
RelatIF: Identifying Explanatory Training Examples via Relative Influence

Elnaz Barshan*
Element AI

Marc-Etienne Brunet*
Element AI

Gintare Karolina Dziugaite
Element AI

Abstract

In this work, we focus on the use of influence functions to identify relevant training examples that one might hope “explain” the predictions of a machine learning model. One shortcoming of influence functions is that the training examples deemed most “influential” are often outliers or mislabelled, making them poor choices for explanation. In order to address this shortcoming, we separate the role of global versus local influence. We introduce RelatIF, a new class of criteria for choosing relevant training examples by way of an optimization objective that places a constraint on global influence. RelatIF considers the local influence that an explanatory example has on a prediction relative to its global effects on the model. In empirical evaluations, we find that the examples returned by RelatIF are more intuitive when compared to those found using influence functions.

1 Introduction

As the use of black-box models becomes widespread, many believe that their predictions must be accompanied by interpretable explanations (Lipton, 2016; Doshi-Velez & Kim, 2017; Goodman & Flaxman, 2017). There is a growing body of research concerned with how to best explain black-box predictions (Kim et al., 2017; Ribeiro et al., 2018; Lundberg & Lee, 2017). A natural way to explain a model prediction is to provide supporting examples from the training data. However, one must *choose* which examples are most relevant. One approach is to trace the model prediction back to the training examples and determine the impact of each individual training example.

This idea underlies so-called *influence functions* (IF), which quantify the impact of an individual example on a prediction in a differentiable model (Jaeckel, 1972; Hampel, 1974). The use of IF to *explain* black-box predictions by training examples with high influence was pioneered by Koh & Liang (2017).

In this work, we revisit the use of IF for identifying examples to explain predictions. Our motivation is the observation that the most influential training examples returned by IF are often mislabelled or outliers, on the account of these examples also incurring high loss (see Fig. 1(a)). This phenomenon is well documented in the literature (Hampel, 1974; Campbell, 1978; Cook & Weisberg, 1982; Croux & Haesbroeck, 2000). As a result, we find that the set of examples deemed “influential” tend to be highly overlapping, in the sense that the predictions for many different inputs are all most affected by the same small set of training examples. In other words, these training examples seem to have a global impact on the model. Fig. 1(b) illustrates these effects. Returning examples from an overlapping set of high loss training examples may be desirable when the recipient of the explanation is attempting to debug or otherwise improve the model. However, these examples can be unsatisfying as an explanation for an end user. Explaining a prediction by returning high loss training examples, which are often unusual or misclassified, does not inspire confidence in the model. Nor does it inspire confidence when, for several seemingly different inputs, the same high-loss examples are returned as an explanation.

We propose RelatIF, a new class of criteria for identifying “relevant” training examples to explain a prediction. RelatIF considers the local influence that an explanatory example has on a prediction relative to its global effects on the model. We formulate these criteria in terms of a constrained optimization problem: for a given allowable global “impact” on the model, which training examples most influence the specific prediction? We argue that this approach is better suited to producing an explanatory training example for an end

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. Copyright 2020 by the author(s).

*Equal contribution, alphabetic order

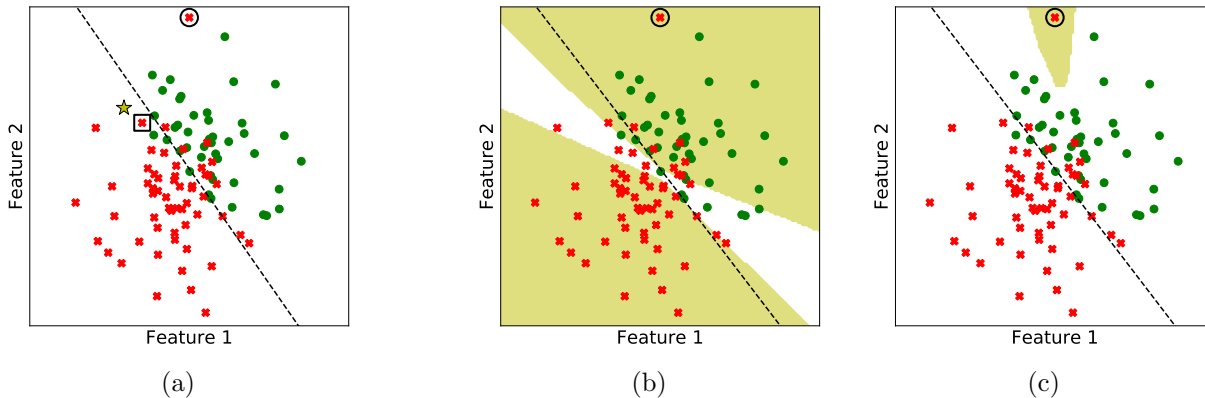


Figure 1: Binary classification by linear decision boundaries (dashed line) to illustrate the difference between IF and RelatIF. (a) The model makes a prediction for the test input (star). As estimated by IF, the most influential training example for the prediction is an outlier (circled). Using RelatIF, the most influential training example is more typical (encased in a square). (b) Using IF, every test input falling within the shaded yellow region is most influenced by the same outlier (circled). Test inputs in the remaining white region are most influenced by one of 5 other high loss examples. (c) Using RelatIF, the region where test inputs are most influenced by the outlier (circled) shrinks. Test inputs in the remaining region are most influenced by one of 65 other examples.

user. Fig. 1(c) illustrates how this objective dampens the influence of an outlier. We show experimentally that RelatIF returns training examples that are more typical as well as more specific to a given input when compared to those identified using IF.

2 Preliminaries

Consider a family of predictors, indexed by parameter values $\theta \in \mathbb{R}^d$, each a map taking input values $x \in \mathcal{X}$ to target values $y \in \mathcal{Y}$. Let $S = (z_1, z_2, \dots, z_n)$ denote the training examples, where $z_i = (x_i, y_i) \in \mathcal{X} \times \mathcal{Y}$ is a pair of an input and target. Let $\ell(z, \theta)$ be the loss for a point z under model parameters θ . Our focus is on algorithms, like stochastic gradient descent (SGD), that return some first-order stationary point $\theta(w_1, \dots, w_n)$ of the objective

$$\sum_{i=1}^n w_i \ell(z_i, \theta), \quad (1)$$

i.e., we are studying the parameters learned by weighted empirical risk minimization, where each w_i is the weight given to the training example z_i . We are primarily interested in the parameter $\theta^* = \theta(1/n, \dots, 1/n)$ obtained by weighting all points equally. Given an unseen input x , our goal is to identify training examples z_i that, through training, had a “meaningful” influence on the prediction that θ^* renders for x .

A simple way to measure the effect of each example z_i on training is to re-weight its contribution, retrain from scratch, and evaluate the change to the learned

parameters. To that end, let $\theta_{i, \epsilon_i}^* = \theta(1/n, \dots, 1/n + \epsilon_i, \dots, 1/n)$, where only the i ’th weight is modified to be $w_i = \frac{1}{n} + \epsilon$. Any deviations of w_i from $\frac{1}{n}$ is denoted by ϵ_i , i.e., $\epsilon_i = w_i - \frac{1}{n}$. Taking $\epsilon_i = -\frac{1}{n}$, equivalently, setting $w_i = 0$, constitutes dropping an example from the training set.

Exhaustive retraining, however, is computationally prohibitive. Instead, we will make an infinitesimal analysis, as was pioneered in the development of *influence functions* (Hampel, 1974; Jaeckel, 1972). To do so, we will assume that the parameter $\theta(w_1, \dots, w_n)$ is a differentiable function of the weights at $w_1 = \dots = w_n = \frac{1}{n}$. Further, we will assume that $\ell(z, \theta)$ is strictly convex, and twice continuously differentiable, in θ .

The method of *influence functions* (IF) allows us to approximate how (some function of) the learned parameters θ^* would change if we were to reweight a training example z_i . The key idea is to make a first-order approximation of change in θ^* around $\epsilon_i = 0$. The same method can be extended to approximate changes in any twice continuously differentiable functions f of the parameters around $\epsilon_i = 0$. Namely, one can approximate how some $f(\theta_{i, \epsilon_i}^*)$ changes with ϵ_i by considering $df(\theta_{i, \epsilon_i}^*)/d\epsilon_i|_{\epsilon_i=0}$. Indeed, our primary interest is the influence on the loss for a test sample.

Definition 2.1. Fix a test sample $z_{\text{test}} = (x_{\text{test}}, y_{\text{test}})$. The *influence of z_i on (the loss of) z_{test}* is defined to be

$$\mathcal{I}_{\text{test}, i} = - \left. \frac{d\ell(z_{\text{test}}, \theta_{i, \epsilon_i}^*)}{d\epsilon_i} \right|_{\epsilon_i=0} = -g_{\text{test}}^T \left. \frac{d\theta_{i, \epsilon_i}^*}{d\epsilon_i} \right|_{\epsilon_i=0},$$

where $g_{\text{test}} = \nabla_{\theta^*} \ell(z_{\text{test}}, \theta^*)$.

Note that the second equality follows from the chain rule and that $\theta_{i,\epsilon_i}^*|_{\epsilon_i=0} = \theta^*$.

Because we consider a negative change in Definition 2.1, a training example z_i having *positive* influence on z_{test} *decreases* the loss $\ell(z_{\text{test}}, \theta_{i,\epsilon_i}^*)$ when upweighted and may thus be considered helpful. A training example having *negative* influence on z_{test} *increases* the loss and may thus be considered harmful.

2.1 Influence for General Loss Functions

Under certain conditions, it can be shown that

$$\left. \frac{d\theta_{i,\epsilon_i}^*}{d\epsilon_i} \right|_{\epsilon_i=0} = -H_{\theta^*}^{-1} g_i, \quad (2)$$

where $H_{\theta^*} = \frac{1}{n} \sum_i \nabla_{\theta^*}^2 \ell(z_i, \theta^*)$ is the Hessian of the objective at θ^* and $g_i = \nabla_{\theta^*} \ell(z_i, \theta^*)$. This result can be obtained by applying the implicit function theorem to the first-order optimality conditions for Eq. (1) (Cook & Weisberg, 1982). In this case, the expression for the influence of z_i on z_{test} becomes

$$\mathcal{I}_{\text{test},i} = g_{\text{test}}^T H_{\theta^*}^{-1} g_i. \quad (3)$$

Again noting that $\theta_{i,\epsilon_i}^*|_{\epsilon_i=0} = \theta^*$, we can form a first-order Taylor series approximation for the change in model parameters due to upweighting z_i ,

$$\theta_{i,\epsilon_i}^* - \theta^* \approx \left. \frac{d\theta_{i,\epsilon_i}^*}{d\epsilon_i} \right|_{\epsilon_i=0} \epsilon_i = -H_{\theta^*}^{-1} g_i \epsilon_i. \quad (4)$$

Similarly, we can approximate the change in loss via

$$\begin{aligned} \ell(z_{\text{test}}, \theta_{i,\epsilon_i}^*) - \ell(z_{\text{test}}, \theta^*) &\approx \frac{d\ell(z_{\text{test}}, \theta_{i,\epsilon_i}^*)}{d\epsilon_i} \epsilon_i \\ &= -g_{\text{test}}^T H_{\theta^*}^{-1} g_i \epsilon_i. \end{aligned} \quad (5)$$

for an infinitesimal change ϵ in w_i . The last equation follows from Eq. (4).

2.2 Influence in Maximum Likelihood

We have presented influence functions as they have been derived for the analysis of a broad class of empirical risk minimization objectives. In this section, we study the special case of cross entropy loss and maximum likelihood estimation (MLE), which leads us to an alternative representation of IF in terms of Fisher information. We use this formulation of IF to derive relative influence in Section 3.

Let $p_\theta(y|x)$ be a parametric statistical model (i.e., likelihood) and let $\ell((x, y), \theta) = -\log p_\theta(y|x)$. Global optimization of Eq. (1) corresponds to maximum likelihood estimation. As shown in Ting & Brochu (2018), in this setting the influence function satisfies

$$\mathcal{I}_{\text{test},i} = g_{\text{test}}^T F_{\theta^*}^{-1} g_i, \quad (6)$$

where

$$F_\theta = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{p_\theta} \nabla_\theta \log p_\theta(y|x_i) \nabla_\theta \log p_\theta(y|x_i)^T \quad (7)$$

is the Fisher information matrix.¹ Then a Taylor series approximation yields

$$\theta_{i,\epsilon_i}^* - \theta^* \approx -F_{\theta^*}^{-1} g_i \epsilon_i, \quad \text{and} \quad (8)$$

$$\ell(z_{\text{test}}, \theta_{i,\epsilon_i}^*) - \ell(z_{\text{test}}, \theta^*) \approx -g_{\text{test}}^T F_{\theta^*}^{-1} g_i \epsilon_i. \quad (9)$$

We briefly discuss the relationship between F_{θ^*} and H_{θ^*} in Appendix A. For a more thorough treatment, see (Martens, 2014).

3 Relative Influence

For a given test input x_{test} , we would like to identify the training examples that serve to explain the target y_{test} . The change in loss on $z_{\text{test}} = (x_{\text{test}}, y_{\text{test}})$ due to re-weighting each z_i in S is described by $\{\mathcal{I}_{\text{test},i}\}_{i=1,\dots,n}$. Returning the top-k examples that maximize $|\mathcal{I}_{\text{test},i}|$ constitutes one form of explanation. However, top influential examples selected via influence functions are often high-loss, i.e., they are either mislabelled, outliers, or otherwise atypical. The sets of top-k examples which maximize $|\mathcal{I}_{\text{test},i}|$ for different test inputs also often overlap. We hypothesize that these types of examples are returned because maximizing $|\mathcal{I}_{\text{test},i}|$ puts no constraint on how re-weighting these examples affects the overall model.

In this section we introduce relative influence, a set of measures which, using influence functions, quantify changes in loss subject to constraints on how the model may change. We show the training examples returned under these constraints are lower loss, more specific to the test input, and that they arguably constitute a more intuitive explanation for an end user.

3.1 θ -Relative Influence

The change in model parameters $\frac{d\theta_{i,\epsilon_i}^*}{d\epsilon_i}$ appearing in Definition 2.1 of an influence function is unconstrained. In contrast, we propose to directly or indirectly constrain the change in the model parameters.

We start with a direct constraint of model parameters. In particular, we aim to identify

$$\begin{aligned} \arg \max_{i \in \{1, \dots, n\}} \max_{\epsilon_i} & |\ell(z_{\text{test}}, \theta_{i,\epsilon_i}^*) - \ell(z_{\text{test}}, \theta^*)| \\ \text{s.t. } & \|\theta_{i,\epsilon_i}^* - \theta^*\|^2 \leq \delta^2. \end{aligned} \quad (10)$$

¹We do not model $p(x)$ and so we replace it with an empirical estimate using training data S

Which is to ask: for some small allowable change in the model parameters δ , which training example z_i should we re-weight to maximally affect the loss on z_{test} ?

Proposition 3.1. Assume that Eq. (4) and Eq. (5) hold with equality. Then Eq. (10) is equivalent to

$$\arg \max_{i \in \{1, \dots, n\}} \frac{|\mathcal{I}_{\text{test}, i}|}{\|H_{\theta^*}^{-1} g_i\|}. \quad (11)$$

The proof is presented in Appendix B.

Definition 3.2. The θ -relative influence of training example z_i on the loss of test sample z_{test} is

$$\frac{\mathcal{I}_{\text{test}, i}}{\|H_{\theta^*}^{-1} g_i\|}.$$

3.2 ℓ -Relative Influence

Here we propose an alternative way to constrain the change in the model in a maximum likelihood setting, where our model outputs $p_{\theta}(y|x)$. In particular, we focus on the expected change in squared loss (log-likelihood). In the appendix we show that such a constraint is equivalent to constraining the change in the Kullback-Leibler (KL) divergence between the original and z_i re-weighted model.

Limiting the allowable expected change in squared loss poses an indirect constraint on θ_{i, ϵ_i}^* . We aim to identify

$$\begin{aligned} \arg \max_{i \in \{1, \dots, n\}} \max_{\epsilon_i} & |\ell(z_{\text{test}}, \theta_{i, \epsilon_i}^*) - \ell(z_{\text{test}}, \theta^*)| \\ \text{s.t. } & \mathbb{E}_{p_{\theta^*}} (\ell(z, \theta_{i, \epsilon_i}^*) - \ell(z, \theta^*))^2 \leq \delta^2. \end{aligned} \quad (12)$$

Which is to ask: for some small, total allowable expected squared change in loss, which training example z_i should we re-weight so as to maximally affect the loss on z_{test} ?

Proposition 3.3. Assume Eq. (8) and Eq. (9) hold with equality. Then Eq. (12) is equivalent to

$$\arg \max_{i \in \{1, \dots, n\}} \frac{|\mathcal{I}_{\text{test}, i}|}{\sqrt{\mathcal{I}_{i, i}}}. \quad (13)$$

The proof is presented in Appendix B.

Definition 3.4. The ℓ -relative influence of training example z_i on the loss of test sample z_{test} is

$$\frac{\mathcal{I}_{\text{test}, i}}{\sqrt{\mathcal{I}_{i, i}}}.$$

3.3 Geometric Interpretation

The influence of a training example z_i on the loss of a test sample z_{test} can be written as

$$\begin{aligned} \mathcal{I}_{\text{test}, i} &= \langle -g_{\text{test}}, \frac{d\theta_{i, \epsilon_i}^*}{d\epsilon_i} \rangle \\ &= \langle -g_{\text{test}}, -H_{\theta^*}^{-1} g_i \rangle, \end{aligned}$$

where $\langle \cdot, \cdot \rangle$ is the inner product operator. The inner product notation emphasizes that $|\mathcal{I}_{\text{test}, i}|$ is equal to the projection length of the change in parameters vector onto the test sample's (negative) loss gradient. A training example z_i which causes a *large magnitude* change in parameters (e.g., an outlier) is likely to have a large magnitude influence on a wide range of test samples, i.e., such a z_i has a global effect. Conversely, a training example z_i which causes a small magnitude change in parameters, will only have a high magnitude influence on test samples for which this change in parameters is *directionally aligned* with g_{test} .

RelatIF considers the influence of a training example relative to its global effects. Geometrically, this switches from a paradigm of identifying relevant training examples using projections to one using cosine similarity². Fig. 2 illustrates this difference.

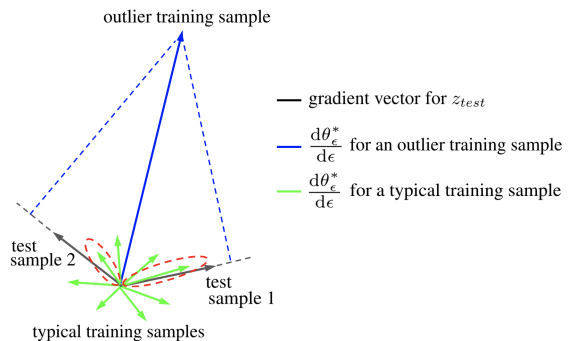


Figure 2: Geometric interpretation of RelatIF. Using IF, the influence of z_i on z_{test} is equal to the projection length of the change in parameters vector onto the test sample's (negative) loss gradient. Observe how the same outlier training example becomes the top influential example for two different test samples due to the very large magnitude effect on the model parameters. RelatIF, instead uses cosine similarity to identify the top influential examples, and thus return those in the red dotted circles.

The equivalence between θ -RelatIF and the cosine similarity between a test point gradient g_{test} and a change in the parameters vector $\frac{d\theta_{i, \epsilon_i}^*}{d\epsilon_i}$ is straightforward. The norm of the test sample gradient $\|g_{\text{test}}\|$ is independent of i , and so it follows that

$$\arg \max_i \frac{\mathcal{I}_{\text{test}, i}}{\|H_{\theta^*}^{-1} g_i\|} = \arg \max_i \frac{\langle -g_{\text{test}}, -H_{\theta^*}^{-1} g_i \rangle}{\|g_{\text{test}}\| \|H_{\theta^*}^{-1} g_i\|}.$$

Further, ℓ -RelatIF can also be interpreted as cosine similarity, but not using the standard euclidean inner product. Instead, it is cosine similarity under the inner product $\langle a, b \rangle = a^T F_{\theta^*}^{-1} b$.

²Recall, $\cos(a, b) = \frac{\langle a, b \rangle}{\|a\| \|b\|}$, and $\|a\| = \sqrt{\langle a, a \rangle}$.

Computational considerations Computing influence can be challenging, especially in larger models. The terms in the denominator introduced by RelatIF pose a further computational challenge. We outline our approach to computing relative influence in larger models in Appendix C. It builds off the strategies used by Koh & Liang (2017), but makes some further approximations, principally motivated by KFAC (Martens & Grosse, 2015)

4 Experimental Results

We empirically assess the use of relative influence (RelatIF) as a technique for explaining model predictions. We compare the examples identified by RelatIF to the baseline examples identified using influence functions (IF) or k-nearest neighbors (k-NN). We consider three models and data sets: logistic regression trained on MNIST handwritten digit data set (LeCun et al., 1998)³ a convolutional neural network (ConvNet) trained on CIFAR10 object recognition data set (Krizhevsky et al., 2009)⁴ and a long short-term memory (LSTM) network trained for classifying names based on their language of origin.^{5,6}

4.1 Quantitative Analysis

It is known that the training examples selected via IF are often high-loss, i.e., they are either mislabelled, outliers, or atypical (Hampel, 1974; Campbell, 1978; Cook & Weisberg, 1982; Croux & Haesbroeck, 2000). We further demonstrate that the sets of top-k examples maximizing $\mathcal{I}_{\text{test},i}$ for different test inputs often have a big overlap. We quantify these properties and contrast them to the examples selected via RelatIF.

Overlap and probability We train a logistic regression model on MNIST, as well as a ConvNet on CIFAR10. In each setup we identify the top-k positively influential examples for 10,000 test inputs using both IF and RelatIF. For each of the selected examples, we examine the class probabilities returned by the model. Since these models are trained based on a negative-log-likelihood loss function, output probabilities are indicative of the loss. We also examine the extent to which these examples overlap by considering

³The test accuracy of this model is 91.93 and the damping coefficient for hessian inversion is 0.001.

⁴The model architecture is 16C-32C-64C-F with max pooling after conv layers. The model test accuracy is 73.22 and the damping coefficient for hessian inversion is 0.1.

⁵The data set is publicly available at <https://download.pytorch.org/tutorial/data.zip>.

⁶We randomly selected 15% of the examples from each class for the test set. The accuracy on this test set is 71.45 and the damping coefficient for hessian inversion is 0.001.

the cardinality of the set of top-k examples across all 10,000 test inputs. The maximum feasible cardinality is $k \times 10,000$ (limited by the number of training examples), which is achieved when unique training examples are identified as explanatory for each of the test samples. A cardinality smaller than this maximum implies that there is an overlap between the influential examples. The results are tabulated in Table 1.

We find that in all of our experiments, the positively influential points identified by RelatIF have little overlap. They also have much higher probability under the model’s conditional distribution, and hence have lower loss. For example, in our ConvNet/CIFAR10 setup, the top-5 highest influential training examples identified by IF across all 10,000 test inputs is a set of only 4,869 examples with mean probability 0.511. In contrast, when identified with RelatIF, the set contains on average 26,430 examples with mean probability 0.928.

Global effects We are also interested in quantifying the global effect of the top influential examples selected by each method. We train a logistic regression model on MNIST to convergence, explicitly seeding all sources of randomness. We then pick 100 test samples uniformly at random. For each of these test samples, we identify the most positively influential training example using IF and RelatIF. We then retrain the model without that training example, maintaining the same random seed, and measure the effects. Specifically, we consider how the removal of that training example affects the root squared error across the training set, as well as how it affects the learned model parameters. The results are tabulated in Table 2.

We find that the examples identified by IF have a much stronger global impact on the model than do those identified by RelatIF. Using IF, the removal of the top influential examples results in a mean change in parameters of 0.039 and a root mean squared change in loss of 0.879. Using RelatIF, these changes are an order of magnitude smaller.

4.2 Qualitative Analysis

The purpose of our qualitative experiments is to evaluate RelatIF in the context of example-based explanations of model predictions. RelatIF is derived by explicitly minimizing the global effect on the model. As a result, the examples selected by RelatIF are specific to the given test sample, and the effect of examples with large global influence is diminished.

We find that examples identified by IF have significantly more global influence. An illustration of this effect for a logistic regression trained on MNIST is presented in Fig. 3. As this figure shows, based on IF, a

Method	Logistic Regression on MNIST				ConvNet on CIFAR10			
	Infl. Set Cardinality		Infl. Set Probability		Infl. Set Cardinality		Infl. Set Probability	
	top 1	top 5	top 1	top 5	top 1	top 5	top 1	top 5
IF	2704	4984	0.168 ± 0.16	0.268 ± 0.21	2279	4869	0.431 ± 0.21	0.511 ± 0.21
θ -RelatIF	8582	29317	0.929 ± 0.16	0.922 ± 0.17	8264	26381	0.918 ± 0.15	0.910 ± 0.16
ℓ -RelatIF	8776	30431	0.938 ± 0.16	0.930 ± 0.17	8336	26550	0.931 ± 0.14	0.920 ± 0.15

Table 1: Comparing the set of top influential training examples recovered using IF and RelatIF for a logistic regression trained on MNIST and a ConvNet trained on CIFAR10 data sets. The number of test samples for each data set is 10000. For each test sample, the top (1 or 5) positively influential training example is recovered and added to the *influential set* (duplicates are removed). The cardinality of this set as well of the mean probability (\pm std-dev) of its elements are reported for IF and different variations of RelatIF. The larger cardinality of the set is an indication of recovering more specific (as opposed to globally) influential examples. Higher likelihood for the elements of the set shows more typical examples (and fewer outliers) are retrieved.

Method	$\Delta\ell_{\text{test}}$	$\ \Delta\theta\ $	$\sqrt{\sum(\Delta\ell_i)^2}$
IF	0.0106 ± 0.0023	0.039 ± 0.003	0.879 ± 0.092
θ -RelatIF	0.0084 ± 0.0027	0.004 ± 0.001	0.080 ± 0.033
ℓ -RelatIF	0.0086 ± 0.0028	0.004 ± 0.002	0.068 ± 0.036
Nearest-N	0.0048 ± 0.0025	0.003 ± 0.001	0.044 ± 0.024

Table 2: The change in loss at a test sample ($\Delta\ell_{\text{test}}$), compared with the norm of the change in parameters ($\|\Delta\theta\|$) and root sum of square change in loss over the training set ($\sqrt{\sum(\Delta\ell_i)^2}$). The results come from removing the mostly positively influential training sample as determined by different methods, then retraining the model (i.e., leave-one-out retraining). The model is a logistic regression trained on MNIST. The experiment is repeated for 100 randomly selected test samples. The mean \pm standard error are reported. While IF identifies the training samples having the largest influence on the test sample’s loss, RelatIF finds comparably influential samples with an order of magnitude smaller global effects. A nearest neighbor baseline is included for comparison.

single high-loss training example is the most positively influential example for a number of different test images. We argue that explaining the model prediction for all of these different test samples using a single high-loss (e.g., outlier) training example is not intuitive and may undermine the user’s trust. RelatIF addresses this issue and enables us to identify examples that are more relevant to each specific test input.⁷

We study the effectiveness of RelatIF for generating example-based explanations under different types of models and data sets. Fig. 4 shows the explanations produced for a logistic regression trained on MNIST. For each test image, the top five positively influential

⁷As shown in Appendix D, θ -RelatIF and RelatIF produce similar example-based explanations and for the experiments in this section we used θ -RelatIF.




Figure 3: The most positively influential training examples suggest by IF and RelatIF. Test samples are in the first row. IF selects the same (high-loss) training example as the top influential example for all the test images. In contrast, RelatIF discovers training examples that are specific to each of the test images.

training examples for the predicted label identified by IF and RelatIF are presented. Most of the examples selected by IF are misclassified training examples (due to their large global influence on the model parameters). In contrast, RelatIF identifies visually similar examples. We understand this to be the result of constraining global influence, thereby putting more importance on the similarity of the test/train gradients.

Now we make the model a bit more complex and try to explain a ConvNet trained on CIFAR10 data set.⁸ Fig. 5 shows the generated explanation for a correctly classified (dog) and a misclassified (truck) test sample. Interestingly, we can see that a set of visually similar dog training images are specifically helpful for predicting the label dog for the given test image. In the case of the red truck, the explanation by RelatIF suggests

⁸This model has 24234 parameters and all of them are used for computing (relative) influence scores

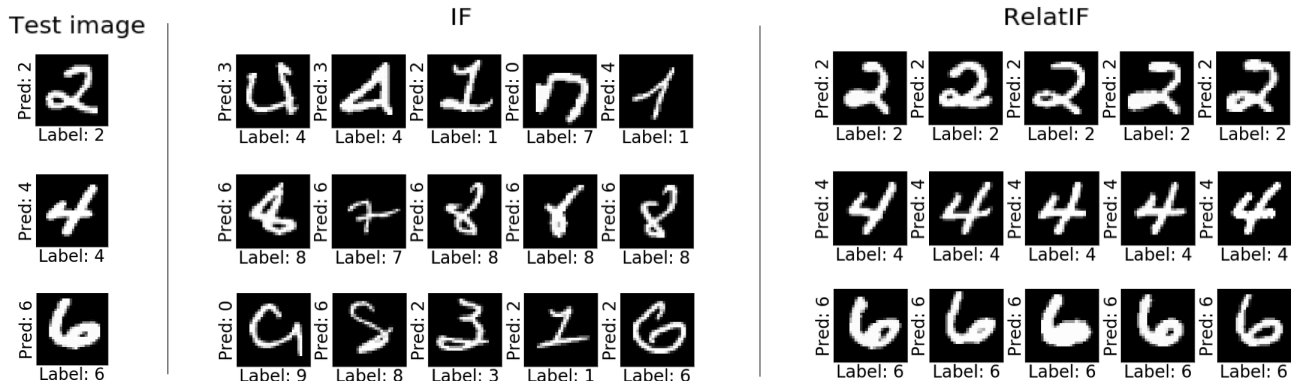


Figure 4: Comparison of the top positively influential training examples recovered by IF and RelatIF for logistic regression on MNIST. Each row shows a test sample and the top five positively influential training examples for the predicted label selected by each method. The true class and the predicted label for each example are provided. Observe how most of the examples selected by IF are unusual and misclassified, despite the training set containing examples with higher visual similarity (as suggested by RelatIF).

that the presence of a number of red cars (with similar shade of color and shape) specifically influences the model to classify it as a car. For more examples of CIFAR10 explanations see Appendix E.

A similar experiment is conducted for a character-level LSTM network trained for classifying names based on their language of origin. The data set consists of 20,000 surnames from 18 different languages. Due to the large number of model parameters, similar to Koh & Liang (2017), in this experiment we used only the last layer parameters for computing influence scores. Table 3 shows the generated explanation for classifying surname “Vasyukov” as Russian. The examples identified by IF are all Russian names that are misclassified as Japanese. In contrast, all of the examples identified by RelatIF are Russian names that end with “kov”, similar to the test input.

The explanations provided for these models enable the end user to assess the reliability of the generated prediction. The user can confirm, reject or modify the model decision by inspecting the extracted evidence from the training set. This is especially useful when the model is not confident in its prediction. For an example of this application see Appendix G.

4.3 Comparison to Nearest-Neighbors

Here we compare a test sample’s k -nearest neighbors (k -NNs) to the samples returned using RelatIF. For a quantitative comparison we use the same MNIST and CIFAR10 experimental setups discussed in Section 4.1. We consider the overlap between the top- k samples identified by RelatIF, and the k -NNs (l_2 distance) for k in $\{1, 5, 10, 20\}$. In the logistic regression

	IF	RelatIF
Test Input: Vasyukov (Russian)	Tokovoi (Japanese)	Gasyukov (Russian)
	Nasikan (Japanese)	Tsayukov (Russian)
	Bakai (Japanese)	Tzarakov (Russian)
	Mihailutsa (Japanese)	Tzayukov (Russian)
	Jugai (Japanese)	Haryukov (Russian)
	Mikhailutsa (Japanese)	Tsarakov (Russian)

Table 3: Generating Explanation for text classification. The model is a character-level RNN trained for classifying names based on their language of origin. The most positively influential training examples selected based on IF and RelatIF are listed. The language in parentheses shows the predicted label. The true label for all examples shown is Russian. Observe how the examples identified by IF are all misclassified as Japanese, while the examples identified by RelatIF end with “kov”, which is similar to the test input.

model on MNIST we find the overlap ranges from 26-35%, increasing with k . A substantial overlap, but still indicative of a fundamentally different set of points. In the CNN on CIFAR10 we find much less overlap, only 2.6-5.6%, decreasing with k . By comparison, the overlap between the top- k samples identified by IF and the k -NNs is 1.3-2.7% in the logistic regression model, and 0.67-1.24% in the CNN. This indicates that the points returned by RelatIF are closer in l_2 distance to the test sample than those returned by IF.

A qualitative comparison finds a significant difference between the k -NNs and the samples returned by RelatIF (see Appendix H). These examples are abundant and easily found for our CNN on CIFAR10. We have included some of these examples in the appendix.



Figure 5: Generating example-based explanations using IF and RelatIF for the predicted label by a ConvNet trained on CIFAR10. Each row shows a test sample and the top five positively influential training examples for the predicted label selected by each method. The true class and the predicted label for each example is marked. Compared to IF, RelatIF is able to retrieve training examples with higher visual similarity to the test image.

5 Related work

Influence functions (IF) were originally proposed by Hampel (1974). Methods based on IF are identical to the infinitesimal jackknife, which was introduced earlier in an independent work on robust statistics by Jaeckel (1972). The connection between IF and infinitesimal jackknife was recognized by Efron (1982).

Koh & Liang (2017) use IF to identify influential training examples in a modern machine learning context. They further demonstrate that influence functions can be used to identify mislabelled data and generate adversarial training examples (Koh et al., 2018). More recently, Koh et al. (2019) use IF to approximate the effects of re-weighting groups of training examples.

A range of approaches have been introduced to identify training examples for interpreting predictions. One such approach was recently proposed by Khanna et al. (2019). The authors use Fisher kernels to identify points that are “most responsible” for a given set of predictions. They demonstrate that in the case of a negative log-likelihood loss function and with a specific choice of hyper-parameters, their approach coincides with IF. Another recent contribution is due to Ghorbani & Zou (2019). The authors propose a new form of data valuation which identifies important examples or subsets of the data using Shapley values.

Other approaches to example-based explanation include the extraction of prototypical data, as well as generating counter-factual examples. Bien & Tibshirani (2011) propose a method to select prototypes for interpretable classification. Kim et al. (2016) introduce a method for explaining the data set by identifying its prototypes and “criticisms” and argue that atypical data must also be extracted. Chang

et al. (2019) introduce a novel method for generating counter-factual images to be used as an explanation.

Finally, the effect of example-based explanations on user trust have been studied in the human-computer interaction community. One of the most recent study is presented in (Zhou et al., 2019). For each test sample, Zhou et al. select examples maximizing influence. They provide these selected training examples along with model prediction and study the effect on the users’ trust in the predictive model. A similar set of experiments can be found in (Cai et al., 2019). In this study, the effect on users’ trust is compared for different types of “explanatory” examples.

6 Discussion

In this work we introduce relative influence (RelatIF) and show that it identifies more intuitive examples than those identified by influence functions (IF). In particular, we demonstrate that, unlike examples identified by IF, which tend to be atypical or misclassified, and otherwise unrelated to the test sample, examples identified by RelatIF seem to be more specific to test sample. As a result, we believe that RelatIF may be superior at identifying explanatory examples that a user may find intuitive and acceptable. The desiderata of example-based explanations are case specific, and therefore one cannot make an absolute comparison of IF versus RelatIF. A data scientist fitting a model to a data set containing spurious samples may very well wish to see the examples identified by *both* IF and RelatIF. We believe a large user study, which is beyond the scope of this paper, may shed light on the relative strengths and weaknesses of RelatIF, and lead to further insight.

References

- Naman Agarwal, Brian Bullins, and Elad Hazan. Second-order stochastic optimization for machine learning in linear time. *The Journal of Machine Learning Research*, 18(1):4148–4187, 2017.
- Jacob Bien and Robert Tibshirani. Prototype selection for interpretable classification. *Ann. Appl. Stat.*, 5(4):2403–2424, 12 2011. doi: 10.1214/11-AOAS495.
- Carrie J. Cai, Jonas Jongejan, and Jess Holbrook. The effects of example-based explanations in a machine learning interface. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, IUI '19, 2019.
- Norm A Campbell. The influence function as an aid in outlier detection in discriminant analysis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 27(3):251–258, 1978.
- Chun-Hao Chang, Elliot Creager, Anna Goldenberg, and David Duvenaud. Explaining image classifiers by adaptive dropout and generative in-filling. In *International Conference on Learning Representations*, 2019.
- R Dennis Cook and Sanford Weisberg. *Residuals and influence in regression*. New York: Chapman and Hall, 1982.
- Christophe Croux and Gentiane Haesbroeck. Principal component analysis based on robust estimators of the covariance or correlation matrix: influence functions and efficiencies. *Biometrika*, 87(3):603–618, 2000.
- Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.
- Bradley Efron. *The jackknife, the bootstrap, and other resampling plans*, volume 38. Society for Industrial and Applied Mathematics, 1982.
- Amirata Ghorbani and James Y. Zou. Data shapley: Equitable valuation of data for machine learning. In *ICML*, 2019.
- Bryce Goodman and Seth Flaxman. European union regulations on algorithmic decision-making and a right to explanation. *AI Magazine*, 38(3):50–57, 2017.
- Frank R Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393, 1974.
- Louis A Jaeckel. *The infinitesimal jackknife*. Bell Telephone Laboratories, Murray Hill, NJ, 1972.
- Rajiv Khanna, Been Kim, Joydeep Ghosh, and Sanmi Koyejo. Interpreting black box predictions using fisher kernels. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pp. 3382–3390. PMLR, 16–18 Apr 2019.
- Been Kim, Rajiv Khanna, and Oluwasanmi O Koyejo. Examples are not enough, learn to criticize! criticism for interpretability. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 29*, pp. 2280–2288. Curran Associates, Inc., 2016.
- Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). *arXiv preprint arXiv:1711.11279*, 2017.
- Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1885–1894. JMLR.org, 2017.
- Pang Wei Koh, Jacob Steinhardt, and Percy Liang. Stronger data poisoning attacks break data sanitization defenses. *arXiv preprint arXiv:1811.00741*, 2018.
- Pang Wei Koh, Kai-Siang Ang, Hubert HK Teo, and Percy Liang. On the accuracy of influence functions for measuring group effects. *arXiv preprint arXiv:1905.13289*, 2019.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Masters thesis, Department of Computer Science, University of Toronto, 2009.
- Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Zachary C Lipton. The mythos of model interpretability. *arXiv preprint arXiv:1606.03490*, 2016.
- Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, pp. 4765–4774, 2017.
- James Martens. New insights and perspectives on the natural gradient method. *arXiv preprint arXiv:1412.1193*, 2014.
- James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International Conference on Machine Learning*, pp. 2408–2417, 2015.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Anchors: High-precision model-agnostic

explanations. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

Daniel Ting and Eric Brochu. Optimal subsampling with influence functions. In *Advances in Neural Information Processing Systems*, pp. 3650–3659, 2018.

Jianlong Zhou, Zhidong Li, Huaiwen Hu, Kun Yu, Fang Chen, Zelin Li, and Yang Wang. Effects of influence on user trust in predictive decision making. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI EA '19, pp. LBW2812:1–LBW2812:6, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-5971-9. doi: 10.1145/3290607.3312962.