CODED-SMOOTHING: CODING THEORY HELPS GENERALIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

We introduce the coded-smoothing module, which can be seamlessly integrated into standard training pipelines, both supervised and unsupervised, to regularize learning and improve generalization with minimal computational overhead. In addition, it can be incorporated into the inference pipeline to randomize the model and enhance robustness against adversarial perturbations. The design of coded-smoothing is inspired by *general coded computing*, a paradigm originally developed to mitigate straggler and adversarial failures in distributed computing by processing linear combinations of the data rather than the raw inputs. Building on this principle, we adapt coded computing to machine learning by designing an efficient and effective regularization mechanism that encourages smoother representations and more generalizable solutions. Extensive experiments on both supervised and unsupervised tasks demonstrate that coded-smoothing consistently improves generalization and achieves state-of-the-art robustness against gradient-based adversarial attacks.

1 INTRODUCTION

Reliable prediction remains a central challenge in modern machine learning. Although deep neural networks have achieved remarkable success across computer vision, natural language processing, and reinforcement learning, their generalization beyond training data remains imperfect, and their reliability under adversarial perturbations is still limited (Szegedy et al., 2013; Goodfellow et al., 2014; Wen et al., 2020; Liu et al., 2020). This vulnerability is largely a consequence of overparameterization combined with limited training data, which makes models prone to overfitting, memorization, and brittle behavior when faced with unseen or corrupted inputs. Regularization techniques therefore play a key role in improving reliability: by guiding models toward simpler and smoother solutions, they reduce generalization error while simultaneously enhancing robustness to adversarial attacks.

Classical regularization strategies such as weight decay (Krogh & Hertz, 1991), dropout (Srivastava et al., 2014), and batch normalization (Ioffe & Szegedy, 2015) have long been established. More recently, data-centric approaches such as label smoothing (Szegedy et al., 2016), mixup and its variations (Zhang et al., 2017; Verma et al., 2019; Berthelot et al., 2019; Yun et al., 2019; Yao et al., 2022; Pinto et al., 2022; Bouniot et al., 2023) have become widely adopted for supervised learning. Nonetheless, data-centric approaches that are broadly applicable to both supervised and unsupervised models, and that simultaneously enhance generalization and adversarial robustness, remain insufficiently investigated.

In this paper, we take a step toward closing this gap, and introduce a new powerful regularization method, using *coded-smoothing module*, which applies seamlessly in both supervised and unsupervised settings. Our approach draws inspiration from an unexpected source: *coded computing*. Originally developed for distributed computing systems to mitigate the effects of straggler servers (Yu et al., 2017; 2020; Dutta et al., 2020; Jahani-Nezhad & Maddah-Ali, 2022; Moradi et al., 2024; Moradi & Maddah-Ali, 2025) and adversarial servers (Yu et al., 2019; Soleymani et al., 2022; Moradi et al., 2025), coded computing injects redundancy into the computational process. In this approach, instead of directly processing raw data and computing the designated results, the servers operate on carefully designed weighted linear combinations of the data, referred to as coded inputs. The number of coded inputs exceeds that of the original raw inputs. This coded redundancy enables the recovery of the original computation through a decoding procedure, even in the presence of missing results from stragglers or corrupted results from adversarial servers. In particular, in *general coded*

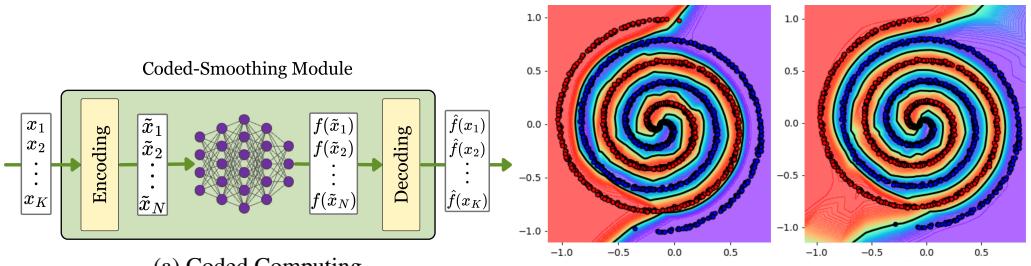

(a) Coded Computing

(b) Classification Boundaries, trained by Mixup (left) versus by coded smoothing module

Figure 1: 1.a. In a coded computing module, instead of directly computing $f(x_1), \ldots, f(x_K)$, the system computes $f(\tilde{x}_1), \ldots, f(\tilde{x}_N)$, where $N > K$ and each coded input $\tilde{x}_i$ is a unique weighted linear combination of the originals. The desired outputs are then reconstructed via a decoding procedure, yielding approximations $\hat{f}(x_1) \approx f(x_1), \ \ldots, \ \hat{f}(x_K) \approx f(x_K)$.

computing (Moradi et al., 2024), the smoother the function representing the computation task, the more accurate the approximated result.

The coded-smoothing module has impactful structure. Given a batch of K input samples, it first generates a new batch of N *coded samples* through an encoding process, where each coded sample is formed as a combination of all inputs in the batch. The network is then evaluated on these coded samples, and a subsequent decoding step reconstructs estimates of the network outputs on the original inputs (see Figure 1a). Importantly, enforcing closeness between these decoded estimates and the true outputs induces local smoothness in the learned network and effectively reduces its complexity. To achieve this, during training we augment the objective with an auxiliary penalty term that encourages the decoded outputs to remain close to their true counterparts (see Fig. 2), thereby guiding the model toward smoother and more generalizable solutions (see Fig. 1b).

Beyond training, using coded-smoothing module offers a striking additional benefit at inference time. Since the coded-smoothing module works independently from the order of data in the input batch, we can inject randomness by applying a random shuffle before encoding and restoring the order after decoding. This simple yet powerful mechanism disrupts gradient-based adversarial attacks such as FGSM (Goodfellow et al., 2014) and PGD (Madry et al., 2017), which rely on precise gradient information to craft adversarial examples. As a result, the model attains substantially improved robustness against adversarial perturbations. Notably, this method imposes negligible computational overhead, making it both effective and practical for real-world deployment.

Our experiments show that the coded-smoothing module consistently improves generalization across a wide range of architectures and benchmarks in both supervised and unsupervised settings. Moreover, coded-smoothing provides substantial gains in adversarial robustness. Compared to mixup (Zhang et al., 2017), it achieves an 8.8% higher accuracy under the FGSM attack ($\epsilon = 8/255$) (Goodfellow et al., 2014), a 31.8% improvement under PGD with 10 steps, and a 37% improvement under PGD with 100 steps (Madry et al., 2017).

Contributions. In summary, this work makes the following key contributions:

- We introduce the *coded-smoothing* module, a novel and computationally efficient regularization mechanism for neural networks inspired by principles of coded computing (Section 3 and Appendix C).
- We provide a theoretical characterization showing that coded-smoothing enforces higher-order local smoothness, thereby acting as a powerful regularizer (Section 4.1).
- We propose a randomized coded inference procedure based on the coded-smoothing module that substantially improves adversarial robustness without requiring adversarial training (Section 5).
- We conduct extensive experiments demonstrating that coded-smoothing consistently enhances both generalization and robustness across datasets and architectures, while incurring minimal computational overhead (Section 6).

Algorithm 1: Pseudo-code for coded-smoothing module

Input: Input tensor X of shape $(K, \cdot)$, where K is the batch size; Computation function f (e.g., a neural network model).

Output: Estimated output tensor $\hat{f}(X)$ of shape $(K, \cdot)$.

```
class CodedSmoothing(nn.Module):
    def __init__(self, K, N):
        super().__init__()
        self.alpha = generate_encoding_points(K)
        self.beta = generate_decoding_points(N)
        self.enc = Spline(knots=alpha)
        self.dec = Spline(knots=beta)

    def forward(self, X, f):
        self.enc.fit(self.alpha, X)
        x_coded = self.enc.predict(self.beta)
        f_coded = f(x_coded)
        self.dec.fit(self.beta, f_coded)
        f_hat = self.dec.predict(self.alpha)
        return f_hat
```

2 Empirical Risk Minimization (ERM)

In the supervised learning setting, let $\mathcal{D} := (x_i, y_i)i = 1^n$ denote a training dataset of size n, sampled from a distribution $\mathbb{P}$, where $x_i \in \mathcal{X}$ is the input and $y_i \in \mathcal{Y}$ is the corresponding label. Here, $\mathcal{X}$ and $\mathcal{Y}$ represent the input and output spaces, respectively, and $\theta \in \Theta$ denotes the parameter space. Given a loss function $\ell(\cdot, \cdot)$, ERM aims to learn a mapping $f\theta : \mathcal{X} \to \mathcal{Y}$ by minimizing the expected loss with respect to the empirical distribution $\mathbb{P}e(x, y) := \frac{1}{n} \sum_{i=1}^{n} \delta(x = x_i, y = y_i)$.

$$\theta^* = \arg\min_{\theta} \mathbb{E}_{\mathbb{P}_e(x,y)}[\ell(f_\theta(x), y)] = \int \ell(f_\theta(x), y) \, \mathrm{d}\mathbb{P}_e(x, y) = \frac{1}{n} \sum_{i=1}^{n} \ell(f_\theta(x_i), y_i) \,. \quad (1)$$

The goal is for the learned model to generalize well to unseen samples drawn from a test distribution $\mathbb{P}_t$, in both in-distribution ($\mathbb{P}_t = \mathbb{P}$) and out-of-distribution ($\mathbb{P}_t \neq \mathbb{P}$) settings.

3 Coded-SmoothingModule

Building on the general coded computing Moradi et al. (2024), we propose the *coded-smoothing* module as a regularization technique to model smoothness. We first describe the architecture of the proposed module, and then explain how coded-smoothing integrates into both the training and inference pipelines. This integration leads to improved generalization as well as enhanced adversarial robustness of the model.

3.1 Architecture

The coded-smoothing module consists of three components: an encoder function $u_{\text{enc}} : \mathbb{R} \to \mathcal{U}$, a computation function $f : \mathcal{U} \to \mathcal{V}$, and a decoder function $u_{\text{dec}} : \mathbb{R} \to \mathcal{V}$. Here, $\mathcal{U}$ and $\mathcal{V}$ are the input and output domains of the function $f(\cdot)$, and f may represent a machine learning model or a set of consecutive layers in a deep neural network. Given a batch of input data $\{x_1, \ldots, x_K\}$, the module produces an estimate of the computation function on these inputs, denoted by $\{\hat{f}(x_i)\}_{i=1}^{K}$.

The end-to-end process proceeds as follows:

(1) **Encoding:** the encoder function u_{enc} is fitted to the set of points $\{(\alpha_i, x_i)\}_{i=1}^{K}$, where $\alpha_1 < \alpha_2 < \cdots < \alpha_K \in [-1, 1]$ are referred to as *encoding points*. Therefore,

$$u_{\text{enc}}(\alpha_i) = x_i, \quad , \forall i \in [K]. \quad (2)$$

Then, N *coded samples* are generated by evaluating the encoder at another fixed set $\{\beta_j\}_{j=1}^N$ with $\beta_1 < \beta_2 < \cdots < \beta_N \in [-1, 1]$, called *decoding points* $\tilde{x}_j = u_{\text{enc}}(\beta_j)$, for $j \in [N]$. We note that each coded sample $\tilde{x}_j$ is a combination of the original input dataset $\{x_i\}_{i=1}^K$.

(2) **Computation:** In this step, $f(\tilde{x}_j)$, for $j = 1, \ldots, N$, are computed.

(3) **Decoding:** In this stage, first, decoder function u_{dec} is fitted to the set of points $\{(\beta_j, f(\tilde{x}_j))\}_{j=1\mathcal{F}}^N$, therefore,

$$u_{\text{dec}}(\beta_j) = f(\tilde{x}_j) = f(u_{\text{enc}}(\beta_j)), \quad \forall j \in [N], \tag{3}$$

where the second equation follows from (2). If the decoder $u_{\text{dec}}(\cdot)$ generalizes well, then $u_{\text{dec}}(z) \approx f(u_{\text{enc}}(z))$, for all $z \in [-1, 1]$. In particular, at the encoding points, we have,

$$u_{\text{dec}}(\alpha_i) \approx f(u_{\text{enc}}(\alpha_i)) = f(x_i), \tag{4}$$

where the first approximation relies on the generalization ability of u_{dec}, and the second equation follows from (2). Thus, $u_{\text{dec}}(\alpha_i)$ approximates $f(x_i)$. We define $\hat{f}(x_i) \triangleq u_{\text{dec}}(\alpha_i)$, for $i \in [K]$.

Algorithm 1 presents PyTorch-style pseudo-code for the coded-smoothing module. As suggested by (Moradi et al., 2024), we use *natural cubic splines* (cubic smoothing splines with smoothing parameter of zero) for both the encoder and decoder.

With a careful choice of encoding and decoding points, the following lemma provides a bound on the approximation error of the coded-smoothing module.

Lemma 1. *For a coded-smoothingmodule with N coded samples, we have:*

$$\frac{1}{K}\sum_{i=1}^{K}\left|\hat{f}(x_i) - f(x_i)\right|^2 \leq \frac{2C}{N^3}\left(\|u''_{enc} \cdot f'! \circ u_{enc}\|\, L^2(\Omega)^2 + \|u''_{enc} \cdot f'! \circ u_{enc}\|\, L^2(\Omega)^2\right), \tag{5}$$

for some constant C.

For proof, see Appendix B. Lemma 1 highlights an important property of coded-smoothing module: The larger the number of coded samples N or the smoother the function f, the smaller the mean squared estimation error.

Spline representation. Let $S_{\vec{t},\vec{y}}(\cdot)$ denote the smoothing spline fitted on $\{(t_i, y_i)\}_{i=1}^n$, where $t_i \in \mathbb{R}$, $y_i \in \mathbb{R}^d$, $\vec{y} := [y_1, \ldots, y_n]^T$, and $\vec{t} := [t_1, \ldots, t_n]^T$. It is well-known that $S_{\vec{t},\vec{y}}(z) = \sum_{i=1}^n y_i \phi(z, t_i)$, where $\phi(.,.)$ is the kernel of the second-order Sobolev space (i.e. functions with square-integrable derivatives up to order two). Thus, $S_{\vec{t},\vec{y}}(\cdot)$ is a linear function of $\vec{y}$ (Wahba, 1975). Therefore, for any evaluation set $\vec{v} := [v_1, \ldots, v_m]^T$, there exists a matrix $A_{\vec{t},\vec{v}} \in \mathbb{R}^{n\times m}$, which depends only on the knot set $\vec{t}$, the evaluation points $\vec{v}$, and the smoothing parameter λ (but not on $\vec{y}$), such that

$$[S_{\vec{t},\vec{y}}(v_1), \ldots, S_{\vec{t},\vec{y}}(v_m)]^T = A_{\vec{t},\vec{v}}^T \vec{y}. \tag{6}$$

Recall that in the coded-smoothing module, both the encoder and decoder are implemented using smoothing splines. Therefore, we have:

$$u_{\text{enc}}(z) = \sum_{i=1}^{K} x_i \phi(z, \alpha_i), \quad u_{\text{dec}}(z) = \sum_{j=1]}^{N} f(\tilde{x}_j)\phi(z, \beta_j). \tag{7}$$

4 Training Regularization using the Coded-Smoothing Module

We now describe how the coded-smoothing module can be integrated into the training pipeline of machine learning models to improve generalization. Since coded-smoothing does not require label information, it can be applied in both supervised (Section 6.1) and unsupervised (Section 6.2) settings.

Figure 2 illustrates the role of coded-smoothing during training. The computation function f may represent the entire network or a part of the network, which we refer to as the *target block*. The

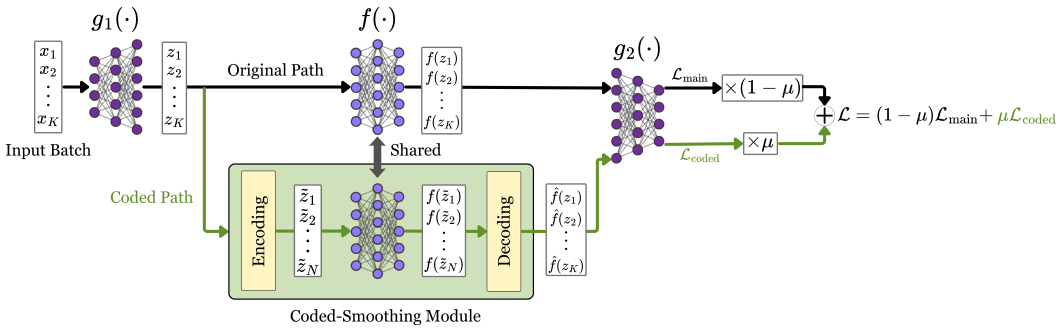

Figure 2: The proposed CODED-SMOOTHING as a regularization in training: the coded path includes a coded-smoothing module and runs in parallel to the original forward pass and contributes to the training objective.

integration of coded-smoothing introduces an additional *coded path* that runs in parallel to the original forward path.

Formally, consider training a deep neural network of the form $\texttt{net}(x) = g_1(f(g_2(x)))$, where $f(\cdot)$ is an intermediate target block. Suppose we apply coded-smoothing to $f(\cdot)$. After the input is passed through $g_1(\cdot)$, we branch it and follow two parallel paths: *the original path* and *the coded path* (see Fig. 2). In the coded path, there is a coded-smoothing module. The encoder generates a set of coded samples $\{\tilde{z}_j\}_{j=1}^N$, which form a new batch and are processed by the target block. The outputs of the target block on the coded samples are then passed through the decoder, producing estimated outputs $\{\hat{f}(z_i)\}_{i=1}^K$, which are approximately equal to $\{f(z_i)\}_{i=1}^K$. These estimated outputs are forwarded to the remainder of the network, denoted by $g_2(\cdot)$. During training, both paths contribute to the loss. Let $\mathcal{L}_{\text{main}}$ denote the loss from the original forward path, i.e., the standard training loss. Similarly, let $\mathcal{L}_{\text{coded}}$ denote the loss from the coded path, which has the same form as $\mathcal{L}_{\text{main}}$ but with the outputs of the original network replaced by those of the auxiliary coded path. The overall objective is then defined as

$$\mathcal{L} = (1-\mu)\mathcal{L}_{\text{main}} + \mu\,\mathcal{L}_{\text{coded}}, \tag{8}$$

where $\mu \in [0,1]$ is a weighting hyperparameter controlling the contribution of two paths. The parameters of the target block are shared between the original and coded paths, and the entire network is optimized with respect to the combined objective.

The second term in the loss function (8) acts as a regularizer, encouraging the coded path to match the predictive performance of the original path. In particular, it drives the coded-smoothing estimations of the target block toward their true outputs $\{f(z_i)\}_{i=1}^K$. Consequently, and in line with Lemma 1, the module implicitly enforces smoothness on the target block $f(\cdot)$. The effect of this regularization depends on the weighting coefficient μ: when $\mu \approx 1$, training is dominated by the coded path, whereas when $\mu \approx 0$, the process reduces to training only with the original loss.

4.1 CODED-SMOOTHINGIS A LOCAL HIGHER-ORDER SMOOTHER

In this subsection, we provide intuition for how the proposed approach encourages smoothness of the function. Recall from (4) that the accuracy of the approximation $\hat{f}(x_i) := u_{\text{dec}}(\alpha_i) \approx f(x_i)$ depends on the quality of the approximation $f(u_{\text{enc}}(z)) \approx u_{\text{dec}}(z)$. Moreover, from (7) we have $u_{\text{dec}}(z) = \sum_{j=1}^N f(\hat{x}_j)\,\phi(z,\beta_j)$. Hence, enhancing the approximation $\hat{f}(x_i) \approx f(x_i)$ is equivalent of improving the approximation

$$f(u_{\text{enc}}(z)) \approx \sum_{j\in[N]} f(\hat{x}_j)\,\phi(z,\beta_j). \tag{9}$$

The right-hand side is a weighted sum of some smooth functions, which implies that during training the regularized loss in (8) promotes smoothness of $f(u_{\text{enc}}(z))$, and consequently enforces smoothness in $f(\cdot)$ itself (see Fig. 3b).

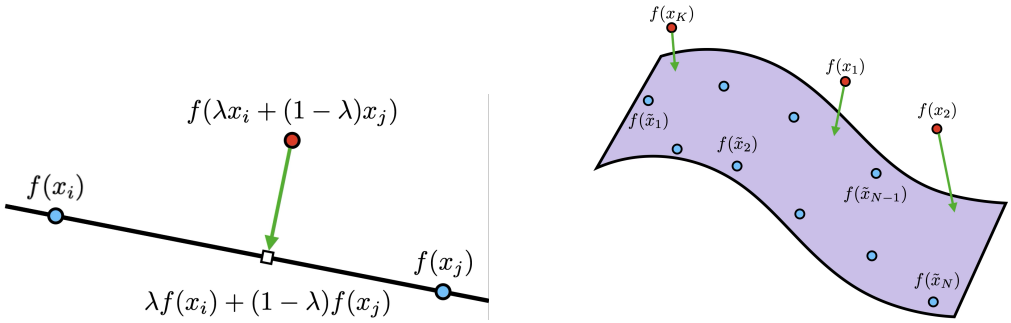

(a) Mixup Regularization (Zhang et al., 2017)

(b) Coded Smoothing

Figure 3: Coded Smoothing versus Mixup

To further clarify the concept, we next discuss the well-known Mixup method (Zhang et al., 2017) and highlight its connection to the proposed approach. In mixup, instead of empirical risk minimization, the model is trained by minimizing the expected loss with respect to a vicinal distribution $\mathbb{P}_v(x, y) := \frac{1}{n}\sum_{i=1}^{n} \delta(x = \bar{x}_i, y = \bar{y}_i)$, where $\bar{x}_i = \lambda x_i + (1-\lambda)x_j$ and $\bar{y}_i = \lambda y_i + (1-\lambda)y_j$ for $\lambda \sim \texttt{Beta}(\alpha, \alpha)$. As a result, the model is encouraged to align the prediction $f(\lambda x_i + (1-\lambda)x_j)$ with the target $\lambda y_i + (1-\lambda)y_j$ for $\lambda \in [0, 1]$. At the endpoints ($\lambda = 0, 1$), this also recovers the original labels, i.e. $f(x_i) \approx y_i$ and $f(x_j) \approx y_j$. Consequently, training implicitly enforces local linearity on the model which regularizes f to vary smoothly along the line segment connecting $f(x_i)$ and $f(x_j)$ (see Figure 3a):

$$f(\lambda x_i + (1-\lambda)x_j) \approx \lambda f(x_i) + (1-\lambda)f(x_j), \quad \lambda \in [0, 1]. \tag{10}$$

Comparing (9) and (10) reveals an intriguing connection between the two schemes. While the coded-smoothing module encourages $f(u_{\text{enc}}(z))$ to approximate a linear combination of smooth functions, mixup explicitly encourages $f(\cdot)$ to behave like a linear function. In other word, coded-smoothing module imposes a higher-order smoothness constraint on f, regularizing it beyond pairwise linearity. Although both approaches promote smoothness in $f(\cdot)$, coded smoothness admits a richer structure and may potentially lead to improved generalization (See Section 6 on experiment results).

5 Robust Inference using a Randomized coded-smoothing module

After training a model with the coded-smoothing module, both the coded path and the original path can be used during inference. Since the coded path generates a smooth approximation of the original outputs, its standalone generalization performance is dominated by that of the original path. However, the coded path possesses a useful property that can be exploited to substantially enhance adversarial robustness.

The key observation is that the proposed module performance does not depend on the order of input samples within a batch: the coded-smoothing module generates a good estimate for each input regardless of its position in the batch. During training, due to random shuffling across epochs, each sample x_i appears at different indices and the network aligns the estimation $\hat{f}(x_i)$ with its true output $f(x_i)$ independently of the sample's index.

Consequently, at inference time, one can introduce additional randomness by applying a random permutation π to the batch before feeding it into the encoder, and subsequently restoring the original order using π^{-1} before passing the outputs to the remainder of the network. We refer to this approach as *Randomized Coded Inference (RCI)*. Figure 4 illustrates this inference approach.

This strategy disrupts adversarial attacks, particularly gradient-based methods such as FGSM (Goodfellow et al., 2014) and PGD (Madry et al., 2017), which rely on precise gradients to craft adversarial examples. The core idea in these methods is to generate an adversarial sample by perturbing the input in the direction of the gradient of the loss with respect to that input. However, since π is chosen uniformly at random from all permutations, with high probability the permutation used by the network at inference differs from the one assumed by the adversary when generating the perturbations.

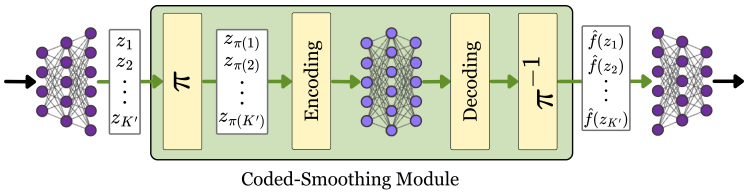

Figure 4: The proposed RANDOMIZED CODED INFERENCE: π represents a random permutation.

As a result, the network's robustness is significantly improved. Note that although coded-smoothing operates in batch mode at inference, the batch size need not match that used during training. In practice, the method is effective for batch sizes as small as $K' \geq 4$, since spline fitting requires at least three points, thereby offering flexibility for deployment (see Table 7 in Appendix G).

6 EXPERIMENTS

In this section, we evaluate the performance of the proposed coded-smoothing training method (using the coded-smoothing module) as well as the randomized coded inference approach, under various settings and across multiple evaluation metrics. We begin with the supervised scenario (Section 6.1), followed by the unsupervised setting (Section 6.2). We then demonstrate how coded-smoothing substantially enhances adversarial robustness during inference (Section 6.3). Finally, we assess its effectiveness under distribution shift, where the test distribution differs from the training distribution (Section 6.4). All experiments are conducted in PyTorch (Paszke et al., 2019) on a single machine equipped with an NVIDIA RTX 5090 GPU.

In all experiments, following Jahani-Nezhad & Maddah-Ali (2022); Moradi et al. (2024), we adopt the first-order Chebyshev points for encoding and the second-order Chebyshev points for decoding, i.e., $\alpha_i = \cos(\frac{(2i-1)\pi}{2K})$ and $\beta_j = \cos(\frac{(j-1)\pi}{N-1})$ for $i \in [K]$, $j \in [N]$. This choice is motivated by their superior empirical performance (Jahani-Nezhad & Maddah-Ali, 2022) and desirable theoretical properties in approximation theory (Phillips, 2003; Trefethen, 2019).

6.1 SUPERVISED

We begin by evaluating the effectiveness of the coded-smoothing module in the supervised setting. Specifically, we assess the generalization gains achieved by coded-smoothing compared to standard empirical risk minimization (ERM) and the widely used mixup regularization method (Zhang et al., 2017).

Datasets and architectures. To ensure a comprehensive evaluation across model families and dataset complexities, we conduct experiments on CIFAR-10, CIFAR-100 (Krizhevsky et al., 2009), and TinyImageNet (Le & Yang, 2015). For CIFAR-10, we use PreActResNet18 (He et al., 2016); for CIFAR-100, we employ WideResNet28-10 (Zagoruyko & Komodakis, 2016); and for TinyImageNet, we adopt ResNet50 (Goyal et al., 2017). These architectures are chosen to capture a range of model complexities while aligning with prior work.

In all supervised experiments, we empirically find that the best performance is achieved when the coded-smoothing module is applied to the full network. Table 1 reports the test performance across datasets and architectures. Each experiment is repeated over 5 independent train-validation splits with different random seeds. We report both the mean and standard deviation. As shown in Table 1, training with coded-smoothing consistently outperforms both mixup and ERM baselines across all benchmarks. Additional experimental details and hyperparameter selection are provided in Appendix D.1.

6.2 UNSUPERVISED

Next, we take one step further and evaluate the effectiveness of the CODED-SMOOTHING training in an unsupervised setting. Specifically, we incorporate coded-smoothing into the training of a WGAN-GP(Gulrajani et al., 2017) which is a variant of WGAN (Arjovsky et al., 2017).

Table 1: Comparisons of accuracies (%) on in-distribution test data.

	CIFAR-10 PARN18	CIFAR-100 WRN28-10	TinyImageNet RN50
ERM	93.8 ± 0.2	76.7 ± 0.3	62.9 ± 0.9
Mixup	95.6 ± 0.2	$\mathbf{80.2 \pm 0.3}$	65.4 ± 1.0
CODED-SMOOTHING (ours)	$\mathbf{95.8 \pm 0.1}$	$\mathbf{79.9 \pm 0.4}$	$\mathbf{67.1 \pm 0.5}$

Table 2: Comparison of FID and IS for generated images for CIFAR-10 and CelebA.

Method	CIFAR-10		CelebA
	IS	FID	FID
WGAN-GP	7.08 ± 0.07	26.93 ± 0.61	28.22 ± 0.17
WGAN-GP + CODED-SMOOTHING	$\mathbf{7.38 \pm 0.06}$	26.94 ± 0.89	$\mathbf{24.58 \pm 0.62}$

Prior work has shown that regularizing the discriminator can improve GAN training stability and performance (Zhang et al., 2017; Verma et al., 2019). However, because mixup and its variants rely on label information, they cannot be directly applied to the generator. Here we use CODED-SMOOTHING training method to regularize the *generator* of a WGAN. Specifically, we use coded-smoothing module with $N = K$ with batchsize $K = 64$ and $\mu = 0.5$. Further experimental details are provided in Appendix D.2. Table 2 reports the Fréchet Inception Distance (FID) (Heusel et al., 2017) and Inception Score (IS) (Salimans et al., 2016) on the CIFAR-10 and CelebA (Liu et al., 2018) datasets, which serve as standard metrics for evaluating generative quality and generalization. As shown in the results, regularizing generator with improves FID and IS, indicating enhanced generalization and higher-quality image generation.

6.3 ADVERSARIAL ROBUSTNESS

We next evaluate the effectiveness of randomized coded inference (RCI) against adversarial attacks on CIFAR-10, focusing on FGSM (Goodfellow et al., 2014) and PGD (Madry et al., 2017) attacks. Since the coded-smoothing module is non-parametric, RCI can be applied to the inference stage of any trained model, with the number of coded samples N adjusted independently of training. Importantly, N can be set relative to the batch size without incurring significant performance degradation (see Table 7 in Appendix G.3 for a sensitivity analysis with respect to batch size).

As shown in Table 3, RCI substantially improves adversarial robustness across all methods, including models already trained with CODED-SMOOTHING, while incurring only a marginal drop in clean (no-attack) accuracy. The strongest results are achieved when models are trained with CODED-SMOOTHING and evaluated with RCI using $N = 1.5K$, where $K = 128$ is the batch size. In this setting, the generalization error increases by only 1%, but robustness gains are significant: improvements of $+8.8\%$ under FGSM ($\epsilon = 8/255$), $+33\%$ under PGD with 10 steps, and $+5.4\%$ under PGD with 100 steps compared to mixup. These results highlight the effectiveness of using RCI in inference for adversarial robustness.

6.4 COVARIATE SHIFT ROBUSTNESS

Finally, we evaluate the performance of the proposed method under distribution shift, where the test distribution differs from the training distribution. For this evaluation, we use CIFAR-10.1 (Recht et al., 2018) and CIFAR-10.2 (Lu et al., 2020), which represent natural covariate shifts of CIFAR-10, as well as CIFAR-10C (Hendrycks & Dietterich, 2019), which introduces 19 types of synthetic corruptions applied at 5 levels of severity to the CIFAR-10 test set. Table 4 in Appendix E compares the performance of our method against ERM and mixup. The coded-smoothing module consistently outperforms both baselines on CIFAR-10.1 and CIFAR-10.2, and achieves comparable performance on CIFAR-10C. For CIFAR-10C, accuracy is reported as the average across all 19 corruption types.

Table 3: Comparison of CIFAR-10 test accuracies under adversarial attacks, contrasting randomized coded inference (RCI) with standard inference. Manifold mixup results are reported from (Verma et al., 2019).

Inference method	Training Method	No Attack -	FGSM $\epsilon = \frac{8}{255}$	PGD 10 steps	PGD 100 steps
Standard inference	ERM	93.7	36.5	5.5	0.0
	Mixup	95.5	71.7	39.9	0.4
	Manifold Mixup	-	77.1*	-	0.0*
	CODED-SMOOTHING (ours)	**95.8**	47.7	8.6	0.0
RSI ($N = 128$)	ERM	55.3	49.1	46.8	19.4
	Mixup	72.4	66.1	64.1	**37.4**
	CODED-SMOOTHING(ours)	72.4	66.2	63.5	27.7
RSI ($N = 190$)	ERM	90.2	75.8	65.7	6.3
	Mixup	93.5	78.2	65.1	9.9
	CSM (ours)	94.8	**80.5**	**72.0**	5.8

7 RELATED WORK

Improving generalization has long been a central challenge in machine learning research. A first class of methods enhances generalization by perturbing hidden representations during training. Classical examples include dropout (Srivastava et al., 2014) and batch normalization (Ioffe & Szegedy, 2015), both of which reduce overfitting by encouraging more robust internal representations.

A second major line of research focuses on data augmentation. Among these, mixup (Zhang et al., 2017) has become a widely adopted regularization strategy. Since its introduction, numerous variants have been proposed to address different limitations of mixup, such as improving generalization (Verma et al., 2019; Yun et al., 2019), adapting it to regression tasks (Yao et al., 2022), enhancing robustness to distribution shift (Pinto et al., 2022), and improving calibration (Bouniot et al., 2023). Despite these extensions, all mixup-style methods fundamentally rely on label information and are thus not applicable in unsupervised settings. The only exception is in GANs (Goodfellow et al., 2020), where mixup regularization has been applied to the supervised discriminator module (Zhang et al., 2017; Verma et al., 2019).

To partially address this limitation, Verma et al. (2022) proposed an unsupervised mixup loss for semi-supervised problems. Their method encourages local linearity by explicitly enforcing the mixup interpolation constraint (see Figure 3). While effective, this approach enforces only pairwise linear constraints, limiting its ability to capture higher-order structures.

In contrast, the proposed coded-smoothingmodule provides a unified regularization framework applicable to both supervised and unsupervised settings with negligible computational overhead. Beyond enforcing linearity, it imposes higher-order smoothness. Moreover, through randomized coded inference, coded-smoothingachieves state-of-the-art robustness against adversarial attacks.

8 CONCLUSION

In this paper, we introduced the coded-smoothing module, a novel regularization framework inspired by coded computing. By enforcing local higher-order smoothness during training, coded-smoothing promotes more generalizable and reliable models. At inference, random shuffling within coded-smoothing, randomized coded inference (RSI), significantly enhances adversarial robustness.

Our method is computationally efficient and applicable to both supervised and unsupervised learning. Across benchmarks and architectures, coded-smoothing improves supervised generalization, outperforming ERM and mixup, while achieving state-of-the-art robustness to adversarial attacks with minimal overhead. In unsupervised settings, applying coded-smoothing to GAN generators boosts generative quality, demonstrating its effectiveness as a label-free regularizer.

REFERENCES

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.

David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32, 2019.

Quentin Bouniot, Pavlo Mozharovskyi, and Florence d'Alché Buc. Tailoring mixup to data for calibration. *arXiv preprint arXiv:2311.01434*, 2023.

Carl De Boor. Calculation of the smoothing spline with weighted roughness measure. *Mathematical Models and Methods in Applied Sciences*, 11(01):33–41, 2001.

Sanghamitra Dutta, Mohammad Fahim, Farzin Haddadpour, Haewon Jeong, Viveck Cadambe, and Pulkit Grover. On the Optimal Recovery Threshold of Coded Matrix Multiplication. *IEEE Transactions on Information Theory*, 66(1):278–301, 2020. ISSN 15579654. doi: 10.1109/TIT.2019.2929328.

Paul HC Eilers and Brian D Marx. Flexible smoothing with b-splines and penalties. *Statistical science*, 11(2):89–121, 1996.

Mohammad Fahim and Viveck R Cadambe. Numerically stable polynomially coded computing. *IEEE Transactions on Information Theory*, 67(5):2758–2785, 2021.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.

Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30, 2017.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456. pmlr, 2015.

Tayyebeh Jahani-Nezhad and Mohammad Ali Maddah-Ali. CodedSketch: A coding scheme for distributed computation of approximated matrix multiplication. *IEEE Transactions on Information Theory*, 67(6):4185–4196, 2021.

Tayyebeh Jahani-Nezhad and Mohammad Ali Maddah-Ali. Berrut approximated coded computing: Straggler resistance beyond polynomial computing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):111–122, 2022.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

Anders Krogh and John Hertz. A simple weight decay can improve generalization. *Advances in neural information processing systems*, 4, 1991.

Yann Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.

Jeremiah Liu, Zi Lin, Shreyas Padhy, Dustin Tran, Tania Bedrax Weiss, and Balaji Lakshminarayanan. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. *Advances in neural information processing systems*, 33:7498–7512, 2020.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Large-scale celebfaces attributes (celeba) dataset. *Retrieved August*, 15(2018):11, 2018.

Shangyun Lu, Bradley Nott, Aaron Olson, Alberto Todeschini, Hossein Vahabi, Yair Carmon, and Ludwig Schmidt. Harder or different? a closer look at distribution shift in dataset reproduction. In *ICML Workshop on Uncertainty and Robustness in Deep Learning*, volume 5, pp. 15, 2020.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

Parsa Moradi and Mohammad Ali Maddah-Ali. General coded computing in a probabilistic straggler regime. *arXiv preprint arXiv:2502.00645*, 2025.

Parsa Moradi, Behrooz Tahmasebi, and Mohammad Maddah-Ali. Coded computing for resilient distributed computing: A learning-theoretic framework. *Advances in Neural Information Processing Systems*, 37:111923–111964, 2024.

Parsa Moradi, Hanzaleh Akbarinodehi, and Mohammad Ali Maddah-Ali. General coded computing: Adversarial settings. *arXiv preprint arXiv:2502.08058*, 2025.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

George M Phillips. *Interpolation and approximation by polynomials*, volume 14. Springer Science & Business Media, 2003.

Francesco Pinto, Harry Yang, Ser Nam Lim, Philip Torr, and Puneet Dokania. Using mixup as a regularizer can surprisingly improve accuracy & out-of-distribution robustness. *Advances in neural information processing systems*, 35:14608–14622, 2022.

Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do cifar-10 classifiers generalize to cifar-10? *arXiv preprint arXiv:1806.00451*, 2018.

Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.

Mahdi Soleymani, Ramy E Ali, Hessam Mahdavifar, and A Salman Avestimehr. ApproxIFER: A model-agnostic approach to resilient and robust prediction serving systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 8342–8350, 2022.

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, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

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*, pp. 2818–2826, 2016.

Lloyd N Trefethen. *Approximation theory and approximation practice, extended edition*. SIAM, 2019.

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 *International conference on machine learning*, pp. 6438–6447. PMLR, 2019.

Vikas Verma, Kenji Kawaguchi, Alex Lamb, Juho Kannala, Arno Solin, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. *Neural Networks*, 145: 90–106, 2022.

Grace Wahba. Smoothing noisy data with spline functions. *Numerische mathematik*, 24(5):383–393, 1975.

Grace Wahba. *Spline models for observational data*. SIAM, 1990.

Yeming Wen, Ghassen Jerfel, Rafael Muller, Michael W Dusenberry, Jasper Snoek, Balaji Lakshminarayanan, and Dustin Tran. Combining ensembles and data augmentation can harm your calibration. *arXiv preprint arXiv:2010.09875*, 2020.

Huaxiu Yao, Yiping Wang, Linjun Zhang, James Y Zou, and Chelsea Finn. C-mixup: Improving generalization in regression. *Advances in neural information processing systems*, 35:3361–3376, 2022.

Qian Yu, Mohammad Maddah-Ali, and Salman Avestimehr. Polynomial codes: an optimal design for high-dimensional coded matrix multiplication. *Advances in Neural Information Processing Systems*, 30, 2017.

Qian Yu, Songze Li, Netanel Raviv, Seyed Mohammadreza Mousavi Kalan, Mahdi Soltanolkotabi, and Salman A Avestimehr. Lagrange coded computing: Optimal design for resiliency, security, and privacy. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1215–1225. PMLR, 2019.

Qian Yu, Mohammad Ali Maddah-Ali, and Amir Salman Avestimehr. Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding. *IEEE Transactions on Information Theory*, 66(3):1920–1933, 2020.

Sangdoo 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/CVF international conference on computer vision*, pp. 6023–6032, 2019.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.