown even to experts of the domain. However, studies from various applications such as [29, 30, 31, 3] have revealed that learning machines can also result in “Clever Hans”-like moments, i.e., human-undesired strategies where the machine exploits artifacts in the dataset.

To “un-Hans” machines, we introduced the novel learning setting of “explanatory interactive learning” (XIL) and illustrated its benefits. XIL adds the scientist into the training loop. She interactively revises the original model via providing feedback on its explanations, used to automatically augment the training with counterexamples, or to modify the model using RRR. Our experimental results demonstrate that users care strongly about “Clever Hans”-like moments in machine learning, and XIL can indeed help avoiding them.

There are several possible avenues for future work to overcome the current limitations of XIL. Acquiring annotations, especially of explanations, can be time-consuming. The number of interactions required in order to reach an acceptable state is an open issue [17]. Hence, one should work on optimal query strategies for XIL that aim at minimizing the interaction efforts. Adapting regret bounds from coactive learning [11] might be an interesting alternative. Moreover, the data at hand may not always allow XIL to fully alleviate wrong reasons without decreasing the network’s predictive performance. One should develop ways to keep the drop as small as possible. Furthermore, XIL relies on two assumptions, namely, (a) faithful explanations can be computed, and (b) the user feedback is faithful, too. Assumption (a) is still subject to very active research, particularly for deep learning methods [32] (see our additional experiments on faithfulness with XIL in the Supplementary Information (Supplementary Figures 2 and 3)). One should improve the quality and robustness of XAI methods and also explore XIL for interpretable models [33]. If the user is rather confident about the right reasons, learning to explain methods such as HINT provide an interesting avenue for future work. Our initial results, see the Supplementary Information, are encouraging. However, even scientific experts do not always know the reasons for predictions. Therefore, one should strive to better understand the effects of wrong feedback and even adversarial attacks [34] on XIL. Additionally, one should turn other interactive learning settings such as coactive [11], active imitation [19], mixed-initiative interactive [20] and guided probabilistic learning [35] into explanatory one. Lastly, because it is not yet clear what makes explanations good for humans [36], one should extend explanatory interactions towards using alternative explanations, multiple modalities, and counterfactuals [37, 38]. In any case, interacting with explanations of machine learning models is an enabler for scientific discoveries for humans and machines in cooperation.

Methods

Active learning. The active learning paradigm targets scenarios where obtaining supervision has a non-negligible cost. Here we cover the basics of pool-based active learning, and refer the reader to two excellent surveys [39, 40] for more details. Let $\mathcal{X}$ be the space of instances and $\mathcal{Y}$ be the set of labels (e.g., $\mathcal{Y} = \{\pm 1\}$). Initially, the learner has access to a small set of labeled examples $\mathcal{L} \subseteq \mathcal{X} \times \mathcal{Y}$ and a large pool of unlabeled instances $\mathcal{U} \subseteq \mathcal{X}$. The learner is allowed to query the label of unlabeled instances (by paying a certain cost) to a user functioning as an annotator, often a human expert. Once acquired, the labeled examples are added to $\mathcal{L}$ and used to update the model. The overall goal is to maximize the model quality while keeping the number of queries or the total cost at a minimum. To this end, the query instances are chosen to be as informative as possible, typically by maximizing some informativeness criterion, such as the expected model improvement [41] or practical approximations thereof. By carefully selecting the instances to be labeled, active learning can enjoy much better sample complexity than passive learning [42, 43]. Prototypical active learners include max-margin [44] and Bayesian approaches [45]; recently, deep variants have been proposed [46]. However, active (showing query data points) and even coactive learning (showing the prediction of the query data point additionally) do not establish trust: informative selection strategies just pick instances where the model is uncertain and likely wrong. There is a trade-off between query informativeness and user “satisfaction”, as noticed and explored

in [47]. To properly modulate trust into the model, we argue it is essential to present explanations, e.g., visual ones, as shown in Extended Data Figure 1.

Local explainers. There are two main strategies for explaining machine learning models. Global approaches aim to explain the model by converting it *as a whole* to a more interpretable format [7],[48]. Local explainers instead focus on the arguably more approachable task of explaining *individual predictions* [9]. While explainable interactive learning can accommodate any local explainer, in our implementations, we used either LIME [8] or GRAD-CAM [21], both described next.

Extended Data Figure 2: **Example of explanations along the spatial and spectral dimensions.** GRAD-CAMs of a hyperspectral sample with spatial and spectral explanations of a corrected network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm and 854-1010 nm.

The idea of LIME (Local Interpretable Model-agnostic Explanations) is simple: even though a classifier may rely on many uninterpretable features, its decision surface around any given instance can be locally approximated by a simple, interpretable *local model*. In LIME, the local model is defined in terms of simple features encoding the presence or absence of *basic components*, such as words in a document or objects in a picture. While not all problems admit explanations in terms of elementary components, many of them do [8]; in this case, LIME assumes these to be provided in advance. An explanation can be readily extracted from such a model by reading off the contributions of the various components to the target prediction and translating them into an interpretable visual artifact. For instance, in document classification, one may highlight the words that support (or contradict) the predicted class.

GRAD-CAMs are a generalization of Class Activation Maps, introduced by [49], and take advantage of the facts that, firstly, deeper layers of a CNN capture higher-level visual constructs and, secondly, that convolutional features retain spatial information. As such, the last convolutional layer represents a trade-off between high visual representation and spatial information. Specifically, a GRAD-CAM is computed by forward passing an image through the network, applying a backpropagation of a one-hot encoding vector that specifies the class label of interest up to the last convolutional layer. The resulting gradients of each channel are global average pooled, multiplied with the corresponding feature maps, summed, and finally passed through a ReLU activation function. In this way, the final feature maps of the convolutional feature extractor are weighted by the importance of these features. The resulting two-dimensional heatmap can finally be interpolated to the original input size for visualization. In case a 3D convolutional network is used to classify hyperspectral data, the resulting heatmap is three dimensional, also showing activations along the spectral dimension of the data, *cf.* Extended Data Figure 1.

Explanatory Interactive Learning with counterexamples. Why is this data augmentation a sensible idea? To see this, consider the case of linear max-margin classifiers. Let $f(x) = \langle \mathbf{w}, \phi(x) \rangle + b$ be a linear classifier over two features, ϕ_1 and ϕ_2 , of which only the first is relevant. Extended Data Figure 2 shows that $f(x)$ (red line) uses ϕ_2 to correctly classify a positive example x_i . In order to obtain a better model (e.g. the green line), the simplest solution would be to enforce an orthogonality constraint $\langle \mathbf{w}, (0, 1)^\top \rangle = 0$ during learning. Counterexamples follow the same principle. In the separable case, the counterexamples $\{\bar{x}_{i\ell}\}_{\ell=1}^c$ amount to additional max-margin constraints [50] of the form $y_i \langle \mathbf{w}, \phi(\bar{x}_{i\ell}) \rangle \geq 1$. The only ones that influence the model are those on the margin, for which strict equality holds. For all pairs of such counterexamples ℓ, ℓ' it holds that $\langle \mathbf{w}, \phi(\bar{x}_{i\ell}) \rangle = \langle \mathbf{w}, \phi(\bar{x}_{i\ell'}) \rangle$, or equivalently $\langle \mathbf{w}, \delta_{i\ell} - \delta_{i\ell'} \rangle = 0$, where $\delta_{i\ell} = \phi(\bar{x}_{i\ell}) - \phi(x_i)$. In other words, the counterexamples encourage orthogonality between $\mathbf{w}$ and the correction vectors $\delta_{i\ell} - \delta_{i\ell'}$, thus approximating the orthogonality constraint above. Most importantly, this data augmentation procedure is model-agnostic, although alternatives indeed exist: (manually) adding a discovered data artifact to samples of other classes [51], contrastive examples [52], feature ranking [53] for SVMs and constraints on the input gradients for differentiable models [4].

We note that due to sampling, LIME may output different explanations for the same prediction. To reduce the variance of the experiments with CE of Table 1(a), we ran it ten times and retained the k components identified most often as relevant by LIME.

fashion-MNIST dataset. The fashion-MNIST dataset, a fashion product recognition dataset, includes 70,000 images over ten classes. All images were corrupted by introducing confounders, that is, 4×4 patches of pixels in randomly chosen corners whose shade is a function of the label in the training set and random in the test set (see [4] for details).

PASCAL VOC 2007 dataset. We used a subset of the PASCAL VOC 2007 dataset in our experiment. This subset includes 1470 train and 782 test images over five classes (horse, cat, bird, bus, dog). Only samples from the *horse* class contain confounding features, i.e., watermark text. We rescale all the images to $224 \times 224 \times 3$ to use the VGG-16 network [54] as a classifier, and we used the ImageNet-pre-trained weights as initial weights, as well as the ADAM optimizer

Extended Data Figure 3: **Mathematical intuition for the counterexample strategy, exemplified for linear classifiers.** Two data features are shown, ϕ_1 and ϕ_2 , of which only the first is truly relevant. (a) The positive example x_i is not enough to disambiguate between the red and green classifiers. (b) Counterexamples $x_{i,\ell}$ are obtained by randomizing the irrelevant feature while keeping the label of x_i . The counterexamples approximate a (local) orthogonality constraint. (c) The red classifier is inconsistent with the counterexamples and eliminated. See the Methods section *Explanatory Interactive Learning with counterexamples* for details. (Best viewed in color).

[55]. We trained a default model without user feedback and a model with user feedback for 2k epochs. The explanation method was instantiated with input gradients (IG).

Sample collection. To demonstrate the significance of XIL, we demonstrate XIL for deep plant phenotyping and plant disease detection, a growing and relevant field of research [56, 57, 58, 59, 60, 61]. To this end, we recorded a scientific, real-world dataset—a plant phenotyping dataset consisting of RGB and hyperspectral images (HS) of healthy and diseased sugar beet leaves. Then, we applied convolutional neural networks to classify the plants’ leaves into the categories *control* (healthy) and *inoculated* (diseased) and investigated the underlying reasons for the network’s predictions. As a model disease, Cercospora leaf spot (CLS) was used. This is caused by *Cercospora beticola* and is the most destructive leaf disease of sugar beet with worldwide economic importance.

The dataset used in this study corresponds to HS and RGB images of leaf discs of sugar beet cv. Isabella (KWS, Einbeck, Germany) inoculated with *Cercospora beticola*. Sugar beet seeds were pre-grown in small pots and piqued after the primary leaves were fully developed. The seedlings were then transferred into plastic pots (diameter of 17 cm) on a commercial substrate (Topfsubstrat 1.5, Balster Erdenwerk, GmbH, Sinntal-Altengronau, Germany) under greenhouse conditions and watered as necessary. After reaching growth stage 16 according to BBCH scale [62] the plants were inoculated with *C. beticola* conidia, which were collected from infested sugar beet leaves after incubation in a moist chamber for 48 hours. A spore suspension of 5×10^5 was sprayed onto leaves before the plants were transferred into plastic bags to achieve 100% RH for 48 hours. For image acquisition, leaf discs were stamped out with a 2 cm diameter cork borer and placed on 10g/l pythoagar (Duchefa Biochemie B.V, Haarlem, Netherlands), containing 0.34 mM benzimidazole, 10 g sucrose and 3 mg kinetin. To observe different symptom classes, sugar beet leaves of 9, 14, and 19 days after inoculation (dai) were used since the first symptoms appeared 9 dai. As a control group, 18 leaf discs of untreated sugar beet plants were measured as well, and five technical replications with six discs each were used for each symptom group.

Each sample, both control and inoculated, was measured daily over five consecutive days such that a sample from 9 dai reappears four further times in the dataset as 10 to 13 dai. A few samples were discarded due to technical issues. The percentage of healthy leaves to unhealthy leaves was approximately 26% to 74%, respectively. For image acquisition, leaf discs on agar were placed on a linear stage at a distance of 53 cm to a Hyperspec VNIR E-series imaging sensor (Headwall Photonics, Bolton, MA, USA) in the range of 380 nm to 1010 nm. The VNIR sensor has a spectral resolution of 2-3 nm and a pixel pitch of 6.5 μm . The sensor was surrounded by eight lamps (Ushio Halogen Lamp J12V-150WA/80 (Marunouchi, Chiyoda-ku, Tokyo, Japan)) and the distance between lamps and leaves was 60 cm with a vertical orientation of 45. Exposure times of 44 ms were used for the VNIR sensor.

The dataset consists of 2410 samples with 504 samples labeled as control and 1906 labeled as inoculated. Control samples were not re-used as inoculated samples. The collected hyperspectral raw data size was around 4 TB. After preprocessing the data by cutting out the leaf discs into hyperspectral cubes, the data has a size of 140 GB. Since there is much redundancy in the wavelength resolution, we further sub-sampled the depth of the data cubes resulting in a final data size of 32 GB.

Data preparation. As mentioned above, each sample was imaged over five consecutive days such that each sample, though slightly differing from day to day, is represented up to 5 times within the full dataset. In this way, a sample from 9 dai would occur for four further days (10-13 dai). To prevent the models from memorizing the structure of the individual leaf samples and correlating this to the corresponding labels, a precaution was taken to exclusively contain all days of one sample in either the training or validation dataset.

Removing confounders for the scientific dataset. It is essential to maintain the underlying assumption that the training and test data are drawn from the same distribution. If this is not the case, changes in accuracy might be due to artifacts of different data, rather than deficits of the model [63]. We applied two variations to the test samples of the HS dataset to remove the confounders: we set the background (everything but the plant tissue) (1) to the per-channel average of the non-tissue regions or (2) the per-channel average of the full images of the training data. We then evaluated the default and RRR revised CNNs on this modified test dataset. We focused here only on the HS data and model, due to the limitations of the RGB model’s performance.

RGB/HS classification. The RGB images used for training the classifiers were generated from the hyperspectral data, by slicing the data at the corresponding RGB channels that were provided by the camera system (*cf.* Fig. 1 (a-right)). Before training the RGB classifiers, the data was standard scaled following $z = (x - u)/s$, where u is the mean and s the standard deviation of the training samples.

To train a classifier on the RGB images of sugar beet leaves, we used a VGG-16 [54] network pre-trained on ImageNet [64] to finetune the network parameters using the RGB plant images. For training a batch size of 32, a learning rate of $1e-4$ and a step learning rate scheduler set to reduce the learning rate at epochs 5 and 15 by a factor of 0.1 were used. Furthermore, the ADAM optimizer was used with L2 regularization $1e-5$. Five separate cross-validation folds were trained until convergence, using a data split of 0.75 for training and 0.25 for testing. Convergence was reached after 30 epochs.

To classify the HS data, we trained a convolutional neural network (CNN) architecture with batch normalization using 3D convolution filters, rather than standard 2D filters, learning features not only along the image dimensions but also over the spectral dimensions. The used network is build up with four residual blocks, each containing one to three convolutional layers. The last two layers are fully connected layers with a final softmax activation function. The other layers use ReLU activations. During training the networks, we used dropout to prevent overfitting. The network’s parameters are trained with a stochastic gradient descent optimizer with momentum using a batch size of 10 HS images, a learning rate of $1e-4$, and an L2 regularization of $1e-5$.

Five separate cross-validation folds were trained until convergence, using a data split of 0.75 for training and 0.25 for testing. Convergence was reached after 100 epochs.

Analyzing classification strategies of the model. Based on the results of [32], in which the authors performed sanity checks over a variety of saliency methods, we chose to investigate our model’s explanations using Gradient-weighted Class Activation Mapping (GRAD-CAM) [21].

To analyze the resulting strategies produced by the layer-wise relevance propagation method (LRP), the authors of [3] revert to using spectral clustering on the resulting heatmaps in a pipeline they termed ‘SpRAY’. This clustering served to receive an overview of the extent of the model’s decision strategies. We apply SpRAY in a similar way. However, rather than using the raw GRAD-CAM heatmaps, we perform a discrete Fourier transformation on these beforehand to better differentiate different strategies that we had previously identified from single samples. In detail, the pipeline is as follows

- Perform a discrete Fourier transform on downsized GRAD-CAM heatmaps.
- Using the Euclidean distance for the RGB data and the Cityblock distance for the HS data compute a k-nearest neighbor graph of the Fourier transformed heatmaps, represented as an adjacency matrix, C .
- Compute the affinity matrix as suggested in [65] as $A = \max(C, C^T)$.
- Perform an eigengap analysis [65] to estimate the number of clusters, k , within the dataset.
- Perform spectral clustering on the affinity matrix, given k from the previous step
- Perform a t-SNE analysis [24] on the similarity matrix, estimated from the affinity matrix as in [3] as $S = \frac{1}{A+\epsilon}$, whereby $\epsilon \in [0, 1]$, here we used $\epsilon = 0.05$.

Applying XIL to CNNs for the scientific dataset. We produced the matrix A (Eq. 1) corresponding to full tissue masks for each sample. Specifically, for each sample, we created a binary mask having values of zero within the tissue and values of one everywhere else, i.e., the background. In this way, during training, the gradients everywhere but on the tissue are to be minimized.

The network models were retrained from the same initial values as in the default training mode (using only the cross-entropy loss), however, now using RRR. To choose the optimal λ_1 value, the resulting explanations were visually assessed. The five cross-validation folds of HS-CNN were thus trained until convergence between 200 and 280 epochs using a $\lambda_1 = 20$ value, with all other hyperparameters as in the default training mode. For training the RGB-CNN with RRR, the learning rate was reduced to a constant learning rate of $5e-05$. Although applying a range of λ_1 values from 0.1 to 1000, using the RGB-CNN, no satisfactory convergence state could be reached in which the regularized model showed acceptable explanations for each cross-validation run. The accuracies in Table 1(b-c) and the strategies presented in Fig. 4(b) and Supplementary Figure 1(b) correspond to GRAD-CAMs of training the five cross-validation folds with $\lambda_1 = 1$ for up to 200 epochs.

Extended related work. Using XIL with CE or RRR, users either introduce counterexamples into the dataset and thus teach the learner not to depend on the irrelevant components or directly penalize the learner as soon as it uses irrelevant components, respectively. One important advantage of XIL is that the user does not have to be certain about

the right reasons and instead can explore the learned reasons of the machine, in contrast to other procedures such as preprocessing the training set.

Recently, Selvaraju *et al.* [22] presented a framework (HINT) similar to RRR but instead of penalizing the wrong reasons it advises the network to use a specific visual area (right reasons). As CE and RRR, the HINT method could be embedded within the introduced XIL framework in case the users are certain about the right reasons. However, in many scientific applications such as the presented plant phenotyping dataset, users are uncertain about what a valid explanation should be. In this case, removing wrong reasons might be preferable to applying right reasons.

The possibility of bi-directional exchange between user and model due to interaction [66] also distinguishes XIL from approaches for feature selection, such as feature masking and approaches that embed prior knowledge into the training process, e.g., [67, 68]. Lastly, interactions also allow that the user can provide incomplete explanations. In other words: only if it is actually required, the user can revise incorrect aspects of a model’s explanation.

Finally, we present the XIL framework here for visual tasks and visual explanations only. With our definition of XIL, it is also applicable to other data domains like natural language processing, see e.g., [4, 17]. However, we experienced that explanations, i.e., right and wrong reasons, are more difficult to define for this modality. In future work, one should generally address the best ways to present explanations, even in multi-modal scenarios.

Details on participant recruitment and study procedure. The presented study is part of an extensive thesis work [16]. It was conducted as an online survey, the link of which was distributed via the social network Facebook and the forum of the student body of the department of computer science at TU Darmstadt. Due to the distribution on these channels, a wide range of people of different ages and different backgrounds was generated. Each participant completed only one of the three test conditions with 33 participants in TC1, 36 participants in TC2, and 37 participants in TC3, totaling 106 participants overall.

The wording of the original TiA was modified by replacing “system” with “artificial intelligence (AI)”. The response format to each question was a 5-point rating scale from strongly disagree to strongly agree.

Statistical analysis of the user study. Samples with missing values were removed from the analysis, and for all tests, a significance level with alpha being 5% was used.

For all tests with the same sample/samples, the alpha level was corrected via the Bonferroni-Holm method. The corrected alpha level will be stated for every analysis. For testing the hypotheses, one multi-factorial analysis of variances (MANOVA) and several one-factorial ANOVAs were conducted. The ANOVA, as well as the MANOVA, requires normal distribution of data, independence of data as well as homogeneity of the variances. To test the latter, a Levene-Test was conducted before every ANOVA and the MANOVA. Normal distribution was presumed due to the sample sizes, and as the samples were drawn randomly, the independence of data was also presumed. A significant result of an ANOVA / MANOVA means that at least two of the groups differ significantly with respect to the dependent variable, but it is not stated which groups differ. Therefore, if the carried out analyses of variances were significant, post hoc tests were carried out to investigate which groups differed exactly. Post hoc tests were selected in this study as the hypotheses did not point out which groups should differ, which is why every possible comparison had to be considered. For post hoc testing, the Tukey-HSD-Test and the Pairwise-Test were performed.

The TiA score of subjects being familiar with AI over the whole sample (all test conditions combined) was higher ($M = 2.82$, $SD = .64$) than the TiA score of subjects being unfamiliar with AI ($M = 2.51$, $SD = .59$). As the conducted Levene-Test ($F(5, 99) = 1.8$, $p = .12$, $\alpha = .05$) was not significant, the homogeneity of variance assumption held. Therefore, the MANOVA was conducted with a significant result for the independent variable test condition ($F(2, 99) = 10.10$, $p < .001$, $\alpha = .025$). The MANOVA was significant for the independent variable familiarity with AI ($F(1, 99) = 7.12$, $p = .009$, $\alpha = .025$). It was not significant for the interaction of the two independent variables ($F(2, 99) = .28$, $p = .75$, $\alpha = .025$).

For Fig. 5(a) in order to determine which test conditions differed significantly in their TiA score, a pairwise test was conducted as a post hoc test. The pairwise test showed significant differences between TC1 and TC3 ($p = .0016$, $\alpha = .05$) as well as between TC2 and TC3 ($p = .0003$, $\alpha = .05$).

For Fig. 5(b) the conducted Levene-Test was not significant ($F(2, 96) = .59$, $p = .56$, $\alpha = .05$). Therefore, an ANOVA was conducted afterwards and showed a significant result ($F(2, 96) = 33.83$, $p < .001$, $\alpha = .0125$). Trust in the correct rule learning by the AI was significantly different between the blocks. The conducted Tukey-HSD test found a significant difference in trust into the correct rule learning only between stage 1 and 3 ($p < .001$, $\alpha = .05$) and between stage 2 and 3 ($p < .001$, $\alpha = .05$).

For Fig. 5(c) the Levene-Test was not significant ($F(2, 104) = .28$, $p = .75$, $\alpha = .05$). The ANOVA was significant ($F(2, 104) = 23.19$, $p < .001$, $\alpha = .0167$). Therefore, a Tukey-HSD test was performed to investigate which blocks differed significantly. The test found only stage 1 and 3 ($p < .001$, $\alpha = .05$) and stage 2 and 3 ($p < .001$, $\alpha = .05$) to differ significantly with respect to trust in correct rule learning by the AI.

For Fig. 5(d) the conducted Levene-Test was not significant ($F(2, 105) = 1.32$, $p = .27$, $\alpha = .05$). The afterwards conducted ANOVA was also not significant ($F(2, 105) = 1.62$, $p = .20$, $\alpha = .05$). Therefore, there was no significant difference in trust into correct rule learning by the AI in TC3 and no post hoc test was performed.

500 Data availability

501 The ML benchmark Fashion-MNIST is available at <https://github.com/zalandoresearch/fashion-mnist>. The
502 PASCAL VOC2007 dataset is available at <http://host.robots.ox.ac.uk/pascal/VOC/voc2007/>. The RGB and
503 hyperspectral data that support the findings of this study are available in the code repository [https://doi.org/10.
504 24433/CO.4559958.v1](https://doi.org/10.24433/CO.4559958.v1) [13]. The user study is available at [https://github.com/ml-research/xil/tree/master/
505 Trust_Study](https://github.com/ml-research/xil/tree/master/Trust_Study).

506 Code availability

507 The code and a fully runnable capsule to reproduce the figures and results of this article, including pre-trained models,
508 can be found at <https://doi.org/10.24433/CO.4559958.v1> [13].

509 Statement of ethical compliance

510 The authors confirm to have complied with all relevant ethical regulations, according to the Ethics Commission of the
511 TU Darmstadt (<https://www.intern.tu-darmstadt.de/gremien/ethikkommission/auftrag/auftrag.en.jsp>). An
512 informed consent was obtained for each participant prior to commencing the user study.

513 Acknowledgments

514 ST and KK thank Antonio Vergari, Andrea Passerini, Samuel Kolb, Jessa Bekker, Xiaoting Shao, and Paolo Morettin for
515 very useful feedback on the conference version of this article. Furthermore, the authors are thankful to Frank Jäkel for
516 support and supervision on the user study, to Cigdem Turan for providing the figure sketches, and to Ulrike Steiner and
517 Stefan Paulus for very useful feedback. PS, AKM, AB and KK acknowledge the support by BMEL funds of the German
518 Federal Ministry of Food and Agriculture (BMEL) based on a decision of the Parliament of the Federal Republic
519 of Germany via the Federal Office for Agriculture and Food (BLE) under the innovation support program, project
520 “DePhenSe” (FKZ 2818204715). WS and KK were also supported by BMEL/BLE funds under the innovation support
521 program, project “AuDiSens” (FKZ 28151NA187). ST acknowledges the supported by the European Research Council
522 (ERC) under the European Unions Horizon 2020 research and innovation programme, grant agreement No. [694980]
523 “SYNTH: Synthesising Inductive Data Models”. XS and KK also acknowledges the support by the German Science
524 Foundation project “CAML” (KE1686/3-1) as part of the SPP 1999 (RATIO). AKM was partially funded by the
525 Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germanys Excellence Strategy - EXC
526 2070 – 390732324

527 Conflict of interest statement

528 The authors declare the following competing interests: HS is employed by LemnaTec GmbH.

529 Author information

530 Affiliations

531
532 Technical University of Darmstadt, Computer Science Department, Artificial Intelligence and Machine Learning Lab,
533 Darmstadt, Germany

534 Patrick Schramowski, Wolfgang Stammer, Franziska Herbert, Xiaoting Shao

535
536 Technical University of Darmstadt, Computer Science Department and Centre for Cognitive Science, Darmstadt, Ger-
537 many

538 Kristian Kersting

539
540 University of Trento, Department of Information Engineering and Computer Science, Trento, Italy
541 Stefano Teso

542
543 University of Bonn, Institute of Crop Science and Resource Conservation (INRES) – Plant Diseases and Plant Pro-
544 tection, Bonn, Germany

545 Anna Brugger

Institute of Sugar Beet Research, Goettingen, Germany
Anne-Katrin Mahlein
LemnaTec GmbH, Aachen, Germany
Hans-Georg Luigs

Author Contributions

PS and WS contributed equally to the work. PS, WS, ST, KK designed the study. ST, KK designed and published (AAAI /ACM Conference on Artificial Intelligence, Ethics, and Society 2019) the preliminary version of this manuscript. PS, WS, XS, ST, and KK developed extensions of the basic XIL methods. PS, WS, AB, AKM, and KK interpreted the data and drafted the manuscript. AB and PS designed the phenotyping dataset. AB and HGL carried out the phenotyping dataset measuring. PS, WS, AB did the biological analysis. FH performed and analyzed the user study. AKM and KK directed the research and gave initial input. All authors read and approved the final manuscript.

Corresponding author

Correspondence to Patrick Schramowski and Wolfgang Stammer.

References

- [1] Guidotti, R. *et al.* A survey of methods for explaining black box models. *ACM computing surveys (CSUR)* **51**, 1–42 (2018).
- [2] Gilpin, L. H. *et al.* Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE International Conference on data science and advanced analytics (DSAA)*, 80–89 (2018).
- [3] Lapuschkin, S. *et al.* Unmasking clever hans predictors and assessing what machines really learn. *Nature communications* **10**, 1096 (2019).
- [4] Ross, A. S., Hughes, M. C. & Doshi-Velez, F. Right for the right reasons: Training differentiable models by constraining their explanations. In *Proceedings of International Joint Conference on Artificial Intelligence*, 2662–2670 (2017).
- [5] Simpson, J. A. Psychological foundations of trust. *Current directions in psychological science* **16**, 264–268 (2007).
- [6] Hoffman, R. R., Johnson, M., Bradshaw, J. M. & Underbrink, A. Trust in automation. *IEEE Intelligent Systems* **28**, 84–88 (2013).
- [7] Buciluă, C., Caruana, R. & Niculescu-Mizil, A. Model Compression. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD*, 535–541 (2006).
- [8] Ribeiro, M. T., Singh, S. & Guestrin, C. Why should I trust you?: Explaining the predictions of any classifier. In *Proceedings of ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144 (ACM, 2016).
- [9] Lundberg, S. & Lee, S. An unexpected unity among methods for interpreting model predictions. *CoRR abs/1611.07478* (2016). URL <http://arxiv.org/abs/1611.07478>.
- [10] Settles, B. Closing the loop: Fast, interactive semi-supervised annotation with queries on features and instances. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 1467–1478 (Association for Computational Linguistics, 2011).
- [11] Shivaswamy, P. & Joachims, T. Coactive learning. *Journal of Artificial Intelligence Research* **53**, 1–40 (2015).
- [12] Kulesza, T. *et al.* Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of International Conference on Intelligent User Interfaces*, 126–137 (2015).
- [13] Schramowski, P., Stammer, W., Teso, S., Herbert, F. Making deep neural networks right for the right scientific reasons by interacting with their explanations. <http://https://doi.org/10.24433/CO.4559958.v1>.

- [14] Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J. & Zisserman, A. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- [15] Lin, T. *et al.* Microsoft COCO: common objects in context. In *Proceedings of European Conference on Computer Vision*, 740–755 (2014).
- [16] Herbert, F. P., Kersting, K. & Jäkel, F. *Why Should I Trust in AI?* Master’s thesis, Technical University Darmstadt (2019).
- [17] Teso, S. & Kersting, K. Explanatory interactive machine learning. In *Proceedings of AAAI/ACM Conference on AI, Ethics, and Society* (AAAI, 2019).
- [18] Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* **1**, 206–215 (2019).
- [19] Judah, K. *et al.* Active imitation learning via reduction to iid active learning. In *AAAI Fall Symposium Series* (2012).
- [20] Cakmak, M. *et al.* Mixed-initiative active learning. *ICML 2011 Workshop on Combining Learning Strategies to Reduce Label Cost* (2011).
- [21] Selvaraju, R. R. *et al.* Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, 618–626 (2017).
- [22] Selvaraju, R. R. *et al.* Taking a hint: Leveraging explanations to make vision and language models more grounded. In *Proceedings of the IEEE International Conference on Computer Vision*, 2591–2600 (2019).
- [23] Xiao, H., Rasul, K. & Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017).
- [24] Maaten, L. v. d. & Hinton, G. Visualizing data using t-SNE. *Journal of machine learning research* **9**, 2579–2605 (2008).
- [25] Körber, M. Theoretical considerations and development of a questionnaire to measure trust in automation. In *Congress of the International Ergonomics Association*, 13–30 (Springer, 2018).
- [26] Jordan, M. I. & Mitchell, T. M. Machine learning: Trends, perspectives, and prospects. *Science* **349**, 255–260 (2015).
- [27] Ghahramani, Z. Probabilistic machine learning and artificial intelligence. *Nature* **521**, 452–459 (2015).
- [28] Silver, D. *et al.* Mastering the game of go without human knowledge. *Nature* **550**, 354–359 (2017).
- [29] Zech, J. R. *et al.* Confounding variables can degrade generalization performance of radiological deep learning models. *CoRR* **abs/1807.00431** (2018). URL <http://arxiv.org/abs/1807.00431>.
- [30] Badgeley, M. A. *et al.* Deep learning predicts hip fracture using confounding patient and healthcare variables. *npj Digital Medicine* **2**, 31 (2019).
- [31] Chaibub Neto, E. *et al.* A permutation approach to assess confounding in machine learning applications for digital health. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 54–64 (ACM, 2019).
- [32] Adebayo, J. *et al.* Sanity checks for saliency maps. In *Proceedings of Advances in Neural Information Processing Systems*, 9505–9515 (2018).
- [33] Chen, C. *et al.* This looks like that: Deep learning for interpretable image recognition. In *Proceedings of Advances in Neural Information Processing Systems*, 8928–8939 (2019).
- [34] Dombrowski, A. *et al.* Explanations can be manipulated and geometry is to blame. In Wallach, H. M. *et al.* (eds.) *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, 13567–13578 (2019).
- [35] Odom, P. & Natarajan, S. Human-guided learning for probabilistic logic models. *Frontiers in Robotics and AI* **5**, 56 (2018).

- [36] Narayanan, M. *et al.* How do humans understand explanations from machine learning systems? an evaluation of the human-interpretability of explanation. *CoRR* **abs/1802.00682** (2018). URL <http://arxiv.org/abs/1802.00682>.
- [37] Kanehira, A. & Harada, T. Learning to explain with complemental examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8603–8611 (2019).
- [38] Huk Park, D. *et al.* Multimodal explanations: Justifying decisions and pointing to the evidence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8779–8788 (2018).
- [39] Settles, B. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* **6**, 1–114 (2012).
- [40] Hanneke, S. *et al.* Theory of disagreement-based active learning. *Foundations and Trends® in Machine Learning* **7**, 131–309 (2014).
- [41] Roy, N. *et al.* Toward optimal active learning through monte carlo estimation of error reduction. *ICML* 441–448 (2001).
- [42] Castro, R. M. *et al.* Upper and lower error bounds for active learning. In *Proceedings of Conference on Communication, Control and Computing*, 2.1, 1 (2006).
- [43] Balcan, M.-F. *et al.* The true sample complexity of active learning. *Machine learning* **80**, 111–139 (2010).
- [44] Tong, S. & Koller, D. Support vector machine active learning with applications to text classification. *Journal of machine learning research* **2**, 45–66 (2001).
- [45] Krause, A. *et al.* Nonmyopic active learning of gaussian processes: an exploration-exploitation approach. In *Proceedings of International Conference on Machine learning*, 449–456 (ACM, 2007).
- [46] Gal, Y. *et al.* Deep bayesian active learning with image data. In *Proceedings of International Conference on Machine learning*, 1183–1192 (2017).
- [47] Schnabel, T. *et al.* Short-term satisfaction and long-term coverage: Understanding how users tolerate algorithmic exploration. In *Proceedings of ACM International Conference on Web Search and Data Mining*, 513–521 (ACM, 2018).
- [48] Bastani, O., Kim, C. & Bastani, H. Interpreting blackbox models via model extraction. *CoRR* **abs/1705.08504** (2017). URL <http://arxiv.org/abs/1705.08504>.
- [49] Zhou, B., Khosla, A., Lapedriza, A., Oliva, A. & Torralba, A. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929 (2016).
- [50] Cortes, C. *et al.* Support-vector networks. *Machine learning* **20**, 273–297 (1995).
- [51] Anders, C. J. *et al.* Analyzing imagenet with spectral relevance analysis: Towards imagenet un-hans’ed. arXiv preprint arXiv:1912.11425 (2019).
- [52] Zaidan, O. *et al.* Using “annotator rationales” to improve machine learning for text categorization. In *Proceedings of Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 260–267 (2007).
- [53] Small, K. *et al.* The constrained weight space svm: learning with ranked features. In *Proceedings of International Conference on Machine learning*, 865–872 (Omnipress, 2011).
- [54] Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *Proceedings of International Conference on Learning Representations* (2015).
- [55] Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. In *Proceedings of International Conference on Learning Representations* (2015).
- [56] Lau, E. High-throughput phenotyping of rice growth traits. *Nature Reviews Genetics* **15**, 778–778 (2014).
- [57] de Souza, N. High-throughput phenotyping. *Nature Methods* 36–36 (2009).
- [58] Tardieu, F., Cabrera-Bosquet, L., Pridmore, T. & Bennett, M. Plant Phenomics, From Sensors to Knowledge. *Current Biology* **27**, R770–R783 (2017).
- [59] Pound, M. P. *et al.* Deep machine learning provides state-of-the-art performance in image-based plant phenotyping. *Gigascience* **6**, gix083 (2017).

- 685 [60] Mochida, K. *et al.* Computer vision-based phenotyping for improvement of plant productivity: a machine learning
686 perspective. *GigaScience* **8**, giy153 (2018).
- 687 [61] Mahlein, A.-K. *et al.* Quantitative and qualitative phenotyping of disease resistance of crops by hyperspectral
688 sensors: seamless interlocking of phytopathology, sensors, and machine learning is needed! *Current opinion in*
689 *Plant Biology* **50**, 156–162 (2019).
- 690 [62] Meier, U. *et al.* Phenological growth stages of sugar beet (*Beta vulgaris* l. ssp.) codification and description according
691 to the general bbch scale (with figures). *Nachrichtenblatt des Deutschen Pflanzenschutzdienstes* **45**, 37–41 (1993).
- 692 [63] Hooker, S., Erhan, D., Kindermans, P. & Kim, B. A benchmark for interpretability methods in deep neural
693 networks. In *Proceedings of Advances in Neural Information Processing Systems*, 9734–9745 (2019).
- 694 [64] Deng, J. *et al.* ImageNet: A Large-Scale Hierarchical Image Database. In *Proceedings of IEEE Conference on*
695 *Computer Vision and Pattern Recognition* (2009).
- 696 [65] Von Luxburg, U. A tutorial on spectral clustering. *Statistics and computing* **17**, 395–416 (2007).
- 697 [66] Abdel-Karim, B. M., Pfeuffer, N., Rohde, G. & Hinz, O. How and what can humans learn from being in the loop?
698 *KI-Künstliche Intelligenz* 1–9 (2020).
- 699 [67] Erion, G. G., Janizek, J. D., Sturmfels, P., Lundberg, S. & Lee, S. Learning explainable models using attribution
700 priors. *CoRR* **abs/1906.10670** (2019). URL <http://arxiv.org/abs/1906.10670>.
- 701 [68] Liu, F. & Avci, B. Incorporating priors with feature attribution on text classification. In *Proceedings of the 57th*
702 *Annual Meeting of the Association for Computational Linguistics*, 6274–6283 (2019).

Supplementary Materials

Strategy analysis of RGB scientific dataset

Figure 1: Cluster analysis of the different decision strategies after training CNNs on the RGB data with the cross-entropy loss (Default) in (a) and with the RRR loss in (b). The images are visualized in a two-dimensional t-SNE embedding and colored by the spectral clustering assignments.

Fig. 1(a) shows the strategies of the CNN trained on the RGB data for data points only in the test set. When CLS were visible the RGB-CNN correctly identifies these as relevant features for classifying the samples as inoculated. However, for many inoculated samples, for which no spots are visible the CNN surprisingly focuses on regions in the background, specifically often on the nutrition solution (agar), which the tissue was embedded in. Also for healthy samples the RGB-CNN focuses on the background.

One can identify different decision strategies, after training with RRR as illustrated in Fig. 1(b). However, even using different hyper-parameters for RRR, we were not able to reach a converged state such that the RGB-CNN fully ignores the background.

Faithfulness of learned explanations

Investigating the faithfulness of an explanation method is a very valid and relevant topic of research ([1], [2], [3]). From our experiments the objection can be made that the revised models have merely learned to produce acceptable explanations, but still focus on wrong features.

Although we have shown the generalization performance in previous experiments using non-confounded test sets, we ran further experiments to investigate the faithfulness of explanations that have been revised using the XIL framework. The questions we wanted to answer were the following: (Q1) are the features learned interactively using XIL more relevant for the original task than the identified features of the default model? We note that an important underlying assumption here is that the user feedback is correct and faithful. (Q2) Is the XIL revised model more strongly influenced by it’s learned explanations in comparison to the default model with it’s unrevised explanations?

We focused on using a more widely used dataset of the ML and computer vision community: MSCOCO 2014 [4]. This dataset presents a multi-label image classification problem of commonly found objects and is completed with a masked segmentation for each class of each sample. For the following experiments we used a subset of the COCO14 classes, focussing on the five classes: elephant, giraffe, cat, dog and truck.

As the MSCOCO dataset is a non-confounded dataset, the task when using XIL with this dataset is therefore to mainly improve the model to focus on right reasons, rather than penalizing it when focusing on wrong reasons. A characteristic of CE and RRR, is that both methods revise an ML model when it is using wrong features for a right prediction, but not when it is *not* using a feature for a right prediction. More specifically, a user might want to direct a model’s attention to features that she finds very relevant (similar to Selvaraju *et al.*’s HINT method [5]). For this reason we adapted the RRR loss of Eq. 1 in the main text to the following, resulting in a HINT-like extension:

$$L(\theta, X, y, A) = \underbrace{\sum_{n=1}^N \sum_{k=1}^K -c_k y_{nk} \log(\hat{y}_{nk})}_{\text{Right answers}} + \underbrace{\lambda_1 \sum_{n=1}^N \sum_{d=1}^D (A_{nd} - \text{expl}(y_n))^2}_{\text{Learning explanations}} + \underbrace{\lambda_2 \sum_i \theta_i^2}_{\text{Weight regularization}}. \quad (1)$$

The notations are largely the same as for Eq. 1 of the main text. $\text{expl}(y_n)$ represents the specific differentiable explanation method that produces an explanation in the dimensions of the input space given the true class label y_n . In comparison to the RRR loss used in previous experiments the matrix A now denotes a value of 1 for features that the user finds relevant, -1 for irrelevant and 0 for features, where the user is indifferent. In our experiments we only used values of 1 and -1 . Additionally, to properly compute the difference between the user and model explanation we rescaled the model explanation to the range $[0, 1]$, thus additionally enforcing the model to ignore irrelevant regions.

For the following experiments we again focused on the GRAD-CAM explanation method. In this case we downscaled the user annotations to the original output dimensions of the GRAD-CAM prior to computing Eq. 1. The user interaction was again simulated, whereby the user annotations corresponded to the ground truth class segmentations provided with the COCO dataset.

Fig. 2 shows several example images (Fig. 2(a)) for which the default explanations are partly correct (Fig. 2(b)). However, it would be valid for a user to be unsatisfied with these explanations, given that only small regions of the to-be-predicted objects are highlighted. These results highlight that even the GRAD-CAM method produces explanations that a human user might not fully accept. With XIL in the form of Eq. 1 these explanations could be refined to coincide more with the user’s explanations (Fig. 2(c)). we note that the default model was trained for as many iterations as the XIL model.

To answer Q1 we applied the method of [6], termed “Remove and Retrain” (ROAR). The idea here is to investigate indeed how relevant the features are that different explanation methods have deemed as important. This is done by removing a certain percentage of relevant features that an explanation method has identified, set these features to the mean of the training data and retrain a model from an initial parameter setting. If the model produces a low prediction performance this is an indication that the features of the explanation method are indeed relevant for the task. If the performance is high this is an indication that there are equally or more relevant features available for the task.

Fig. 3 shows the results of ROAR where the initial on ImageNet-pretrained VGG-16 [7] was retrained until convergence using the modified datasets. This was repeated for random explanations as a baseline, the default trained model GRAD-CAM explanations and the XIL revised GRAD-CAM explanations. One can indeed observe that given the assumption of relevant and faithful user feedback, with XIL it is possible for a differentiable model to improve its explanations to focus on more relevant features.

With the previous experiment we could show that with a human in the loop a model can be revised to focus on more relevant features, which accord more strongly with the user’s explanations, even if the model’s original explanations were not considered as entirely wrong. ROAR, however, was developed to test explanation methods which were not explicitly trained to improve their explanations. This is different in the XIL setting. Due to that, for ROAR, the same model is retrained over all conditions, we have not yet shown, that the revised model actually

Figure 2: Several examples from MSCOCO 2014. The left column (a) presents the original images, the middle column (b) presents the explanations (GRAD-CAM) after training without user feedback (default), the right column (c) presents the explanations after training with user feedback (XIL). Also here, as in the main text, light regions represent relevant regions for the model’s decision, dark regions represent irrelevant regions. Here the user annotations were the complete class segmentations to illustrate that XIL can also aid in improving the explanations for non-confounded data.

focuses more strongly on it’s learned explanations in comparison to the default model which had not optimized it’s explanations. In other words it remains open to show that the explanations of the XIL revised model are more faithful to the model’s decisions than the explanations of the default model are to the default model.

We therefore evaluated both models, the default and XIL revised model, on the test set where, similar to the ROAR experiment, we replaced a certain percentage of relevant features with the per channel mean. Particularly each model (default and XIL revised) was evaluated on the test set, where features were removed based on their explanations. Importantly this was set in comparison to evaluation on the test set, where random features had been removed.

The results can be found in Fig. 4, where also here a lower accuracy indicates a feature’s importance for the specific model. One can observe that there is little difference between the evaluations of both models on the random-explanation-modified test set (baseline). There are however strong differences between evaluations on the test sets modified by their respective explanations, where the accuracy strongly drops for the XIL revised model based on its explanations, than the default model, even when taking the difference between baseline evaluations into account. Thus indicating that the learned explanations of the XIL revised model are more faithful to the model’s decisions.

Figure 3: Three-fold cross-validation ROAR results for the MSCOCO 2014 dataset. A model was trained from an initial parameter setting with altered datasets where a certain percentage of the most relevant features were removed in the training and test set. The relevance of each feature is indicated by the explanations of the XIL trained and default trained model as well as random explanations. The lower the accuracy, the more likely it is that the removed features are informative for the original model.

Figure 4: Three-fold cross validation of default and XIL trained models evaluated on the test set. Varying percentages of relevant pixels are removed from the test set, whereby the relevance is indicated by the explanations from the trained models (default, or XIL) or random pixel assignments.

Data augmentation and training based on prior knowledge

To compare to the setting of removing the confounders in the datasets based on prior knowledge before training, we train a model based on variants without confounders and test the resulting model on both variants, the original dataset (with confounders) and the dataset with removed confounders (w/o confounders). We here focused on the

HS data setting.

To remove confounders in the decoy fashion MNIST dataset we reverted to the original dataset (w/o confounders). Regarding the scientific plant phenotyping dataset, no variant without confounders is available. Therefore, we removed everything that is not belonging to the user annotated tissue and replaced it with the average reflectance of each hyperspectral channel. The trained models were tested on both dataset variants (w/o confounders and with confounders).

For both datasets, one can observe a similar behavior. Training with the XIL framework and training the model on the dataset variant w/o confounders results in a similar score when testing the model on the dataset without confounders. However, when testing on the original dataset one can observe that the model trained without confounders has an extreme accuracy drop. Instead, the models trained with the XIL framework generalized to not using confounders and perform well on both variants. The accuracy drop of XIL applied to the dataset without confounders can be explained by the change of data distributions. By a stronger weight of the right reasons one might be able to correct this even further.

	Fashion-MNIST		Scientific Dataset (HS)	
	trained w/o confounders	XIL	trained w/o of confounders	XIL
w/o confounders	85%	84%	86%	82%
with confounders	76%	85%	56%	95%

Table 1: Comparison of training on augmented data based on prior knowledge and XIL. The trained models were tested on both dataset variants (w/o confounders and with confounders).

Questionnaire and example online survey used in User Study

The competence of a classifier can be assessed by monitoring its behavior and beliefs over time, directability can be achieved by allowing the user to actively teach the model how to act and what to believe, while understandability can be approached by explaining the models decisions. To investigate how interaction with a machine and augmenting this interaction with explanations influences the trust of the user into the model we designed a questionnaire based on a binary classification toy problem. The Questionnaire document, as well as an example of the online survey for TC2, can be found at <https://github.com/ml-research/XIL>

Strategy analysis classification errors

Fig. 5 shows the class prediction versus underlying (ground truth) class for each sample of all four CNN training versions (from upper left to bottom right: default RGB-CNN, revised RGB-CNN, default HS-CNN, revised HS-CNN). Each sample is plotted in the embedding of Fig. 3 and Fig. 4 (main article). Particularly the HS-CNN, regardless of the training configuration, shows strong differences in the decision strategies for the two classes.

Explanations along hyperspectral data dimension

Below are several detailed examples of the explanations from the HS-CNNs. Figures 2-9 present the explanations from the default HS-CNN, showing samples from varying stages of disease progression (see captions for details). Figures 10-17 present the explanations from the revised HS-CNN, showing samples from varying stages of disease progression (see captions for details). Each Figure 2-17 depicts the sample in the leftmost panel, followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Default RGB-CNN Strategy Analysis

Human-revised RGB-CNN
Strategy Analysis

Default HS-CNN Strategy Analysis

Human-revised HS-CNN
Strategy Analysis

Figure 5: T-SNE embeddings of the different image type and training loss configurations colored by the ground truth (gt) and prediction (pred) labels of each sample. The top row depicts the results of training with RGB images, the bottom row with hyperspectral images. The left column shows the results of training only with the cross-entropy loss, whereas the right column shows the results of training with the right for the RRRLoss.

Figure 6: GRAD-CAMS with spatial and spectral activations from healthy sample of unrevised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 7: GRAD-CAMS with spatial and spectral activations from healthy sample of unrevised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 8: GRAD-CAMS with spatial and spectral activations from healthy sample of unrevised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 9: GRAD-CAMS with spatial and spectral activations from healthy sample of not regularized network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 10: GRAD-CAMS with spatial and spectral activations from inoculated sample with single visible symptoms of unrevised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 11: GRAD-CAMS with spatial and spectral activations from inoculated sample with single visible symptoms of unrevised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 12: GRAD-CAMS with spatial and spectral activations from inoculated sample with multiple visible symptoms of not regularized network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 13: GRAD-CAMS with spatial and spectral activations from inoculated sample with multiple visible symptoms of unrevised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 14: GRAD-CAMS with spatial and spectral activations from healthy sample of revised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 15: GRAD-CAMS with spatial and spectral activations from healthy sample of revised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 16: GRAD-CAMS with spatial and spectral activations from inoculated sample without visible symptoms of revised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 17: GRAD-CAMS with spatial and spectral activations from inoculated sample without visible symptoms of revised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 18: GRAD-CAMS with spatial and spectral activations from inoculated sample with single visible symptoms of revised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 19: GRAD-CAMS with spatial and spectral activations from inoculated sample with single visible symptoms of revised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 20: GRAD-CAMS with spatial and spectral activations from inoculated sample with multiple visible symptoms of revised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

Figure 21: GRAD-CAMS with spatial and spectral activations from inoculated sample with multiple visible symptoms of revised network. Leftmost image shows the sample followed by the corresponding spatial activations maps mapped to four different hyperspectral areas. The areas are 380-537 nm, 538-695 nm, 696-853 nm, and 854-1010 nm.

References

- [1] Dombrowski, A. *et al.* Explanations can be manipulated and geometry is to blame. In Wallach, H. M. *et al.* (eds.) *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, 13567–13578 (2019).
- [2] Adebayo, J. *et al.* Sanity checks for saliency maps. In *Proceedings of Advances in Neural Information Processing Systems*, 9505–9515 (2018).
- [3] Sixt, L., Granz, M. & Landgraf, T. When explanations lie: Why modified BP attribution fails. *CoRR* **abs/1912.09818** (2019). URL <http://arxiv.org/abs/1912.09818>.
- [4] Lin, T. *et al.* Microsoft COCO: common objects in context. In *Proceedings of European Conference on Computer Vision*, 740–755 (2014).
- [5] Selvaraju, R. R. *et al.* Taking a hint: Leveraging explanations to make vision and language models more grounded. In *Proceedings of the IEEE International Conference on Computer Vision*, 2591–2600 (2019).
- [6] Hooker, S., Erhan, D., Kindermans, P. & Kim, B. A benchmark for interpretability methods in deep neural networks. In *Proceedings of Advances in Neural Information Processing Systems*, 9734–9745 (2019).
- [7] Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *Proceedings of International Conference on Learning Representations* (2015).