

On Mixup Regularization

Luigi Carratino*

MaLGA - University of Genova, Italy

LUIGI.CARRATINO@DIBRIS.UNIGE.IT

Moustapha Cissé

Google Research - Brain team, Accra

MOUSTAPHACISSE@GOOGLE.COM

Rodolphe Jenatton

Google Research - Brain team, Berlin

RJENATTON@GOOGLE.COM

Jean-Philippe Vert[†]

Google Research - Brain team, Paris

JEAN-PHILIPPE.VERT@M4X.ORG

Editor: Simon Lacoste-Julien

Abstract

Mixup is a data augmentation technique that creates new examples as convex combinations of training points and labels. This simple technique has empirically shown to improve the accuracy of many state-of-the-art models in different settings and applications, but the reasons behind this empirical success remain poorly understood. In this paper we take a substantial step in explaining the theoretical foundations of Mixup, by clarifying its regularization effects. We show that Mixup can be interpreted as standard empirical risk minimization estimator subject to a combination of data transformation and random perturbation of the transformed data. We gain two core insights from this new interpretation. First, the data transformation suggests that, at test time, a model trained with Mixup should also be applied to transformed data, a one-line change in code that we show empirically to improve both accuracy and calibration of the prediction. Second, we show how the random perturbation of the new interpretation of Mixup induces multiple known regularization schemes, including label smoothing and reduction of the Lipschitz constant of the estimator. These schemes interact synergistically with each other, resulting in a self calibrated and effective regularization effect that prevents overfitting and overconfident predictions. We corroborate our theoretical analysis with experiments that support our conclusions.

Keywords: regularization, generalization, approximation, deep learning, mixup

1. Introduction

Regularization is an essential component of machine learning models and plays an even more important role in deep learning (Goodfellow et al., 2016). Regularization mechanisms can take various forms. They can be *explicitly* enforced by: (i) applying various penalties to the parameters of the models (Hinton, 1987; Krogh and Hertz, 1991; Bartlett et al., 2017; Neyshabur et al., 2015; Sedghi et al., 2019; Arjovsky et al., 2017), (ii) injecting noise to the internal representations of the network (Srivastava et al., 2014; Gal and Ghahramani, 2016) and/or to its outputs (Szegedy et al., 2016; Müller et al., 2019), or (iii) normalizing the activations (He et al., 2016; Salimans and Kingma, 2016). Or they can be *implicit*

*. This work was done in part at Google Brain, Paris.

†. Now at Owkin, Paris, France

thanks to: (j) parameter sharing in architectures such as convolutional networks (LeCun et al., 1998), (jj) the choice of the optimization algorithm (Neyshabur, 2017), e.g., stochastic gradient descent converging to small norm solutions (Arora et al., 2019), or (jjj) through data augmentation and transformation (Goodfellow et al., 2016). There is a large body of work explaining the effects of the numerous explicit and implicit regularization procedures existing in the literature. For instance, explicit regularization schemes usually proceed from analysis aiming to control specific characteristics of a model such as robustness (Hein and Andriushchenko, 2017; Cissé et al., 2017) or calibration (Guo et al., 2017; Müller et al., 2019), while the forms of implicit regularization are often understood through the angle of generalization (Neyshabur, 2017; Arora et al., 2019). However, the regularization effects of modern data augmentation procedures are less theoretically understood.

Data augmentation is a core ingredient for successful deep learning pipelines. It helps to alleviate sample size issues and prevent overfitting. In simple cases, there are known equivalences between data augmentation and other existing explicit regularization procedures, e.g., training with additional noisy points in least-squares regression is equivalent to Tikhonov regularization (Bishop, 1995). Similar analysis have recently been performed to explain the regularization effect of dropout (Srivastava et al., 2014; Wager et al., 2013; Wei et al., 2020). In this work, we focus on *Mixup* (Zhang et al., 2018; Tokozume et al., 2018), a recently introduced data-augmentation technique that consists in generating examples as random convex combinations of data points and labels from the training set (as illustrated in Figure 1). Despite its simplicity, Mixup has been shown to substantially improve generalization on a broad range of tasks ranging from computer vision (Zhang et al., 2018; Tokozume et al., 2018) to natural language processing (Guo, 2020) and semi-supervised learning (Berthelot et al., 2019). The success of Mixup has triggered several variations such as adaptive Mixup (Guo et al., 2019), manifold Mixup (Verma et al., 2019) and Cutmix (Yun et al., 2019), but the reasons why Mixup and its variants work so well in practice remain poorly understood.

Mixup’s primary motivation was to alleviate overfitting in training deep neural networks (Zhang et al., 2018). However, previous studies have also empirically noticed other desirable regularization effects it induces. These include improved calibration (Thulasidasan et al., 2019), robustness to input adversarial noise (Zhang et al., 2018), and robustness to label corruption (Zhang et al., 2018). Zhang et al. (2018) also showed it helps stabilize notoriously difficult learning problems such as generative adversarial networks. Traditionally, separate regularization methods are applied to induce the above effects. For example, label smoothing (Szegedy et al., 2016; Müller et al., 2019) leads to better calibration, while dropout improves generalization (Srivastava et al., 2014; Wager et al., 2013) and robustness to label corruption (Arpit et al., 2017). Lipschitz regularization helps stabilize the training of generative adversarial networks (Arjovsky et al., 2017; Gulrajani et al., 2017). It also leads to increased robustness to adversarial perturbations (Hein and Andriushchenko, 2017; Cissé et al., 2017). Table 1 shows a comparison of various regularization procedure proposed in the literature, and the effect they are known to induce on the model. Although all these desirable regularization effects have been observed empirically, no theoretical explanation has been given yet.




Fig. 1: Illustration of how training a model with Mixup (second plot) differs from training a model on original data (first plot), the fourth plot highlighting the discrepancy between the Bayes classifiers in both situations (black vs red). To explain this difference, we show in this paper that the model trained with Mixup can be interpreted as a regularized version of a model trained on modified data (third plot, blue curve on the zoom plot), and characterize both the data modification (from black to blue) and the regularization effect (from blue to red). Both effects interact synergistically to confer Mixup strong regularization properties, which may explain its good empirical behavior in a variety of tasks.

In this work, we propose the first theoretical analysis of Mixup¹ to better understand the reasons for its empirical success. We show that Mixup can be analyzed through the lenses of empirical risk minimization with random perturbations, and exploit ingredients from previous analysis of dropout (Wager et al., 2013; Khalfaoui et al., 2019; Wei et al., 2020) to derive a regularized objective function that sharply captures the regularization effects of Mixup. In particular, our analysis sheds some light on the multiple effects that Mixup borrows from the popular regularization mechanisms listed above such as label smoothing (Pereyra et al., 2017) (output noise) or dropout (Srivastava et al., 2014) (input noise), and how it uniquely combines them to improve calibration and smooth the Jacobian of the model. We further show that this analysis points out a missing step in learning with Mixup, and we present how applying a simple transformation when evaluating at test time the function learned with Mixup can improve accuracy and calibration. More precisely, we make the following contributions (illustrated in Figure 1):

- We show that Mixup can be reinterpreted as a standard empirical risk minimization procedure, applied to a transformation of the original data perturbed by random perturbations, and give explicit formulas for the data transformation and the perturbations.
- In particular, we show that the Mixup transformations shrinks both the inputs and the outputs towards their mean, the later creating a form of regularization by label smoothing. We notably give a formal description of the effect of label smoothing in the case of the cross-entropy loss where it translates into an increase in the entropy of the predictions.
- We show that Mixup learns functions from a modified version of the input space of the training points to a modified version of the output space of the training points. Thus, we present how to properly evaluate the learned functions to further improve accuracy and calibration.

1. After we published a first version of this work (Carratino et al., 2020), Zhang et al. (2021) independently derived a similar and complementary analysis of Mixup; we summarize in Section 5 the main differences between both works.

- We characterize the random perturbations induced by Mixup on both the inputs and the outputs, as well as their dependency and their correlation structure.
- We deduce an approximation of the regularization induced by Mixup, and highlight in particular how it regularizes both the model and its derivatives. We discuss in details the specific cases of classification with cross-entropy loss, and least squares regression.
- We provide empirical support for our interpretation of Mixup regularization.

Method	Calibration	Jacobian Reg.	Robustness Label Noise	Input Normaliz.
Label smooth. (Szegedy et al., 2016)	✓			
Spectral Reg. (Cissé et al., 2017)		✓		
Dropout (Wager et al., 2013)	✓	✓	✓	
Temperat. scaling (Guo et al., 2017)	✓			
Mixup (Zhang et al., 2018)	✓	✓	✓	✓

Table 1: Summary of the effects induced by various regularizers. Absence of checkmark means the corresponding effect is not known for this regularizer.

The rest of the paper is organized as follows. In Section 2, we introduce notations used throughout the paper and describe the setting of empirical risk minimization and learning with Mixup. In Section 3, we show how Mixup can be interpreted as an empirical risk minimization on modified data with random perturbations. In Section 4 we analyze the regularization effect of Mixup through a quadratic Taylor approximation of the formulation derived in Section 3. In Section 5, we discuss in detail several aspects of Mixup that the theoretical analysis in Section 4 and Section 3 suggest, and confront them to experimental validations. The proofs of all results are detailed in the Annex, together with additional experimental results.

2. Notations and setting

Notations. For any $n \in \mathbb{N}$, $[n] = \{1, \dots, n\}$ is the set of nonzero integers up to n , $\mathbf{1}_n \in \mathbb{R}^n$ is the n -dimensional vector of ones, and $\mathbf{0}_n$ and $\mathbf{I}_n \in \mathbb{R}^{n \times n}$ are the n -dimensional null and identity matrices, respectively. For any two matrices Z, Z' of equal size we note $\langle Z, Z' \rangle = \text{Trace}(Z^\top Z')$ their Frobenius inner product, and with $\|Z\|_F = \sqrt{\langle Z, Z \rangle}$ the Frobenius norm. For any vector $x \in \mathbb{R}^n$, matrix $M \in \mathbb{R}^{n \times m}$ and positive semi-definite matrix $Z \in \mathbb{R}^{n \times n}$, we denote by $\|x\|_Z^2 = x^\top Z x$ the squared semi-norm of x with metric Z , and with $\|M\|_Z^2 = \langle M, ZM \rangle = \text{Trace}(M^\top ZM)$ the squared Frobenius norm with metric Z . For any function $f : \mathbb{R}^a \rightarrow \mathbb{R}^b$ and vector $x \in \mathbb{R}^a$, we denote respectively by $\nabla f(x) \in \mathbb{R}^{b \times a}$ and $\nabla^2 f(x) \in \mathbb{R}^{b \times a \times a}$ the Jacobian and Hessian of f at x , i.e., if $f(x) = (f_1(x_1, \dots, x_a), \dots, f_b(x_1, \dots, x_a))$, then $[\nabla f(x)]_{i,j} = \partial f_i / \partial x_j(x)$ and $[\nabla^2 f(x)]_{i,j,k} = \partial^2 f_i / \partial x_j \partial x_k(x)$, for $(i, j, k) \in [b] \times [a] \times [a]$. Note in particular that if $f : \mathbb{R}^a \rightarrow \mathbb{R}$, then the gradient of f is a row vector $\nabla f(x) \in \mathbb{R}^{1 \times a}$. When f has several arguments and we wish to take partial derivatives with respect to some of the arguments, we explicitly name the different arguments as $f(u, v)$ and then indicate as a subscript to the ∇ sign the argument(s) according to which we take derivatives, e.g., if $u \in \mathbb{R}^{a_u}$ and $v \in \mathbb{R}^{a_v}$, then $\nabla_u f(u, v) \in \mathbb{R}^{b \times a_u}$ is the Jacobian of f with respect to u ,

and $\nabla_{uv}^2 f \in \mathbb{R}^{b \times a_u \times a_v}$ is the tensor of second derivatives of f of the form $[\nabla_{uv}^2 f(u, v)]_{i,j,k} = \partial^2 f_i / \partial u_j \partial v_k(u, v)$ for $(i, j, k) \in [b] \times [a_u] \times [a_v]$. We recall that if $f : \mathbb{R}^{a_u + a_v} \rightarrow \mathbb{R}$ is twice continuously differentiable, then $\nabla_{uv} f = \nabla_{vu} f^\top$, by Schwarz's theorem. For any random variable X and measurable function f , we denote by $\mathbb{E}_X f(X)$ the expectation of $f(X)$, or simply $\mathbb{E} f(X)$ when no confusion is possible. For any shape parameters $\alpha, \beta > 0$, and any interval $[a, b] \subset [0, 1]$, $\text{Beta}_{[a,b]}(\alpha, \beta)$ denotes the truncated Beta distribution on $[a, b]$, i.e., the distribution of a random variable with values in $[a, b]$ and density proportional to $x^{\alpha-1}(1-x)^{\beta-1}$ on $[a, b]$. We simply write $\text{Beta}(\alpha, \beta) = \text{Beta}_{[0,1]}(\alpha, \beta)$ for the usual Beta distribution. For any $p \in [0, 1]$, $\text{Ber}(p)$ denotes the Bernoulli distribution with parameter p . For any $c \in \mathbb{N}$, we denote by

$$\Delta_c = \{u \in \mathbb{R}^c; u^\top \mathbf{1}_c = 1 \text{ and for } j \in [c], u_j \geq 0\}$$

the simplex in \mathbb{R}^c , and for any $p \in \Delta_c$, we denote by $\mathcal{Z}(p) = -\sum_{j=1}^c p_j \log(p_j)$ the entropy of a categorical distribution with parameter p .

Learning problem. We consider a training set $S_n = \{(x_1, y_1), \dots, (x_n, y_n)\}$ made of n input/output pairs, where for each pair $i \in [n]$, $x_i \in \mathcal{X} \subset \mathbb{R}^d$ and $y_i \in \mathcal{Y} \subset \mathbb{R}^c$. This covers in particular the regression or binary classification settings, where $c = 1$, or the multivariate regression and multiclass classification setting, where y_i is an embedding of the class of x_i in \mathbb{R}^c , e.g., the one-hot encoding by taking c equal to the total number of classes and letting $y_i \in \{0, 1\}^c$ be the binary vector with all entries equal to zero except for the one corresponding to the class of x_i . We further denote the mean input and output as

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i, \quad \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i,$$

and the empirical variance and covariance matrices of inputs and outputs as

$$\Sigma_{xx} = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(x_i - \bar{x})^\top, \quad \Sigma_{xy} = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})^\top, \quad \Sigma_{yy} = \frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})(y_i - \bar{y})^\top.$$

Our goal is to learn from S_n a function $f : \mathcal{X} \rightarrow \mathbb{R}^c$ to predict the output corresponding to any new input $x \in \mathcal{X}$ via $\rho(f(x))$, where $\rho : \mathbb{R}^c \rightarrow \mathcal{Y}$ maps an \mathbb{R}^c -valued prediction to an element of \mathcal{Y} ; standard mappings include the identity $\rho(y) = y$ for regression problems, and the softmax operator $\rho(y)_i = e^{y_i} / \left(\sum_{j=1}^c e^{y_j} \right)$ for multiclass classification problems. For that purpose, we formulate the inference problem as an optimization problem:

$$\min_{f \in \mathcal{H}} \mathcal{E}(f), \tag{1}$$

where \mathcal{H} is a class of candidate functions, such as linear functions or deep neural networks, and $\mathcal{E}(f)$ is a risk functional that depends on S_n . The most standard risk used in machine learning is the empirical risk, defined for any loss function $\ell : \mathcal{Y} \times \mathbb{R}^c \rightarrow \mathbb{R}$ by:

$$\mathcal{E}^{\text{Empirical}}(f) = \frac{1}{n} \sum_{i=1}^n \ell(y_i, f(x_i)). \tag{2}$$

Solving (1) with the empirical risk (2) is often called *empirical risk minimization* (ERM), and is typically performed in practice by first-order numerical optimization such as stochastic gradient descent (Bottou and Bousquet, 2008). Standard losses ℓ include the squared error (in regression) and the cross-entropy loss applied to the softmax mapping (in classification, assuming that $\forall y \in \mathcal{Y}, y^\top \mathbf{1}_c = 1$, which is true for one-hot encoded classes and their convex combinations):

$$\forall (y, u) \in \mathcal{Y} \times \mathbb{R}^c, \quad \ell^{\text{SE}}(y, u) = \frac{1}{2} \|y - u\|^2, \quad \ell^{\text{CE}}(y, u) = \log \left(\sum_{i=1}^c e^{u_i} \right) - y^\top u. \quad (3)$$

Mixup. Instead of minimizing the empirical risk (2), Mixup (Zhang et al., 2018) creates new random input/output samples by taking convex combinations of pairs of training samples, and minimizes the corresponding empirical risk. With our notations, Mixup therefore minimizes the following *Mixup risk* over $f \in \mathcal{H}$:

$$\mathcal{E}^{\text{Mixup}}(f) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}_\lambda \ell(\lambda y_i + (1 - \lambda)y_j, f(\lambda x_i + (1 - \lambda)x_j)), \quad (4)$$

where $\lambda \sim \text{Beta}(\alpha, \alpha)$, and α is a parameter of Mixup. The minimization of (4) is typically performed by stochastic gradient descent, where λ is sampled at each iteration to obtain a stochastic gradient. In practice, Zhang et al. (2018) suggest to sample minibatches of training pairs, and generate Mixup random pairs within the minibatch, which also produces a stochastic gradient of (4).

3. Mixup as a perturbed ERM

The Mixup risk (4) is defined as a sum over pairs of samples, making a comparison with standard ERM approaches (2) not direct. The following result shows that the Mixup risk can be equivalently rewritten as a standard empirical risk, over modified input/output pairs (as in the third plot of Figure 1), subject to random perturbations.

Theorem 1 *Let $\theta \sim \text{Beta}_{[\frac{1}{2}, 1]}(\alpha, \alpha)$ and $j \sim \text{Unif}([n])$ be two random variables with $\alpha > 0$, $n > 0$ and let $\bar{\theta} = \mathbb{E}_\theta \theta$. For any training set \mathcal{S}_n , let $(\tilde{x}_i, \tilde{y}_i)$ for any $i \in [n]$ be the modified input/output pair given by*

$$\begin{cases} \tilde{x}_i &= \bar{x} + \bar{\theta}(x_i - \bar{x}), \\ \tilde{y}_i &= \bar{y} + \bar{\theta}(y_i - \bar{y}), \end{cases} \quad (5)$$

and $(\delta_i, \varepsilon_i)$ be the random perturbations given by:

$$\begin{cases} \delta_i &= (\theta - \bar{\theta})x_i + (1 - \theta)x_j - (1 - \bar{\theta})\bar{x}, \\ \varepsilon_i &= (\theta - \bar{\theta})y_i + (1 - \theta)y_j - (1 - \bar{\theta})\bar{y}. \end{cases} \quad (6)$$

Then for any $i \in [n]$, $\mathbb{E}_{\theta, j} \delta_i = \mathbb{E}_{\theta, j} \varepsilon_i = 0$, and for any function $f \in \mathcal{H}$,

$$\mathcal{E}^{\text{Mixup}}(f) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\theta, j} \ell(\tilde{y}_i + \varepsilon_i, f(\tilde{x}_i + \delta_i)). \quad (7)$$

Both δ_i and ε_i are random vectors because they are functions of θ and j in (6), which are themselves random variables. We hence use the notation $\mathbb{E}_{\theta,j}$ in (7). Note also $\bar{\theta} \in [1/2, 1]$ meaning that the transformation from (x_i, y_i) to $(\tilde{x}_i, \tilde{y}_i)$ in (5) shrinks the inputs and the outputs towards their mean.

Theorem 1 and the expression (7) of the Mixup risk allow us to re-interpret Mixup as a combination of two standard techniques: (i) transforming each input/output pair (x_i, y_i) into $(\tilde{x}_i, \tilde{y}_i)$, and (ii) adding zero-mean random perturbations $(\delta_i, \varepsilon_i)$ to each transformed pair, before minimizing the empirical risk. This helps us to understand the effects of training a model with Mixup by studying each technique and their interaction. In particular, perturbing input data is a classical approach to regularize ERM estimators (Bishop, 1995; Srivastava et al., 2014; Wager et al., 2013; Wei et al., 2020), and we study in detail in the next section the particular regularization induced by the Mixup perturbations on both inputs and outputs, before interpreting the resulting regularization aspects of Mixup due to both data transformation and perturbation in Section 5.

4. The regularization effects of Mixup

We now study the effect of the random perturbations $(\delta_i, \varepsilon_i)$ for $i \in [n]$ in the Mixup risk (7). While perturbing inputs with additive or multiplicative noise (e.g., dropout), and independently perturbing outputs (resulting, e.g., in label smoothing) have been widely studied, the Mixup perturbation (7) is unique in the sense that it is applied to both inputs and outputs simultaneously, and that the input and output perturbations are not independent from each other by (6). In order to study the regularization effect of these perturbations, we first characterize the covariance structure among the input and output perturbations.

Lemma 2 *Let $\bar{\theta}$ and σ^2 be respectively the mean and variance of a $\text{Beta}_{[1/2, 1]}(\alpha, \alpha)$ distributed random variable, and $\gamma^2 = \sigma^2 + (1 - \bar{\theta})^2$. For any $i \in [n]$, let*

$$\begin{aligned}\Sigma_{\tilde{x}\tilde{x}}^{(i)} &= \frac{\sigma^2(\tilde{x}_i - \bar{x})(\tilde{x}_i - \bar{x})^\top + \gamma^2 \Sigma_{\tilde{x}\tilde{x}}}{\bar{\theta}^2}, \\ \Sigma_{\tilde{y}\tilde{y}}^{(i)} &= \frac{\sigma^2(\tilde{y}_i - \bar{y})(\tilde{y}_i - \bar{y})^\top + \gamma^2 \Sigma_{\tilde{y}\tilde{y}}}{\bar{\theta}^2}, \\ \Sigma_{\tilde{x}\tilde{y}}^{(i)} &= \frac{\sigma^2(\tilde{x}_i - \bar{x})(\tilde{y}_i - \bar{y})^\top + \gamma^2 \Sigma_{\tilde{x}\tilde{y}}}{\bar{\theta}^2}.\end{aligned}\tag{8}$$

Then, for any $i \in [n]$, the random perturbations defined in (6) satisfy

$$\mathbb{E}_{\theta,j} \delta_i \delta_i^\top = \Sigma_{\tilde{x}\tilde{x}}^{(i)}, \quad \mathbb{E}_{\theta,j} \varepsilon_i \varepsilon_i^\top = \Sigma_{\tilde{y}\tilde{y}}^{(i)}, \quad \text{and} \quad \mathbb{E}_{\theta,j} \delta_i \varepsilon_i^\top = \Sigma_{\tilde{x}\tilde{y}}^{(i)}.\tag{9}$$

Following recent lines of work that interpret various random perturbations such as dropout as regularization (Wager et al., 2013; Wei et al., 2020), we can now introduce and study an approximate Mixup risk:

$$\mathcal{E}_Q^{\text{Mixup}}(f) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\theta,j} \ell_Q^{(i)}(\tilde{y}_i + \varepsilon_i, f(\tilde{x}_i + \delta_i)),\tag{10}$$

obtained by replacing the loss function $\ell(\tilde{y}, f(\tilde{x}))$ by a second-order quadratic Taylor approximation near each modified input/output training pairs $(\tilde{x}_i, \tilde{y}_i)$, namely, for any $i \in [n]$ and $(\delta, \varepsilon) \in \mathcal{X} \times \mathcal{Y}$:

$$\begin{aligned} \ell_Q^{(i)}(\tilde{y}_i + \varepsilon, f(\tilde{x}_i + \delta)) &= \ell(\tilde{y}_i, f(\tilde{x}_i)) + \nabla_y \ell(\tilde{y}_i, f(\tilde{x}_i)) \varepsilon + \nabla_u \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla_x f(\tilde{x}_i) \delta \\ &\quad + \frac{1}{2} \left\langle \delta \delta^\top, \nabla f(\tilde{x}_i)^\top \nabla_{uu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) + \nabla_u \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla^2 f(\tilde{x}_i) \right\rangle \\ &\quad + \frac{1}{2} \left\langle \varepsilon \varepsilon^\top, \nabla_{yy}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \right\rangle + \left\langle \varepsilon \delta^\top, \nabla_{yu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) \right\rangle, \end{aligned} \quad (11)$$

assuming both ℓ and f are twice continuously differentiable. Due to its quadratic form as a function of input and output perturbations, the approximate Mixup risk (10) can be re-expressed as a regularized ERM risk, as shown in the next result. We note that the expression we derive is in fact valid for *any* joint perturbation of the inputs and outputs with covariance structure given in (9).

Theorem 3 *For any twice continuously differentiable loss $\ell(y, u)$, the approximate Mixup risk at any twice differentiable $f \in \mathcal{H}$ satisfies*

$$\mathcal{E}_Q^{\text{Mixup}}(f) = \frac{1}{n} \sum_{i=1}^n \ell(\tilde{y}_i, f(\tilde{x}_i)) + R_1(f) + R_2(f) + R_3(f) + R_4(f), \quad (12)$$

where

$$\begin{aligned} R_1(f) &= \frac{1}{2n} \sum_{i=1}^n \left\| \left(\nabla f(\tilde{x}_i) - J^{(i)} \right)^\top \left(\nabla_{uu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \right)^{\frac{1}{2}} \right\|_{\Sigma_{\tilde{x}\tilde{x}}^{(i)}}^2, \\ R_2(f) &= \frac{1}{2n} \sum_{i=1}^n \left\langle \Sigma_{\tilde{x}\tilde{x}}^{(i)}, \nabla_u \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla^2 f(\tilde{x}_i) \right\rangle, \\ R_3(f) &= -\frac{1}{2n} \sum_{i=1}^n \left\| \Sigma_{\tilde{y}\tilde{y}}^{(i)} \nabla_{yu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \left(\nabla_{uu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \right)^{-\frac{1}{2}} \right\|_{\left(\Sigma_{\tilde{x}\tilde{x}}^{(i)} \right)^{-1}}^2, \\ R_4(f) &= \frac{1}{2n} \sum_{i=1}^n \left\langle \Sigma_{\tilde{y}\tilde{y}}^{(i)}, \nabla_{yy}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \right\rangle, \end{aligned}$$

and

$$\forall i \in [n], \quad J^{(i)} = - \left(\nabla_{uu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \right)^{-1} \nabla_{uy}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \Sigma_{\tilde{y}\tilde{x}}^{(i)} \left(\Sigma_{\tilde{x}\tilde{x}}^{(i)} \right)^{-1}. \quad (13)$$

Theorem 3 captures the effect of the random perturbations in Mixup as a sum of four penalty terms $R_i(f)$ for $i \in [4]$. They regularize the simple ERM risk applied on the modified inputs \tilde{x}_i and smoothed outputs \tilde{y}_i . Before discussing the accuracy and practical consequences of this reformulation of Mixup as regularized empirical risk minimization on modified data in the next Section, we now derive the details of this approximation for the cross-entropy, logistic and squared error losses. We begin by presenting the results for the cross-entropy loss:

Corollary 4 Let $\mathcal{S} : \mathbb{R}^c \rightarrow \mathbb{R}^c$ be the softmax operator, i.e., for any $i \in [c]$ and $u \in \mathbb{R}^c$, $\mathcal{S}(u)_i = e^{u_i} / \sum_{j=1}^c e^{u_j}$, and let $H(u) = \text{diag}(\mathcal{S}(u)) - \mathcal{S}(u)\mathcal{S}(u)^\top \in \mathbb{R}^{c \times c}$. The approximate Mixup risk for the cross-entropy loss satisfies

$$\mathcal{E}_Q^{\text{Mixup}}(f) = \frac{1}{n} \sum_{i=1}^n \ell^{\text{CE}}(\tilde{y}_i, f(\tilde{x}_i)) + R_1^{\text{CE}}(f) + R_2^{\text{CE}}(f) + R_3^{\text{CE}}(f),$$

where

$$\begin{aligned} R_1^{\text{CE}}(f) &= \frac{1}{2n} \sum_{i=1}^n \left\| \left(\nabla f(\tilde{x}_i) - J^{(i)} \right)^\top H(f(\tilde{x}_i))^{\frac{1}{2}} \right\|_{\Sigma_{\tilde{x}\tilde{x}}^{(i)}}^2, \\ R_2^{\text{CE}}(f) &= \frac{1}{2n} \sum_{i=1}^n \left\langle \Sigma_{\tilde{x}\tilde{x}}^{(i)}, (\mathcal{S}(f(\tilde{x}_i)) - \tilde{y}_i)^\top \nabla^2 f(\tilde{x}_i) \right\rangle, \\ R_3^{\text{CE}}(f) &= -\frac{1}{2n} \sum_{i=1}^n \left\| \Sigma_{\tilde{x}\tilde{y}}^{(i)} H(f(\tilde{x}_i))^{-\frac{1}{2}} \right\|_{\left(\Sigma_{\tilde{x}\tilde{x}}^{(i)}\right)^{-1}}^2, \end{aligned}$$

with

$$\forall i \in [n], \quad J^{(i)} = H(f(\tilde{x}_i))^{-1} \Sigma_{\tilde{y}\tilde{x}}^{(i)} \left(\Sigma_{\tilde{x}\tilde{x}}^{(i)} \right)^{-1}. \quad (14)$$

In the binary classification setting, minimizing the empirical cross-entropy risk over $f : \mathcal{X} \rightarrow \mathbb{R}^2$ after one-hot encoding of the two possible classes in \mathbb{R}^2 as $(0, 1)^\top$ and $(1, 0)^\top$ is equivalent to minimizing the following well-known logistic loss over $f : \mathcal{X} \rightarrow \mathbb{R}$ after encoding the two classes in \mathbb{R} as 0 and 1:

$$\ell^{\text{LR}}(y, u) = \log(1 + e^u) - yu. \quad (15)$$

The regularization effect of Mixup in that case is detailed in the following result:

Corollary 5 Let $s : \mathbb{R} \rightarrow \mathbb{R}$ be the sigmoid operator, i.e., for any $u \in \mathbb{R}$, $s(u) = (1 + e^{-u})^{-1}$, and let $v(u) = s(u)(1 - s(u)) \in \mathbb{R}$. The approximate Mixup risk for the logistic regression loss satisfies

$$\mathcal{E}_Q^{\text{Mixup}}(f) = \frac{1}{n} \sum_{i=1}^n \ell^{\text{LR}}(\tilde{y}_i, f(\tilde{x}_i)) + R_1^{\text{LR}}(f) + R_2^{\text{LR}}(f) + R_3^{\text{LR}}(f),$$

where

$$\begin{aligned} R_1^{\text{LR}}(f) &= \frac{1}{2n} \sum_{i=1}^n v(f(\tilde{x}_i)) \left\| \nabla f(\tilde{x}_i) - J^{(i)} \right\|_{\Sigma_{\tilde{x}\tilde{x}}^{(i)}}^2, \\ R_2^{\text{LR}}(f) &= \frac{1}{2n} \sum_{i=1}^n (s(f(\tilde{x}_i)) - \tilde{y}_i) \left\langle \Sigma_{\tilde{x}\tilde{x}}^{(i)}, \nabla^2 f(\tilde{x}_i) \right\rangle, \\ R_3^{\text{LR}}(f) &= -\frac{1}{2n} \sum_{i=1}^n v(f(\tilde{x}_i))^{-1} \Sigma_{\tilde{y}\tilde{x}}^{(i)} \left(\Sigma_{\tilde{x}\tilde{x}}^{(i)} \right)^{-1} \Sigma_{\tilde{x}\tilde{y}}^{(i)}, \end{aligned}$$

with

$$\forall i \in [n], \quad J^{(i)} = \frac{\Sigma_{\tilde{y}\tilde{x}}^{(i)} \left(\Sigma_{\tilde{x}\tilde{x}}^{(i)} \right)^{-1}}{v(f(\tilde{x}_i))}. \quad (16)$$

The next result summarizes the form of the approximate Mixup risk in the case of the squared error loss, and shows in particular that Mixup has no effect for linear least-squares regression models.

Corollary 6 *The approximate Mixup risk for the squared error loss satisfies*

$$\mathcal{E}_Q^{Mixup}(f) = \frac{1}{n} \sum_{i=1}^n \ell^{SE}(\tilde{y}_i, f(\tilde{x}_i)) + R_1^{SE}(f) + R_2^{SE}(f) + C, \quad (17)$$

where C is a constant independent of f and

$$R_1^{SE}(f) = \frac{1}{2n} \sum_{i=1}^n \|\nabla f(\tilde{x}_i) - J^{(i)}\|_{\Sigma_{\tilde{x}\tilde{x}}^{(i)}}^2 \quad \text{and} \quad R_2^{SE}(f) = \frac{1}{2n} \sum_{i=1}^n \left\langle \Sigma_{\tilde{x}\tilde{x}}^{(i)}, (f(\tilde{x}_i) - \tilde{y}_i)^\top \nabla^2 f(\tilde{x}_i) \right\rangle,$$

with

$$\forall i \in [n], \quad J^{(i)} = \Sigma_{\tilde{y}\tilde{x}}^{(i)} \left(\Sigma_{\tilde{x}\tilde{x}}^{(i)} \right)^{-1}. \quad (18)$$

In particular, when we consider linear models with intercept of the form $f_{W,b}(x) = Wx + b$ for $(W, b) \in \mathbb{R}^{c \times d} \times \mathbb{R}^c$, then the exact Mixup risk satisfies

$$\mathcal{E}^{Mixup}(f_{W,b}) = \frac{2\sigma^2 + 2\bar{\theta}^2 + (1 - \bar{\theta})^2}{2n} \sum_{i=1}^n \ell^{SE}(y_i, f_{W,\bar{b}}(x_i)) + \|b - \bar{b}\|^2 + C, \quad (19)$$

where C is a constant that does not depend on (W, b) and $\bar{b} = \bar{y} - W\bar{x}$. Consequently, the linear model that minimizes \mathcal{E}^{Mixup} is the standard multivariate ordinary least squares (MOLS) predictor that minimizes \mathcal{E}^{ERM} on the original data, i.e., Mixup has no effect on linear least-squares regression.

5. Discussion and experiments

Let us now discuss how our analysis relates to a recent similar study by Zhang et al. (2021), and empirically assess the validity of our analysis and the regularization properties of Mixup it suggests. To support our discussion, we provide empirical results on CIFAR-10/100 and ImageNet for different networks (LeNet, ResNet-34/50). For each experimental result we report mean and 95% confidence interval using 10 repetitions (unless stated otherwise). All details about experiments are provided in Appendix B, together with other experiments on the simpler setting of learning on the two-moon dataset with random features.

Comparison with related work. Zhang et al. (2021) independently published a similar analysis of the regularization effect of Mixup. Both works provide complementary and coherent views of the effect of Mixup on generalization and robustness, and differ in a few technical aspects that we clarify here. First, we provide an analysis valid for any loss $\ell(y, u)$, where y and u are respectively the predicted and true outputs, while Zhang et al. (2021)

restrict their analysis to the losses of the form $\ell(y, u) = h(u) - yu$ for some function h . Second, while both works use a second-order Taylor expansion to approximate the loss at a Mixup example, the expansions are different since we do a Taylor expansion near the expected value of the Mixup example, while Zhang et al. (2021, Lemma 3.1) do a Taylor expansion near a non-Mixup example. One consequence is that our Taylor approximation converges to the exact Mixup risk when the amount of Mixup reduces (i.e., when $\alpha \rightarrow +\infty$ in the $Beta(\alpha, \alpha)$ distribution of the mixing parameter), while the one used by Zhang et al. (2021) does not². Third, the generalization analysis of Zhang et al. (2021) for generalized linear models (GLM) is performed for a variant of Mixup where the mixed input is $x_i + (1/\lambda - 1)x_j$, while we focus on the standard Mixup $\lambda x_i + (1 - \lambda)x_j$. A consequence of this difference is that we identify the importance of rescaling data at test time for standard Mixup (see below), while no such rescaling is needed in the variant considered by Zhang et al. (2021). Finally, Zhang et al. (2021) explore the link between Mixup-induced regularization and adversarial robustness and generalization through Rademacher complexity analysis, which we do not explore in this work but could be done in a similar manner.




Fig. 2: From left to right: train loss, test loss and accuracy during optimization of LeNet on CIFAR-10 with Mixup and approximate Mixup.

Validity of the Taylor approximation. To analyze the regularization effect of Mixup, we used a quadratic approximation of the loss function (11). We note that compared to similar approximations that have been proposed to study the regularization effect of input perturbation only, such as dropout (Wager et al., 2013; Wei et al., 2020), we must include in the Taylor expansion all second-order terms involving the input perturbation only (term with $\delta\delta^\top$), the output perturbation only (term with $\varepsilon\varepsilon^\top$), and their interaction (term with $\varepsilon\delta^\top$). In the absence of output perturbation (e.g., in the case of dropout), only the term in $\delta\delta^\top$ matters, and in the absence of correlation between input and output perturbation (e.g., dropout combined with independent label smoothing), then the term in $\varepsilon\delta^\top$ does not matter either. Mixup is unique in the correlation it creates between input and output perturbations, which is captured by the interaction term with $\varepsilon\delta^\top$ in (11). Regarding the validity of the Taylor approximation, we note that, as for similar work on input perturbation, the approximate Mixup risk (10) is only a good approximation to the Mixup risk for “small” perturbations; as noted by Wei et al. (2020, Annex A.2), though, this often remains valid even for “large”

2. More precisely, the Taylor approximation remainder in Zhang et al. (2021, Lemma 3.1) does not vanish since it is a Taylor expansion near 0 of a random variable that follows a $Beta(\alpha + 1, \alpha + 1)$ distribution, i.e., that takes values close to 1 with probability 1/2 when α goes to $+\infty$.

input perturbation followed by a linear transformation layer. To support empirically the validity of the approximation, Figure 2 shows the training and test performance of training a LeNet on CIFAR-10 using Mixup (minimizing (4)), and using the approximate Mixup formulation (minizing (12)), where we dropped the term $R_2(f)$ in the regularization since it empirically induces numerical instability due to its non-convexity (see also Wei et al., 2020, for a discussion about discarding the Hessian regularization). We can see how when training




Fig. 3: Evaluations of the regularization terms of the mixup approximation (left) and of the loss on modified data (right) for functions learned with Mixup, ERM and ERM on modified data.

using the approximate Mixup formulation, we learn functions which mimic, both in training and test, the performance of functions learned when training with the standard Mixup procedure. To further mark the validity of the approximation and decouple the contributions of the data transformation (5) and the input perturbation (6), we evaluate for functions learned with Mixup (4), ERM (2) and ERM on modified data:

$$\min_{f \in \mathcal{Z}} \frac{1}{n} \sum_{i=1}^n \ell(\tilde{y}_i, f(\tilde{x}_i)), \quad (20)$$

the regularization terms of the approximation (12) and the loss on modified data (20). From Figure 3 we see that the functions learned with Mixup are the ones with the smallest values of the regularization terms but not the smallest loss on modified data, which confirms that the model trained with Mixup finds a trade-off between empirical risk and the regularization we study.

Data modification at test time. By Theorem 1, we see that Mixup implicitly shrinks inputs towards their mean since the Mixup risk involves the empirical risk over modified inputs \tilde{x}_i and outputs \tilde{y}_i . In particular, this means that the functions that the standard Mixup procedure learns are not functions from the space of input points \mathcal{X} to the output points \mathcal{Y} , but it learns functions from $\tilde{\mathcal{X}}$ to $\tilde{\mathcal{Y}}$, which are spaces defined by the training points and the hyperparameter α of Mixup as

$$\begin{cases} \tilde{\mathcal{X}} = \{\bar{x} + \bar{\theta}(x - \bar{x}) \mid \forall x \in \mathcal{X}\}, \\ \tilde{\mathcal{Y}} = \{\bar{y} + \bar{\theta}(y - \bar{y}) \mid \forall y \in \mathcal{Y}\}, \end{cases} \quad (21)$$

where \bar{x} and \bar{y} are respectively the average training inputs and outputs points, and $\bar{\theta} = \mathbb{E}_\theta \theta$ as defined in Theorem 1. An important consequence is that the function $f : \mathcal{X} \rightarrow \mathcal{Y}$ estimated

by Mixup, should ideally be applied at test time to transformed data, to map the test point $x_{\text{test}} \in X$ to $\tilde{x}_{\text{test}} \in \tilde{X}$, and the evaluation of the function $f(\tilde{x}_{\text{test}}) \in \tilde{Y}$ should be mapped back to \mathcal{Y} . In details, the prediction for point x_{test} should be

$$\text{pred}_f(x_{\text{test}}) = \bar{y} (1 - 1/\bar{\theta}) + 1/\bar{\theta} \cdot f(\bar{\theta}x_{\text{test}} + (1 - \bar{\theta})\bar{x}). \quad (22)$$

Algorithm 1 python code to evaluate according to (22) functions learned with mixup

Input: X_{test} *Tensor* of test points to evaluate, trained_model *Model* trained using mixup, α *float* hyperparameter used by mixup during training, $X_{\text{train}}, Y_{\text{train}}$ *Tensor* of points used for training with one-hot encoded labels.

```
import scipy.special as sc
import torch

X_bar = torch.mean(X_train, dim=0, keepdim=True)
Y_bar = torch.mean(Y_train, dim=0, keepdim=True)

# expectation of a truncated beta distribution in [0.5, 1]
theta_bar = 1. - sc.betainc(alpha + 1., alpha, .5)

def predict_mixup(X_test, trained_model):
    f_X = trained_model.forward((1. - theta_bar) * X_bar + theta_bar * X_test)
    return Y_bar * (1. - 1. / theta_bar) + f_X / theta_bar
```

For centered training data ($\bar{x} = \bar{y} = 0$) and homogeneous functions ($f(ux) = uf(x)$ for any $(u, x) \in \mathbb{R}^+ \times \mathcal{X}$, e.g., linear models or neural networks with ReLU activation and linear transformations), this has no impact as $\text{pred}_f(x) = f(x)$ in that case. For more general models, however, (22) may be a better predictor than f . For example, we clearly see in Figure 1 that the asymptotically Bayes optimal classifier under the Mixup distribution matches the one under the empirical distribution of the modified data (up to regularization effects), and not of the original data. Interestingly, when the classes are balanced, i.e., $\bar{y} = \frac{1}{c}\mathbf{1}_c$, the transformation in (22) adds the same constant to each of the c entries of f . In particular, in the multi-class setting, since the softmax is invariant to a constant in the logits, (22) becomes equivalent to a scaling of the logits, commonly referred to as temperature scaling (Guo et al., 2017). While temperature scaling is traditionally tuned with a validation set (Guo et al., 2017), mixup automatically sets this value, according to the distribution of θ . To point out the advantages of using (22), we compare in Figure 4 the performance of ERM, standard Mixup (for different α values), the same Mixup but with the proper data transformation (22) at test time and the same data transformation applied to the ERM estimator. The trained networks are ResNet-34 for CIFAR-10 and CIFAR-100, and ResNet-50 for ImageNet. For CIFAR-10 and 100, we observe overall benefits of using the data transformation for evaluating the functions learned with Mixup: higher test accuracy, lower test loss, and lower expected calibration error (ECE). For ImageNet we have the same benefits, with the only exception of the test loss and the ECE for very low values of α . Notice indeed that for small values of α the data transformation (22) has a smaller impact than for bigger values of α , as $\lim_{\alpha \rightarrow 0} \bar{\theta} = 1$ while $\lim_{\alpha \rightarrow +\infty} \bar{\theta} = \frac{1}{2}$. Thus, as observed empirically,




Fig. 4: Test accuracy, test cross-entropy loss and test expected calibration error (first, second and third column respectively) on the CIFAR-10, CIFAR-100, ImageNet datasets (first, second and third row respectively). We report the mean and the standard error over 10 repetitions for CIFAR-10 and CIFAR-100 and 5 repetitions for ImageNet. The training was done with ResNet-34 for CIFAR-10, CIFAR-100, and with ResNet-50 for ImageNet.

the “correction” (22) is more important, as it brings bigger improvements, when training Mixup with bigger α . Finally, we can observe that when the data transformation (22) is applied to ERM, performance always deteriorate. This supports that (22) is a mixup specific improvement. Algorithm 1 shows the few lines of codes that implement the new prediction procedure (22).

Data modification for out-of-distribution data. Even though our theoretical analysis holds only for in-distribution data, we now empirically investigate whether the benefits of rescaling data at test time that we observe for models trained with Mixup also hold when the test time data come from a different distribution, a setting called out-of-distribution (OOD) data (Hendrycks and Gimpel, 2017). Indeed, standard Mixup is known to provide some benefits in out-of-distribution settings compared to models simply trained by ERM, so it is interesting to assess whether the data modification scheme we propose can further boost the accuracy and calibration of models trained with Mixup in that setting. In particular we consider CIFAR-10-C, CIFAR-100-C, ImageNet-C (Hendrycks and Dietterich, 2019), ImageNet-A (Hendrycks et al., 2021), ImageNetV2 (Recht et al., 2019), ImageNet-Vid-Robust, YTBB-Robust (Shankar et al., 2021), ObjectNet (Barbu et al., 2019). These benchmark datasets are designed by systematically perturbing in a controlled way the in-distribution data




Fig. 5: Test accuracy, test cross-entropy loss and test expected calibration error (first, second and third column respectively) on the CIFAR-10-C, CIFAR-100-C (first, and second row respectively). We report the mean and the standard error over 10 repetitions. The training was done with ResNet-34 on the standard CIFAR-10, CIFAR-100.

(CIFAR-10, CIFAR-100 and ImageNet), e.g., by adding noise or applying transformations to the images. Table 2 details the performance of models trained with Mixup for $\alpha = 0.25$




Fig. 6: Difference in test performance between Mixup and rescaled Mixup on CIFAR-10-C, CIFAR-100-C for the different corruption intensities. The training was done with ResNet-34 on the standard CIFAR-10, CIFAR-100 dataset and the performance metrics are accuracy, cross-entropy loss, expected calibration error.

and $\alpha = 0.5$, with or without data modification at test time, on out-of-distribution datasets derived from ImageNet, while Figure 5 shows similar results for benchmarks derived from

CIFAR-10 and CIFAR-100. We see that for most datasets the rescaling improves accuracy and some of the other metrics, but for others and in particular CIFAR-10-C, CIFAR-100-C and ImageNet-C, the rescaling worsen almost all metrics with respect of the standard Mixup. We now investigate how the performance of Mixup with or without rescaling differs with respect to the intensity of the noise.

ImageNet	Metric	mixup(rescaled)		mixup	
		$\alpha = 0.25$		$\alpha = 0.5$	
standard	Acc	0.7764 \pm 0.0022	0.7754 \pm 0.0021	0.7780 \pm 0.0006	0.7766 \pm 0.0009
	CE-loss	0.8820 \pm 0.0092	0.8964 \pm 0.0113	0.8852 \pm 0.0046	0.9482 \pm 0.0188
	ECE	0.0291 \pm 0.0031	0.0236 \pm 0.0058	0.0280 \pm 0.0040	0.0707 \pm 0.0127
A	Acc	0.0165 \pm 0.0011	0.0165 \pm 0.0011	0.0229 \pm 0.0004	0.0225 \pm 0.0005
	CE-loss	7.3395 \pm 0.0884	6.8075 \pm 0.0687	7.1010 \pm 0.0879	6.3738 \pm 0.0506
	ECE	0.4322 \pm 0.0072	0.3629 \pm 0.0075	0.4318 \pm 0.0110	0.3161 \pm 0.0117
C	Acc	0.4638 \pm 0.0021	0.4662 \pm 0.0019	0.4719 \pm 0.0015	0.4761 \pm 0.0011
	CE-loss	2.8050 \pm 0.0140	2.7374 \pm 0.0094	2.7842 \pm 0.0164	2.7206 \pm 0.0136
	MCE	0.6824 \pm 0.0025	0.6795 \pm 0.0022	0.6735 \pm 0.0018	0.6685 \pm 0.0014
	ECE	0.1040 \pm 0.0046	0.0652 \pm 0.0019	0.1096 \pm 0.0054	0.0793 \pm 0.0061
V2	Acc	0.6560 \pm 0.0028	0.6543 \pm 0.0029	0.6593 \pm 0.0007	0.6575 \pm 0.0011
	CE-loss	1.5208 \pm 0.0118	1.5048 \pm 0.0122	1.5130 \pm 0.0057	1.5295 \pm 0.0170
	ECE	0.0746 \pm 0.0053	0.0344 \pm 0.0030	0.0715 \pm 0.0079	0.0417 \pm 0.0096
Vid-Robust	Acc (pm-k)	0.5254 \pm 0.0043	0.5244 \pm 0.0036	0.5224 \pm 0.0023	0.5216 \pm 0.0035
YTBB-Robust	Acc (pm-k)	0.4738 \pm 0.0037	0.4717 \pm 0.0035	0.4705 \pm 0.0021	0.4692 \pm 0.0017
ObjectNet	Acc	0.2756 \pm 0.0035	0.2751 \pm 0.0034	0.2837 \pm 0.0010	0.2825 \pm 0.0011

Table 2: Test performance of Mixup and rescaled Mixup on different OOD versions of ImageNet. The training was done with ResNet-50 on the standard ImageNet dataset and the performance metrics are accuracy, cross-entropy loss, expected calibration error and mean corruption error (Hendrycks and Dietterich, 2019). For each metric means and standard errors over 10 repetitions are reported.

In details, the CIFAR-10-C and CIFAR-100-C data are 95 different corrupted versions each of the original CIFAR-10, CIFAR-100 test sets: 19 different corruption types, for 5 different growing corruption intensities. The results that we reported in Figure 5 are the averages of the performance of these 95 test sets. In Figure 6 instead, for Mixup and rescaled Mixup, we compute the average of the different metrics (test accuracy, test cross-entropy loss, ECE) over the 19 corruption types for each corruption intensity, and we report the difference in performance between Mixup and rescaled Mixup. We can see that as the noise intensity grows the gap between the two methods always grows in favor of standard Mixup, and that only for a few metrics and settings with low noise the rescaled version is better than the standard one. This observations encourage future investigations of the effect of Mixup on OOD data³.

3. Codes for reproducing results for ImageNet and various OOD variants are available at <https://github.com/google/uncertainty-baselines/tree/master/baselines/imagenet>

Label smoothing. The transformation that modifies the original labels y_i onto \tilde{y}_i acts as some form of *label smoothing*, a technique known to often improve accuracy and calibration (Szegedy et al., 2016; Müller et al., 2019). The transformed labels \tilde{y}_i are indeed pulled towards the average label \bar{y} . Recall from Szegedy et al. (2016) that label smoothing consists in training a model on the perturbed version of the training labels defined as $y_i^{LS} = (1 - \varepsilon)y_i + \varepsilon u(i)$, where ε is a fixed scalar in $[0, 1]$ and $u(i)$ is a fixed distribution over the labels. It is easy to see that for $\varepsilon = (1 - \bar{\theta})$ and $u(i) = \bar{y}$ the two formulations coincide. This implies that Mixup implicitly performs label smoothing, and can benefit from this technique in terms of accuracy or calibration. In the following Proposition 7, we formally prove that, in the case of the cross entropy and linear models, label smoothing translates into an *increase* in the average entropy of the predictions, or, in other words, that predictions become less certain, as observed in practice. Like in Theorem 4, we use the softmax operator $\mathcal{S} : \mathbb{R}^c \mapsto \Delta_c$ defined for $j \in [c]$ by $\mathcal{S}(u)_j = e^{u_j} / (\sum_{k=1}^c e^{u_k})$:

Proposition 7 *Let us consider the following two classification problems with a cross-entropy loss and linear model $f(x) = Wx$ parameterized by $W \in \mathbb{R}^{c \times d}$,*

$$\min_{W \in \mathbb{R}^{c \times d}} \frac{1}{n} \sum_{i=1}^n \ell^{CE}(y_i, Wx_i) \tag{23}$$

and

$$\min_{W \in \mathbb{R}^{c \times d}} \frac{1}{n} \sum_{i=1}^n \ell^{CE}(\tilde{y}_i, Wx_i) \tag{24}$$

defined without and with label smoothing respectively, i.e., with $\tilde{y}_i = \bar{y} + \bar{\theta}(y_i - \bar{y}) \in \Delta_c$ for $i \in [n]$. Let us denote by W^* and W_{ls}^* a solution of (23) and (24) respectively, together with

$$p_i = \mathcal{S}(W^*x_i) \quad \text{and} \quad \tilde{p}_i = \mathcal{S}(W_{ls}^*x_i). \tag{25}$$

It holds that the average entropy of the predictions of W_{ls}^* is lower bounded as follows

$$\bar{\theta} \frac{1}{n} \sum_{i=1}^n \mathcal{Z}(p_i) + (1 - \bar{\theta}) \mathcal{Z}(\bar{y}) \leq \frac{1}{n} \sum_{i=1}^n \mathcal{Z}(\tilde{p}_i).$$

If predicting with W^* also reduces the entropy of the average predictor, i.e., $\frac{1}{n} \sum_{i=1}^n \mathcal{Z}(p_i) \leq \mathcal{Z}(\bar{y})$, then label smoothing increases the average entropy of the predictions:

$$\frac{1}{n} \sum_{i=1}^n \mathcal{Z}(p_i) \leq \frac{1}{n} \sum_{i=1}^n \mathcal{Z}(\tilde{p}_i).$$

To illustrate how both Mixup and label smoothing increase the entropy of predictions compared to ERM, we show on Figure 7 the histograms of the confidence of the estimators' predictions on test points, for a LeNet neural network trained on CIFAR-10. From the first plot, we notice how standard ERM produces very confident predictions, how label smoothing helps decreasing ERM confidence at test time, and how Mixup naturally produces even less confident predictions. From the second plot, we see that that approximate Mixup, like Mixup, produces less confident prediction, which confirms that the Mixup approximation we study captures well this behavior of Mixup.




Fig. 7: Histograms of confidence of predictions for models trained with different techniques.

Jacobian regularization. The first implicit regularization term $R_1(f)$ in Theorem 3 penalizes the discrepancy between $\nabla f(\tilde{x}_i)$ and $J^{(i)}$ given by (13). We recognize in $J^{(i)}$ the Jacobian of the standard MOLS model trained in the input space on the modified training set, with an increased weight for sample $(\tilde{x}_i, \tilde{y}_i)$ in $J^{(i)}$. Compared to, e.g., dropout regularization which penalizes the norm of ∇f at the training points, we therefore see that Mixup also regularizes the Jacobian of f but with a different and more informative implicit bias, namely, to mimic a good linear model in the input space. Furthermore, we note from the proof of Theorem 3 that this implicit bias results from the correlation between input and output noise, which may explain why independent Mixup in the input and output performs more poorly than standard Mixup (Zhang et al., 2018). While this regularization is similar across all points in the squared loss setting (Corollary 6), it is weighted by the Hessian $H(p(f(\tilde{x}_i)))$ in the cross-entropy loss (Corollary 4). Similar to dropout, this implies that this regularization vanishes when the prediction $p(f(\tilde{x}_i))$ is confidently near 0 or 1. In the Mixup case, though, the label smoothing effect discussed in the previous paragraph tends to prevent over-confident predictions on the training point (see Proposition 7 for a formal description of that property), therefore ensuring that the Jacobian regularization in $R_1(f)$ remains active even for “easy” points. This interaction between label smoothing (due to output Mixup) and Jacobian regularization (due to input Mixup) may explain why Mixup on inputs only performs poorly compared to Mixup on both inputs and outputs (Thulasidasan et al., 2019).

6. Conclusions

In this paper we have proposed the first theoretical analysis that explains the multiple regularization effects of Mixup. We have proved that training with Mixup is equivalent to learning on modified data with the injection of structured noise. Through a Taylor approximation, we have further shown that Mixup amounts to empirical risk minimization on modified points plus multiple regularization terms. Fascinatingly, our derivation shows that Mixup induces varied and complex effects, e.g., calibration, Jacobian regularization, label noise and normalization, while being a simple and cheap data augmentation technique. Further, we have shown how this analysis points out a missing rescaling procedure required to evaluate functions learned with Mixup, and we brought empirical evidence that implementing it improves accuracy and calibration. More broadly, we have studied how a specific combination

of data modification and noise injection leads to certain regularizers. An interesting research question is whether we can *reverse-engineer* this process, namely starting from possibly expensive regularizers and design the corresponding data augmentation technique emulating their effects at a lower computational cost.

Acknowledgments

The authors thank Jeremiah Zhe Liu for his feedback on an early version of this work.

Appendix

Appendix A. Proofs

A.1 Proof of Theorem 1

Proof To simplify notations, let us denote, for any $i, j \in [n]$ and $u \in [0, 1]$,

$$m_{ij}(u) = \ell(uy_i + (1-u)y_j, f(ux_i + (1-u)x_j)) .$$

The Mixup risk (4) can then be written as

$$\mathcal{E}^{\text{Mixup}}(f) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}_{\lambda} m_{ij}(\lambda), \quad \lambda \sim \text{Beta}(\alpha, \alpha). \quad (26)$$

We now separate the values of λ depending on whether or not they are above $1/2$ by expressing it as

$$\lambda = \pi\lambda_0 + (1-\pi)\lambda_1, \quad \lambda_0 \sim \text{Beta}_{[0, \frac{1}{2}]}(\alpha, \alpha), \quad \lambda_1 \sim \text{Beta}_{[\frac{1}{2}, 1]}(\alpha, \alpha), \quad \pi \sim \text{Ber}\left(\frac{1}{2}\right). \quad (27)$$

By symmetry of the $\text{Beta}(\alpha, \alpha)$ distribution around $1/2$, it is clear that λ defined in (27) follows a $\text{Beta}(\alpha, \alpha)$ distribution, and furthermore that $\lambda'_1 = 1 - \lambda_0$ follows, like λ_1 , a $\text{Beta}_{[\frac{1}{2}, 1]}(\alpha, \alpha)$ distribution. For any $i, j \in [n]$, we therefore get

$$\begin{aligned} \mathbb{E}_{\lambda} m_{ij}(\lambda) &= \mathbb{E}_{\lambda_0, \lambda_1, \pi} m_{ij}(\pi\lambda_0 + (1-\pi)\lambda_1) \\ &= \frac{1}{2} [\mathbb{E}_{\lambda_0} m_{ij}(\lambda_0) + \mathbb{E}_{\lambda_1} m_{ij}(\lambda_1)] \\ &= \frac{1}{2} [\mathbb{E}_{\lambda'_1} m_{ji}(\lambda'_1) + \mathbb{E}_{\lambda_1} m_{ij}(\lambda_1)] , \end{aligned}$$

where we used the fact that $m_{ij}(1-u) = m_{ji}(u)$ to get the third equality. Plugging this equality back into (26), we finally get

$$\begin{aligned}
 \mathcal{E}^{\text{Mixup}}(f) &= \frac{1}{2n^2} \sum_{i=1}^n \sum_{j=1}^n \left[\mathbb{E}_{\lambda'_1} m_{ji}(\lambda'_1) + \mathbb{E}_{\lambda_1} m_{ij}(\lambda_1) \right] \\
 &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}_{\lambda_1} m_{ij}(\lambda_1) \\
 &= \frac{1}{n} \sum_{i=1}^n \left[\frac{1}{n} \sum_{j=1}^n \mathbb{E}_{\lambda_1} m_{ij}(\lambda_1) \right] \\
 &= \frac{1}{n} \sum_{i=1}^n \ell_i,
 \end{aligned} \tag{28}$$

where

$$\ell_i = \mathbb{E}_{\theta,j} \ell(\theta y_i + (1-\theta)y_j, f(\theta x_i + (1-\theta)x_j)), \quad \theta \sim \text{Beta}_{[\frac{1}{2},1]}(\alpha, \alpha), \quad j \sim \text{Unif}([n]).$$

We now easily see that \tilde{x}_i and \tilde{y}_i defined in (5) satisfy

$$\begin{cases} \tilde{x}_i = \mathbb{E}_{\theta,j} [\theta x_i + (1-\theta)x_j], \\ \tilde{y}_i = \mathbb{E}_{\theta,j} [\theta y_i + (1-\theta)y_j], \end{cases}$$

and furthermore that δ_i and ε_i defined in (6) satisfy

$$\begin{cases} \tilde{\delta}_i = \theta x_i + (1-\theta)x_j - \mathbb{E}_{\theta,j} [\theta x_i + (1-\theta)x_j], \\ \tilde{\varepsilon}_i = \theta y_i + (1-\theta)y_j - \mathbb{E}_{\theta,j} [\theta y_i + (1-\theta)y_j], \end{cases}$$

from which we deduce that $\mathbb{E}_{\theta,j} \delta_i = \mathbb{E}_{\theta,j} \varepsilon_i = 0$ and

$$\ell_i = \mathbb{E}_{\theta,j} \ell(\tilde{y}_i + \tilde{\varepsilon}_i, f(\tilde{x}_i + \tilde{\delta}_i)).$$

Plugging this equality back into (28) concludes the proof. ■

A.2 Proof of Lemma 2

Proof From the definition of δ_i in (6), we easily get:

$$\begin{aligned}
 \mathbb{E}_{\theta,j} \delta_i \delta_i^\top &= \mathbb{E}_\theta (\theta - \bar{\theta})^2 x_i x_i^\top + \mathbb{E}_\theta [(1-\theta)^2] \mathbb{E}_j [x_j x_j^\top] + (1-\bar{\theta})^2 \bar{x} \bar{x}^\top \\
 &\quad + \mathbb{E}_\theta [(\theta - \bar{\theta})(1-\theta)] \mathbb{E}_j [x_i x_j^\top + x_j x_i^\top] - (1-\bar{\theta}) \mathbb{E}_\theta [1-\theta] \mathbb{E}_j [x_j \bar{x}^\top + \bar{x} x_j^\top] \\
 &= \sigma^2 x_i x_i^\top + \gamma^2 \mathbb{E}_j [x_j x_j^\top] + (1-\bar{\theta})^2 \bar{x} \bar{x}^\top - \sigma^2 (x_i \bar{x}^\top + \bar{x} x_i^\top) - 2(1-\bar{\theta})^2 \bar{x} \bar{x}^\top \\
 &= \sigma^2 (x_i x_i^\top - x_i \bar{x}^\top - \bar{x} x_i^\top) + \gamma^2 (\Sigma_{xx} + \bar{x} \bar{x}^\top) - (1-\bar{\theta})^2 \bar{x} \bar{x}^\top \\
 &= \sigma^2 (x_i - \bar{x})(x_i - \bar{x})^\top + \gamma^2 \Sigma_{xx},
 \end{aligned} \tag{29}$$

where we used the independence between θ and j in the first equality; for the second, the facts that $\mathbb{E}_\theta(\theta - \bar{\theta})^2 = \sigma^2$, $\mathbb{E}_\theta[(1 - \theta)^2] = \sigma^2 + (1 - \bar{\theta})^2 = \gamma^2$, $\mathbb{E}_\theta[(\theta - \bar{\theta})(1 - \theta)] = \bar{\theta}^2 - \mathbb{E}_\theta[\theta^2] = -\sigma^2$, and $\mathbb{E}_\theta[1 - \theta] = 1 - \bar{\theta}$; for the third, we reorganized the terms and used the equality $\mathbb{E}_j x_j x_j^\top = \Sigma_{xx} + \bar{x}\bar{x}^\top$ by definition of the empirical covariance matrix Σ_{xx} ; the last equality is obtained by reorganizing the terms and using the definition of γ^2 . In order to write this covariance matrix in terms of modified inputs, we notice that by definition (5) we have $x_i - \bar{x} = (\tilde{x}_i - \bar{x})/\bar{\theta}$ and $\mathbb{E}_j \tilde{x}_j = \mathbb{E}_j x_j = \bar{x}$, which implies that the empirical covariance matrix of the modified inputs is $\Sigma_{\tilde{x}\tilde{x}} = \bar{\theta}^2 \Sigma_{xx}$. Combining these equalities with (29) gives the first equality in (9). The two other equalities can be proved exactly the same way. \blacksquare

A.3 Proof of Theorem 3

Proof Given a modified input/output pair $(\tilde{x}, \tilde{y}) \in \mathcal{X} \times \mathcal{Y}$ and a function $f \in \mathcal{H}$, the second-order Taylor approximation of the loss $G(\tilde{x}, \tilde{y}) = \ell(\tilde{y}, f(\tilde{x}))$ is, for any $(\delta, \varepsilon) \in \mathcal{X} \times \mathcal{Y}$:

$$\begin{aligned} G_Q(\tilde{x} + \delta, \tilde{y} + \varepsilon) &= G(\tilde{x}, \tilde{y}) + \nabla_{\tilde{x}} G(\tilde{x}, \tilde{y})\delta + \nabla_{\tilde{y}} G(\tilde{x}, \tilde{y})\varepsilon \\ &\quad + \frac{1}{2}\delta^\top \nabla_{\tilde{x}\tilde{x}}^2 G(\tilde{x}, \tilde{y})\delta + \frac{1}{2}\varepsilon^\top \nabla_{\tilde{y}\tilde{y}}^2 G(\tilde{x}, \tilde{y})\varepsilon + \varepsilon^\top \nabla_{\tilde{y}\tilde{x}}^2 G(\tilde{x}, \tilde{y})\delta. \end{aligned} \quad (30)$$

Using this quadratic approximation at each training point $i \in [n]$ in (7), and using the fact that $\mathbb{E}_{\theta,j}\delta_i = \mathbb{E}_{\theta,j}\varepsilon_i = 0$, we get

$$\begin{aligned} \mathbb{E}_{\theta,j} G_Q(\tilde{x}_i + \delta_i, \tilde{y}_i + \varepsilon_i) &= G(\tilde{x}_i, \tilde{y}_i) + \frac{1}{2} \left\langle \mathbb{E}_{\theta,j} \delta_i \delta_i^\top, \nabla_{\tilde{x}\tilde{x}}^2 G(\tilde{x}_i, \tilde{y}_i) \right\rangle \\ &\quad + \frac{1}{2} \left\langle \mathbb{E}_{\theta,j} \varepsilon_i \varepsilon_i^\top, \nabla_{\tilde{y}\tilde{y}}^2 G(\tilde{x}_i, \tilde{y}_i) \right\rangle + \left\langle \mathbb{E}_{\theta,j} \varepsilon_i \delta_i^\top, \nabla_{\tilde{y}\tilde{x}}^2 G(\tilde{x}_i, \tilde{y}_i) \right\rangle, \end{aligned} \quad (31)$$

which we can rewrite as follows by expressing the derivatives of $G(x, y) = \ell(y, f(x))$ in terms of derivatives of $\ell(y, u)$ and $f(x)$:

$$\begin{aligned} \mathbb{E}_{\theta,j} \ell_Q(\tilde{y}_i + \varepsilon_i, f(\tilde{x}_i + \delta_i)) &= \ell(\tilde{y}_i, f(\tilde{x}_i)) \\ &\quad + \frac{1}{2} \left\langle \mathbb{E}_{\theta,j} \delta_i \delta_i^\top, \nabla f(\tilde{x}_i)^\top \nabla_{uu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) + \nabla_u \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla^2 f(\tilde{x}_i) \right\rangle \\ &\quad + \frac{1}{2} \left\langle \mathbb{E}_{\theta,j} \varepsilon_i \varepsilon_i^\top, \nabla_{yy}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \right\rangle + \left\langle \mathbb{E}_{\theta,j} \varepsilon_i \delta_i^\top, \nabla_{yu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) \right\rangle. \end{aligned} \quad (32)$$

Replacing the expectations in this equation by their values given by Lemma 2 gives:

$$\begin{aligned} \mathbb{E}_{\theta,j} \ell_Q(\tilde{y}_i + \varepsilon_i, f(\tilde{x}_i + \delta_i)) &= \ell(\tilde{y}_i, f(\tilde{x}_i)) \\ &\quad + \frac{1}{2} \left\langle \Sigma_{\tilde{x}\tilde{x}}^{(i)}, \nabla f(\tilde{x}_i)^\top \nabla_{uu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) \right\rangle \\ &\quad + \frac{1}{2} \left\langle \Sigma_{\tilde{y}\tilde{y}}^{(i)}, \nabla_u \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla^2 f(\tilde{x}_i) \right\rangle \\ &\quad + \frac{1}{2} \left\langle \Sigma_{\tilde{y}\tilde{y}}^{(i)}, \nabla_{yy}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \right\rangle \\ &\quad + \left\langle \Sigma_{\tilde{y}\tilde{x}}^{(i)}, \nabla_{yu}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) \right\rangle. \end{aligned} \quad (33)$$

We now use the following fact, true for any square symmetric and invertible matrices A and C and rectangular matrices B and Y (such that the matrix multiplications below make sense):

$$\left\langle A, Y^\top C Y \right\rangle - 2 \left\langle B, Y \right\rangle = \left\langle A, (Y - Z)^\top C (Y - Z) \right\rangle - \left\langle A^{-1}, B^\top C^{-1} B \right\rangle, \quad (34)$$

where $Z = C^{-1} B A^{-1}$, to combine the second and fifth terms together. Indeed, the fifth term (33) can be rewritten as

$$\left\langle \Sigma_{\tilde{y}\tilde{x}}^{(i)}, \nabla_{y u}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) \right\rangle = \left\langle \nabla_{u y}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \Sigma_{\tilde{y}\tilde{x}}^{(i)}, \nabla f(\tilde{x}_i) \right\rangle,$$

so plugging into (34) the following matrices:

$$\begin{cases} A &= \Sigma_{\tilde{x}\tilde{x}}^{(i)}, \\ B &= -\nabla_{u y}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \Sigma_{\tilde{y}\tilde{x}}^{(i)}, \\ C &= \nabla_{u u}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)), \\ Y &= \nabla f(\tilde{x}_i), \end{cases}$$

gives

$$\begin{aligned} & \frac{1}{2} \left\langle \Sigma_{\tilde{x}\tilde{x}}^{(i)}, \nabla f(\tilde{x}_i)^\top \nabla_{u u}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) \right\rangle + \left\langle \Sigma_{\tilde{y}\tilde{x}}^{(i)}, \nabla_{y u}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla f(\tilde{x}_i) \right\rangle \\ &= \frac{1}{2} \left\langle \Sigma_{\tilde{x}\tilde{x}}^{(i)}, \left(\nabla f(\tilde{x}_i) - J^{(i)} \right)^\top \nabla_{u u}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \left(\nabla f(\tilde{x}_i) - J^{(i)} \right) \right\rangle \\ & \quad - \frac{1}{2} \left\langle \Sigma_{\tilde{y}\tilde{x}}^{(i)} \left(\Sigma_{\tilde{x}\tilde{x}}^{(i)} \right)^{-1} \Sigma_{\tilde{x}\tilde{y}}^{(i)}, \nabla_{y u}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \nabla_{u u}^2 \ell(\tilde{y}_i, f(\tilde{x}_i))^{-1} \nabla_{u y}^2 \ell(\tilde{y}_i, f(\tilde{x}_i)) \right\rangle \end{aligned} \quad (35)$$

where $J^{(i)}$ is defined in (13). Theorem 3 then follows by merging the second and fifth terms in (33) using (35), and summing over i . \blacksquare

A.4 Proof of Corollary 4

Proof For the definition (3) of the cross-entropy loss $\ell^{\text{CE}}(y, u)$, we easily get:

$$\begin{aligned} \nabla_y \ell^{\text{CE}}(y, u) &= -u^\top, \\ \nabla_u \ell^{\text{CE}}(y, u) &= (\mathcal{S}(u) - y)^\top, \\ \nabla_{y y}^2 \ell^{\text{CE}}(y, u) &= \mathbf{0}_c, \\ \nabla_{y u}^2 \ell^{\text{CE}}(y, u) &= -\mathbf{I}_c, \\ \nabla_{u u}^2 \ell^{\text{CE}}(y, u) &= H(u). \end{aligned}$$

Plugging these results back in the four regularization terms in Theorem 3 we conclude the proof. \blacksquare

A.5 Proof of Corollary 5

Proof For the definition (15) of the logistic regression loss $\ell^{\text{LR}}(y, u)$, we easily get:

$$\begin{aligned}\nabla_y \ell^{\text{LR}}(y, u) &= -u, \\ \nabla_u \ell^{\text{LR}}(y, u) &= s(u) - y, \\ \nabla_{yy}^2 \ell^{\text{LR}}(y, u) &= 0, \\ \nabla_{yu}^2 \ell^{\text{LR}}(y, u) &= -1, \\ \nabla_{uu}^2 \ell^{\text{LR}}(y, u) &= v(u).\end{aligned}$$

Plugging these results back in the four regularization terms in Theorem 3 we conclude the proof. \blacksquare

A.6 Proof of Corollary 6

Proof For the definition (3) of the squared error loss $\ell^{\text{SE}}(y, u)$, we easily get:

$$\begin{aligned}\nabla_y \ell^{\text{SE}}(y, u) &= (y - u)^\top, \\ \nabla_u \ell^{\text{SE}}(y, u) &= (u - y)^\top, \\ \nabla_{yy}^2 \ell^{\text{SE}}(y, u) &= \nabla_{uu}^2 \ell^{\text{SE}}(y, u) = \mathbf{I}_c, \\ \nabla_{yu}^2 \ell^{\text{SE}}(y, u) &= -\mathbf{I}_c.\end{aligned}$$

Plugging these results back in the 4 regularization terms in Theorem 3 proves (17).

When f is a linear function with intercept of the form $f_{W,b}(x) = Wx + b$, then we first note that $\ell^{\text{SE}}(y, f_{W,b}(x))$ is a quadratic function of (x, y) , so the second-order Taylor approximation (11) is exact in that case: $\ell_Q^{\text{SE}(i)}(y, f_{W,b}(x)) = \ell^{\text{SE}}(y, f_{W,b}(x))$ for any $i \in [n]$ and $(x, y) \in \mathcal{X} \times \mathcal{Y}$, and consequently:

$$\forall (W, b) \in \mathbb{R}^{c \times d} \times \mathbb{R}^c, \quad \mathcal{E}_Q^{\text{Mixup}}(f_{W,b}) = \mathcal{E}^{\text{Mixup}}(f_{W,b}).$$

Applying (17) to the case of a linear function $f_{W,b}$ gives us immediately $R_2^{\text{SE}}(f_{W,b}) = 0$, because $\nabla^2 f_{W,b} = 0$. For the first regularization term, we compute

$$\begin{aligned}
 R_1^{\text{SE}}(f_{W,b}) &= \frac{1}{2n} \sum_{i=1}^n \|\nabla f(\tilde{x}_i) - J^{(i)}\|_{\Sigma_{\tilde{x}\tilde{x}}^{(i)}}^2 \\
 &= \frac{1}{2n} \sum_{i=1}^n \|W - J^{(i)}\|_{\Sigma_{\tilde{x}\tilde{x}}^{(i)}}^2 \\
 &= \frac{1}{2n} \sum_{i=1}^n \left(\langle W^\top W, \Sigma_{\tilde{x}\tilde{x}}^{(i)} \rangle - 2 \langle W, \Sigma_{\tilde{y}\tilde{x}}^{(i)} \rangle \right) + C \\
 &= \frac{1}{2n} \sum_{i=1}^n \left[\frac{\gamma^2}{n\bar{\theta}^2} \sum_{j=1}^n \|W(\tilde{x}_j - \bar{x}) - (\tilde{y}_j - \bar{y})\|^2 + \frac{\sigma^2}{\theta^2} \|W(\tilde{x}_i - \bar{x}) - (\tilde{y}_i - \bar{y})\|^2 \right] + C \\
 &= \frac{\sigma^2 + \gamma^2}{2n\bar{\theta}^2} \sum_{i=1}^n \|W(\tilde{x}_i - \bar{x}) - (\tilde{y}_i - \bar{y})\|^2 + C \\
 &= \frac{2\sigma^2 + (1 - \bar{\theta})^2}{2n} \sum_{i=1}^n \|W(x_i - \bar{x}) - (y_i - \bar{y})\|^2 + C.
 \end{aligned} \tag{36}$$

As for the empirical risk term, we can also rewrite it as

$$\begin{aligned}
 \frac{1}{n} \sum_{i=1}^n \ell^{\text{SE}}(\tilde{y}_i, f_{W,b}(\tilde{x}_i)) &= \frac{1}{n} \sum_{i=1}^n \|W\tilde{x}_i + b - \tilde{y}_i\|^2 \\
 &= \frac{1}{n} \sum_{i=1}^n \|W(\tilde{x}_i - \bar{x}) - (\tilde{y}_i - \bar{y}) + (b - \bar{b})\|^2 \\
 &= \frac{1}{n} \sum_{i=1}^n \|W(\tilde{x}_i - \bar{x}) - (\tilde{y}_i - \bar{y})\|^2 + \|b - \bar{b}\|^2 \\
 &= \frac{\bar{\theta}^2}{n} \sum_{i=1}^n \|W(x_i - \bar{x}) - (y_i - \bar{y})\|^2 + \|b - \bar{b}\|^2
 \end{aligned} \tag{37}$$

Plugging (36) and (37) into (17) finally gives (19).

To see that the minimizer of (19) is the standard MOLS solution, we notice that the obvious solution for b is $b = \bar{b}$, which is the intercept of MOLS, while the solution for W should minimize the sum of squared errors over centered points, which is exactly what MOLS does. ■

A.7 Proof of Proposition 7

Proof We first start by deriving a dual formulation for (24). The derivation for (23) follows along the same lines, up to the replacement of \tilde{y}_i by y_i .

Introducing primal variables $u_i \in \mathbb{R}^c$ for $i \in [n]$ with equality constraints $Wx_i = u_i$ and the dual variables $\xi_i \in \mathbb{R}^c$, we obtain the following Lagrangian (see Boyd and Vandenberghe

(2004)):

$$\begin{aligned} \mathcal{L}(\{u_i, \xi_i\}_{i=1}^n, W) &= \frac{1}{n} \sum_{i=1}^n \{ \ell^{\text{CE}}(\tilde{y}_i, u_i) + \xi_i^\top (Wx_i - u_i) \} \\ &= \frac{1}{n} \sum_{i=1}^n \left\{ \log \left(\sum_{j=1}^c e^{(u_i)_j} \right) - (\xi_i + \tilde{y}_i)^\top u_i \right\} + \left\langle W, \frac{1}{n} \sum_{i=1}^n \xi_i x_i^\top \right\rangle. \end{aligned}$$

We recall that the entropy is concave and is the Fenchel conjugate, up to a sign flip, of the log-sum-exp function

$$\mathcal{Z}(p) = - \max_{t \in \mathbb{R}^c} \left\{ p^\top t - \log \left(\sum_{j=1}^c e^{t_j} \right) \right\}, \quad (38)$$

as for instance detailed in Example 5.5 of (Boyd and Vandenberghe, 2004). We therefore derive the dual function using (38)

$$\min_{\{u_i\}_{i=1}^n, W} \mathcal{L}(\{u_i, \xi_i\}_{i=1}^n, W) = \begin{cases} \frac{1}{n} \sum_{i=1}^n \mathcal{Z}(\xi_i + \tilde{y}_i) & \text{if } \sum_{i=1}^n \xi_i x_i^\top = 0 \text{ and } \xi_i + \tilde{y}_i \in \Delta_c, \\ -\infty & \text{otherwise} \end{cases}$$

and the dual problem

$$\max_{\{\nu_i\}_{i=1}^n} \frac{1}{n} \sum_{i=1}^n \mathcal{Z}(\nu_i) \quad \text{subject to } \nu_i \in \Delta_c, \quad i \in [n], \quad \text{and} \quad \sum_{i=1}^n (\nu_i - \tilde{y}_i) x_i^\top = 0 \quad (39)$$

where we have made the change of variables $\nu_i = \xi_i + \tilde{y}_i$. The dual problem for (23) is identical to (39) up to the replacement of \tilde{y}_i by y_i .

Recalling the definitons in (25) and exploiting the first-order optimality conditions of (23) and (24), we have

$$\sum_{i=1}^n (p_i - y_i) x_i^\top = 0 \quad \text{and} \quad \sum_{i=1}^n (\tilde{p}_i - \tilde{y}_i) x_i^\top = 0,$$

so that $\{\tilde{p}_i\}_{i=1}^n$ is feasible for the dual problem (39). Since strong duality applies (Boyd and Vandenberghe, 2004), it also holds that $\{\tilde{p}_i\}_{i=1}^n$ maximize (39).

Let us consider $\{q_i\}_{i=1}^n$ defined for $i \in [n]$ by

$$q_i = \bar{\theta} p_i + (1 - \bar{\theta}) \bar{y}.$$

We can easily observe that $q_i \in \Delta_c$ as convex combination of $p_i, \bar{y} \in \Delta_c$ and that

$$\sum_{i=1}^n (q_i - \tilde{y}_i) x_i^\top = \bar{\theta} \sum_{i=1}^n (p_i - y_i) x_i^\top + (1 - \bar{\theta}) \sum_{i=1}^n (\bar{y} - \tilde{y}_i) x_i^\top = 0.$$

This implies that $\{q_i\}_{i=1}^n$ is feasible for (39) and we have

$$\frac{1}{n} \sum_{i=1}^n \mathcal{Z}(q_i) \leq \frac{1}{n} \sum_{i=1}^n \mathcal{Z}(\tilde{p}_i).$$

We get the advertised result by using the concavity of \mathcal{Z} so that $\bar{\theta} \mathcal{Z}(p_i) + (1 - \bar{\theta}) \mathcal{Z}(\bar{y}) \leq \mathcal{Z}(q_i)$.
■

Appendix B. Experiments

B.1 CIFAR-10 and CIFAR-100

To learn on CIFAR-10, CIFAR-100 we use ResNet-34 and LeNet. For both architectures the optimizer used for training is SGD with momentum 0.9 for 200 epochs with mini-batch size 128, weight-decay 5×10^{-4} . For ResNet-34 the learning rate is 0.1 reduced by a factor 10 at epoch 60, 120, 160. For LeNet the learning rate is 0.01 reduced by a factor 10 at epoch 100.

B.2 ImageNet

To learn on the ImageNet we use ResNet-50. The optimizer used for training is SGD with Nesterov and momentum 0.9 for 200 epochs with mini-batch size 4096 (32 numbers of cores \times 128 per core batch-size), weight-decay 5×10^{-5} . The learning rate is 1.6 reduced by a factor 10 at epoch 66, 133, 177.

B.3 Two Moons with Random Features

We report in Figure 8 and Figure 9 some more results on a synthetic binary classification problem (noisy two-moon problem), where we train a logistic regression model with random Fourier features (Rahimi and Recht, 2008). This allowed us to get rid of convergence issues due to the convexity of the problem, but still work with nonlinear models of the input points. For each experimental result we report mean and 95% confidence interval using 30 repetitions over 30 different instances of the data.

What we notice from these results is similar behaviors as reported for CIFAR-10, CIFAR-100, ImageNet in the main paper.




Fig. 8: From left to right: train loss, train and test accuracy during optimization of a logistic regression model trained on the noisy two-moon problem with Mixup, Mixup (rescaled), approximate Mixup and ERM risks.

Data generation. To generate the data we use the `sklearn.datasets.make_moons` function from the `scikit-learn` library. We create $n = 300$ points with `noise = 0.01`, and split them in 50% for train and 50% for test. We then randomly flip 20% of the training labels to make the learning task more difficult. We repeat this pipeline 30 times for 30 different random seeds.

Function space. Let $M = 1000$, $w \in \mathbb{R}^M$ and $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^M$ be the feature map defined as $\phi(x) = \frac{1}{\sqrt{M}} \cos(Sx + B)$, where $\cos : \mathbb{R}^M \rightarrow \mathbb{R}^M$ is the element-wise *cosine* function,




Fig. 9: Histograms of confidence of predictions on test points, for models trained with different techniques.

$S \in \mathbb{R}^{M \times d}$ is the random matrix s.t. $S_{i,j} \sim \mathcal{N}(0, \sigma^2), \forall i \in [M], j \in [d]$ with $\sigma = 10$, and $B \in \mathbb{R}^M$ s.t. $B_i \sim \text{Unif}(0, 2\pi), \forall i \in [M]$. The space of candidate solutions \mathcal{H} we consider is the class of functions of the form $f(x) = w^\top \phi(x)$.

Optimization. To minimize any functional we use stochastic gradient descent with mini-batching, with mini-batch size $b = 50$ and step-size $\gamma = 5$.

Mixup hyperparameter. We consider the Beta distribution in Mixup and its approximation to be $\text{Beta}(\alpha, \alpha)$ with $\alpha = 1$.

References

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 7411–7422. Curran Associates, Inc., 2019.
- Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 233–242. JMLR.org, 2017.
- Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. ObjectNet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 9453–9463. Curran Associates, Inc., 2019.
- Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 6241–6250. Curran Associates Inc., 2017.

- David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. MixMatch: A holistic approach to semi-supervised learning. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Curran Associates Inc., 2019.
- Chris M Bishop. Training with noise is equivalent to Tikhonov regularization. *Neural computation*, 7(1):108–116, 1995.
- Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20, pages 161–168. Curran Associates, Inc., 2008.
- Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- Luigi Carratino, Moustapha Cissé, Rodolphe Jenatton, and Jean-Philippe Vert. On mixup regularization. Technical Report 2006.06049, arXiv, 2020.
- Moustapha Cissé, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. Parseval networks: Improving robustness to adversarial examples. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 854–863. JMLR.org, 2017.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059. PMLR, 2016.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of Wasserstein GANs. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 5767–5777. Curran Associates, Inc., 2017.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1321–1330. JMLR.org, 2017.
- Hongyu Guo. Nonlinear Mixup: Out-of-manifold data augmentation for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4044–4051, 2020.
- Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3714–3722, 2019.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Matthias Hein and Maksym Andriushchenko. Formal guarantees on the robustness of a classifier against adversarial manipulation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 2263–2273. Curran Associates Inc., 2017.
- Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 15262–15271. Computer Vision Foundation / IEEE, 2021.
- Geoffrey E Hinton. Learning translation invariant recognition in a massively parallel networks. In *International Conference on Parallel Architectures and Languages Europe*, pages 1–13. Springer, 1987.
- Beyrem Khalfaoui, Joseph Boyd, and Jean-Philippe Vert. Asni: Adaptive structured noise injection for shallow and deep neural networks. *arXiv preprint arXiv:1909.09819*, 2019.
- Anders Krogh and John A Hertz. A simple weight decay can improve generalization. In J. Moody, S. Hanson, and R.P. Lippmann, editors, *Advances in neural information processing systems*, volume 4, pages 950–957. Morgan-Kaufmann, 1991.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 4694–4703. Curran Associates, Inc., 2019.
- Behnam Neyshabur. Implicit regularization in deep learning. *arXiv preprint arXiv:1709.01953*, 2017.
- Behnam Neyshabur, Russ R Salakhutdinov, and Nati Srebro. Path-SGD: Path-normalized optimization in deep neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, pages 2422–2430, 2015.

- Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in neural information processing systems*, volume 20, pages 1177–1184. Curran Associates, Inc., 2008.
- Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do ImageNet classifiers generalize to ImageNet? In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 5389–5400. PMLR, 2019.
- Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in neural information processing systems*, volume 29, pages 901–909. Curran Associates, Inc., 2016.
- Hanie Sedghi, Vineet Gupta, and Philip M Long. The singular values of convolutional layers. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- Vaishal Shankar, Achal Dave, Rebecca Roelofs, Deva Ramanan, Benjamin Recht, and Ludwig Schmidt. Do image classifiers generalize across time? In *ICCV*, pages 9641–9649. IEEE, 2021.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- Sunil Thulasidasan, Gopinath Chennupati, Jeff A Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Yuji Tokozume, Yoshitaka Ushiku, and Tatsuya Harada. Between-class learning for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5486–5494, 2018.
- Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of*

- Machine Learning Research*, pages 6438–6447, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v97/verma19a.html>.
- Stefan Wager, Sida Wang, and Percy S Liang. Dropout training as adaptive regularization. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in neural information processing systems*, volume 26, pages 351–359. Curran Associates, Inc., 2013.
- Colin Wei, Sham Kakade, and Tengyu Ma. The implicit and explicit regularization effects of dropout. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 10181–10192. PMLR, 13–18 Jul 2020.
- Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6023–6032, 2019.
- Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.
- Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Zou. How does mixup help with robustness and generalization? In *International Conference on Learning Representations*. OpenReview.net, 2021.