Gaussian MRF Covariance Modeling for Efficient Black-Box Adversarial Attacks

Anit Kumar Sahu Bosch Center for Artificial Intelligence anit.sahu@gmail.com Satya Narayan Shukla College of Information and Computer Sciences University of Massachusetts Amherst snshukla@cs.umass.edu

J. Zico Kolter Bosch Center for Artificial Intelligence Carnegie Mellon University zkolter@cs.cmu.edu

Abstract

We study the problem of generating adversarial examples in a black-box setting, where we only have access to a zeroth order oracle, providing us with loss function evaluations. Although this setting has been investigated in previous work, most past approaches using zeroth order optimization implicitly assume that the gradients of the loss function with respect to the input images are *unstructured*. In this work, we show that in fact substantial correlations exist within these gradients, and we propose to capture these correlations via a Gaussian Markov random field (GMRF). Given the intractability of the explicit covariance structure of the MRF, we show that the covariance structure can be efficiently represented using the Fast Fourier Transform (FFT), along with low-rank updates to perform exact posterior estimation under this model. We use this modeling technique to find fast one-step adversarial attacks, akin to a black-box version of the Fast Gradient Sign Method (FGSM), and show that the method uses fewer queries and achieves higher attack success rates than the current state of the art. We also highlight the general applicability of this gradient modeling setup.

1 Introduction

Most methods for adversarial attacks on deep learning models operate in the so-called *white-box* setting [7], where the model being attacked, and its gradients, are assumed to be fully known. Recently, however there has also been considerable attention given to the *black-box* setting as well, where the model is unknown and can only be queried by a user, and which much better captures the "typical" state by which an attacker can interact with a model [5, 26, 11, 17]. And several past methods in this area have conclusively demonstrated that, given sufficient number of queries, it is possible to achieve similarly effective attacks in the black-box setting akin to the white-box setting. However, as has also been demonstrated by past work [11, 10, 2], the efficiency of these black-box attacks (the number of queries need to find an adversarial example) is fundamentally limited unless they can exploit the spatial correlation structure inherent in the model's gradients. Yet, at the same time, most previous methods have used rather ad-hoc methods of modeling such correlation structure, such as using "tiling" bases and priors over time [11] or explicit square regions of the input space [2] that require attack vectors be constant over large regions, or by other means such as using smoothly-varying perturbations [10] to estimate these gradients.

In this work, we present a new, more principled approach to model the correlation structure of the gradients within the black-box adversarial setting. In particular, we propose to model the gradient

of the model loss function with respect to the input image using a Gaussian Markov Random Field (GMRF). This approach offers a number of advantages over prior methods: 1) it naturally captures the spatial correlation observed empirically in most deep learning models; 2) using the model, we are able to compute exact posterior estimates of the true gradient given observed data, while also fitting the parameters of the GMRF itself via an expectation maximization (EM) approach; and 3) the method provides a natural alternative to uniformly sampling perturbations, based upon the eigenvectors of the covariance matrix. Although representing the joint covariance over the entire input image may seem intractable for large-scale images, we can efficiently compute necessary terms for very general forms of grid-based GMRFs using the Fast Fourier Transform (FFT).

We evaluate our approach by attempting to find adversarial examples, over multiple different data sets and model architectures, using the GMRF combined with a very simple greedy zeroth order search technique; the method effectively forms a "black-box" version of the fast gradient sign method (FGSM), by constructing an estimate of the gradient at the input image itself, then taking a single signed step in this direction. Despite its simplicity, we show that owing to the correlation structure provided by the GMRF model, the approach outperforms more complex gradient-based approaches such as the BANDITS-TD [11] or PARSIMONIOUS [17] methods, especially for small query budgets.

2 Related Work

Most black-box adversarial attack methods catered towards untargeted attacks either rely on multi iteration optimization schemes or transfer-style attacks using substitute models. Black-box adversarial attacks can be broadly categorized across a few different dimensions: optimization-based versus transfer-based attacks, and score-based versus decision-based attacks.

In the optimization-based adversarial setting, the adversarial attack is formulated as the problem of maximizing some loss function (e.g., the accuracy of the classifier or some continuous variant) using a zeroth order oracle, i.e., by making queries to the classifier. And within this optimization setting, there is an important distinction between score-based attacks, which directly observe a traditional model loss, class probability, or other continuous output of the classifier on a given example, versus decision-based attacks, which only observe the hard label predicted by the classifier. Decision based attacks have been studied by much past research [3, 4, 5, 23], and (not surprisingly) typically require more queries to the classifier than the score-based setting. Most black-box adversarial attacks can be categorized into score based and decision based attacks. Multi iteration optimization based methods typically involve algorithms which access a zeroth order oracle, i.e., make queries to the model to obtain loss function evaluations in order to maximize the loss function so as to cause misclassification. Decision-based attacks, where the attacker only has access to the predicted label of the input by the model has been studied in Brendel et al. [3], Chen et al. [4].

In the setting of score-based attacks, the first such iterative attack on a class of binary classifiers was first studied in [19]. A real-world application of black-box attacks to fool a PDF malware classifier was demonstrated by [27], for which a genetic algorithm was used. Narodytska and Kasiviswanathan [18] demonstrated the first black-box attack on deep neural networks. Subsequently black-box attacks based on zeroth order optimization schemes, using techniques such as KWSA [12] and RDSA [20] were developed in [5, 10]. Though Chen et al. [5] generated successful attacks attaining high attack accuracy, the method was found to be extremely query hungry which was then remedied to an extent by Ilyas et al. [10]. In [11], the authors exploit correlation of gradients across iterations by setting a prior and use a piece wise constant perturbation, i.e., tiling to develop a query efficient black-box method. In [1], the authors focused on estimating the signed version of the gradient to generate efficient black-box attacks. Recently Moon et al. [17] used a combinatorial optimization perspective to address the black-box adversarial attack problem.

A concurrent line of work [21] has considered the transfer-based setting. These approaches create adversarial attacks by training a substitute network with the aim to mimic the target model's decisions, which are then obtained through black-box queries. However, substitute network based attack strategies have been found to have a higher query complexity than those based on gradient estimation.

The exploitation of the structure of the input data space so as to append a regularizer has been recently found to be effective for robust learning. In particular, in [15] showed that by using Wasserstein-2 geometry to capture semantically meaningful neighborhoods in the space of images helps to learn discriminative models that are robust to in-class variations of the input data.

Setting of this work In this paper, we are specifically focused on the optimization-based, scorebased setting, following most directly upon the work of Chen et al. [5], Ilyas et al. [10, 11], Moon et al. [17]. However, our contribution is also largely orthogonal to the methods presented in these prior works, and indeed could be combined with any gradient-based method. Specifically, we show that by modeling the covariance structure of the gradient using a Gaussian MRF, a very simple approach achieves performance that is competitive with the best current methods in low query budgets. We further emphasize that this GMRF approach can be applied to other black-box search strategies as well.

3 Preliminaries: Adversarial Attacks

In the context of classifiers, adversarial examples are carefully crafted inputs to the classifier which have been perturbed by an additive perturbation so as to cause the classifier to misclassify the input. Formally, we define a classifier $C : \mathcal{X} \mapsto \mathcal{Y}$ and its corresponding classification loss function $L(\mathbf{x}, y)$ (typically taken to be the cross-entropy loss on the class logits produced by the classifier), where $\mathbf{x} \in \mathcal{X}$ is the input to the classifier, $y \in \mathcal{Y}$, \mathcal{X} is the set of inputs and \mathcal{Y} is the set of labels. The objective of generating a misclassified example can be posed as an optimization problem. In particular, the aim is to generate an adversarial example \mathbf{x}' for a given input \mathbf{x} which maximizes $L(\mathbf{x}', y)$ but still remains ϵ_p -close in terms of a specified metric to the original input. Thus, the generation of an adversarial attack can be formalized as a constrained optimization as follows:

$$\mathbf{x}' = \underset{\mathbf{x}':\|\mathbf{x}'-\mathbf{x}\|_{p} \le \epsilon_{p}}{\arg\max} L(\mathbf{x}', y).$$
(1)

We give a brief overview of adversarial attacks categorized in terms of access to information namely, white-box and black-box attacks.

White-box attacks. White-box settings assume access to the entire classifier and the analytical form of the (possibly non-convex) classifier's loss function. One of the original methods for generating such attacks is the Fast Gradient Sign Method (FGSM) [7], which remains relatively successful at attacking undefended models. FGSM forms the attack using a single steepest descent update under the ℓ_{∞} norm (similar analogues exist for other norms, though are technically then no longer gradient "sign" method), which corresponds to the update

$$\mathbf{x}' = \mathbf{x} + \epsilon_p \operatorname{sign} \left(\nabla L(\mathbf{x}, y) \right). \tag{2}$$

A stronger and slightly more general attack uses the PGD algorithm to generate adversarial examples [16] (which was first introduced in the adversarial examples setting as the Basic Iterative method [13] in the ℓ_{∞} setting), which corresponds to repeating the the updates

$$\mathbf{x}_{l} = \Pi_{B_{p}(\mathbf{x},\epsilon)}(\mathbf{x}_{l-1} + \eta \mathbf{s}_{l}) \text{ with } \mathbf{s}_{l} = \Pi_{\partial B_{p}(0,1)} \nabla_{x} L(\mathbf{x}_{l-1}, y),$$
(3)

where Π_S denotes the projection onto the set S, $B_p(\mathbf{x}', \varepsilon')$ is the ℓ_p ball of radius ε' centered at \mathbf{x}', η denotes the step size, and ∂U is the boundary of a set U (projection onto this set, e.g. takes the sign of the gradient for the ℓ_∞ ball, or normalizes the graident for the ℓ_p ball). By making, \mathbf{s}_l to be the projection of the gradient $\nabla_x L(\mathbf{x}_{l-1}, y)$ at x_{l-1} onto the unit ℓ_p ball, it is ensured that \mathbf{s}_l is the unit ℓ_p -norm vector that has the largest inner product with $\nabla_x L(\mathbf{x}_{l-1}, y)$. However, in most real world deployments, it is impractical to assume complete access to the classifier and analytic form of the corresponding loss function, which makes black-box settings more realistic.

Black-box atacks. In a typical *black-box* setting, the adversary's access is limited to the value of the loss function $L(\mathbf{x}, y)$ for an input (\mathbf{x}, y) . The crux of black-box methods can be divided into two classes. The first class consists of methods which operate in a derivative free manner, i.e., directly seek for the next iterate without estimating gradients. Evolution strategies such as CMA [8] or alternatively Bayesian optimization method [22] could be used for such attacks. The second class more directly attempts estimate the gradient through, where our method is situated.

The main building block of gradient-based black-box methods is the use of some difference schemes to estimate gradients. A typical two-sided finite difference scheme, called the random direction stochastic approximation (RDSA), for instance, computes the estimate

$$\widehat{\nabla}_{x}L(\mathbf{x}, y) = \frac{1}{m} \sum_{k=1}^{m} \mathbf{z}_{k} \left(L(\mathbf{x} + \delta \mathbf{z}_{k}, y) - L(\mathbf{x}, y) \right) / \delta, \tag{4}$$



Figure 1: Gradient autocorrelation shown for a 9x9 spatial window for (left) 100 MNIST images with a LeNet-like architecture; and (right) 50 ImageNet images on the VGG16 model for one of the color channels.

where directions $\{\mathbf{z}_k\}_{k=1}^m$ can be canonical basis vectors, or normally distributed vectors, or the Fast Fourier Transform basis (FFT) basis. The step size $\delta > 0$ is a key parameter of choice; a higher δ could lead to extremely biased estimates, while a lower δ can lead to an unstable estimator. Owing to the biased gradient estimates, a finite difference based black-box attack turns out to be query hungry. In particular, in order to ensure sufficient increase of the objective at each iteration, the query complexity scales with dimension and hence is prohibitively large.

In this setting, the crux of our approach we present next is a method for *decreasing* the variance of these gradient estimates, by exploiting the spatial correlation that arises due to the structure of deep networks. We then apply our approach within the context of a simple FGSM-like method, illustrating that even a simple algorithm, when combined with these more efficient gradient methods, can outperform many existing approaches in the low-query setting. In principle, however, our gradient estimation method can be combined with any adversarial attack that attempts to estimate gradients.

4 Gradient Correlation and Query Efficient Black-Box Attacks

In this section, we present the main algorithmic contribution of our work, a method for efficiently exploiting correlation to estimate network gradients, plus an application to a single-step FGSM-like attack.

4.1 Gradient Correlation

In most black-box adversarial attacks, the gradient terms across different images are implicitly assumed to be independent from each other. However, empirically we find that the gradients across different images exhibit substantial correlation across dimensions and across images. In Figure 1, we plot the autocorrelation over a 9×9 window for the true gradients from 100 images sampled from the MNIST test set on a LeNet model, for 50 images sampled from the ImageNet validation set for the VGG16-bn model. In both cases it is evident that there exists substantial correlation between nearby points.

We propose to exploit and model these correlations using a Gaussian Markov random field so as to reduce the variance of our gradient estimate. Formally, letting x be the input to a classifier, and $\mathbf{g} = \nabla L(\mathbf{x}, y)$ the gradient of the loss function with respect to the input, we are attempting to query loss function values and estimate the gradient, then we aim to put a prior distribution over \mathbf{g}

$$\mathbf{g} \sim \mathcal{N}(0, \Sigma) \tag{5}$$

where Σ is a non-identity covariance matrix modeling the correlation between terms. Following common practice we model the inverse covariance matrix $\Lambda = \Sigma^{-1}$, a setting also known as the Gaussian Markov random field (GMRF) setting, given that the non-zero entries in Λ correspond exactly to the edges in a graphical model describing the distribution. Furthermore, we use a parameterized Gaussian MRF with relatively few free parameters. For example, if x is a 2D image, then we may have one parameter α governing the diagonal terms $\Lambda_{i,i} = \alpha, \forall i$, and another governing adjacent pixels $\Lambda_{i,j} = \beta$ for i, j corresponding to indices that are neighboring in the original image. We will jointly refer to all the parameters of this model as θ , so in this case $\theta = (\alpha, \beta)$, and we refer to the resulting Λ as $\Lambda(\theta)$.

4.2 Learning Gradient Correlation

We first consider the problem of fitting a parameterized MRF model to given inputs $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(m)}$. We further apply *n* random directions $\mathbf{u}^{(1)}, \ldots, u^{(n)}$ to each image to produce multiple RDSA-like gradient estimates for each image $\mathbf{G} = [\mathbf{g}^{(1)}, \ldots, \mathbf{g}^{(mn)}]$ where

$$\mathbf{g}^{(ij)} = \mathbf{u}^{(j)} (L(\mathbf{x}^{(i)} + \delta \mathbf{u}^{(j)}, y^{(i)}) - L(\mathbf{x}^{(i)}, y^{(i)})) / \delta.$$
(6)

Given this data, the maximum likelihood estimation for the MRF parameters is simply the standard Gaussian maximum likelihood estimate, given by the optimization problem

$$\min_{\theta} \operatorname{tr}(\mathbf{S}\Lambda(\theta)) - \operatorname{logdet}(\Lambda(\theta)), \tag{7}$$

where $\mathbf{S} = \frac{1}{mn} \sum_{i} \mathbf{g}^{(i)} \mathbf{g}^{(i)^{\top}}$ is the sample covariance and logdet denotes the log determinant. While this is a trivial problem for the case of general covariance, when we use a parameterized form of Λ , it becomes less clear how to solve this optimization problem efficiently for large inputs.

As we show, however, this optimization problem can be easily solved using the Fourier Transform; we focus for simplicity of presentation on the 2D case, but the method is easily generalizable to three dimensional convolutions to capture color channels in addition to the spatial dimensions itself. First, we focus on evaluating the trace term. The key idea here is that the Λ operator can be viewed as a (circular) convolution¹

$$\mathbf{K}(\theta) = \begin{bmatrix} 0 & \beta & 0\\ \beta & \alpha & \beta\\ 0 & \beta & 0 \end{bmatrix},\tag{8}$$

which then lets us compute $tr(S\Lambda(\theta))$ as sum of the elements of the product of **G** and the zero padded 2D convolution of **G** and $\mathbf{K}(\theta)$. For the log determinant term, we again use the fact that Λ is a convolution operator. Specifically, because it is a convolution, it can be diagonalized using the fast Fourier transform

$$\Lambda(\theta) = \mathbf{Q}^H \mathbf{D} \mathbf{Q}$$

where \mathbf{Q} is the the Fast Fourier Transform (FFT) basis, and the eigenvalues being the diagonal elements of \mathbf{D} can be found by a FFT to the zero-padded convolution operator; thus, we can compute the log determinant term by simply taking the sum of the log of the FFT-computed eigenvalues. Once the efficient computation of the objective function is in place, we then employ Newton's method to optimize the objective, with the gradient and Hessian terms all computed via automatic differentiation (since there typically very few parameters, even computing the full Hessian matrix is straightforward). The entire procedure of estimating the GMRF parameters is depicted in Algorithm 1. Other appropriate optimization schemes could also be used instead of Newton's method, if desired.

For a $N \times N$ sized image, the dominating cost for the procedure is the $O(N^2 \log N)$ computation of the 2D FFT; this contrasts with the $O(N^6)$ naive complexity of the eigen decomposition of the $N^2 \times N^2$ inverse covariance. Thus, the use of the properties of convolution operator and FFT basis makes the optimization in equation 7 to be feasible: for example, computing the log determinant term for the VGG16-bn architecture for ImageNet would involve performing eigenvalue decomposition of a matrix with dimensions 150528×150528 , which is computationally prohibitive without using the FFT basis.

4.3 Application to Estimating Adversarial Directions

With the GMRF estimation framework in place, our next task is to estimate adversarial directions which involves estimating the posterior mean. Next, we apply the GMRF-based gradient estimate to the setting of constructing black-box attacks in a query efficient manner. As we show later, exploiting the correlation across gradients substantially reduces query complexity without compromising on

¹The FFT operation technically operates circular convolutions (meaning the convolution wraps around the image), and thus the covariance naively models a correlation between, e.g., the first and last rows of an image. However, this is a minor issue in practice since: 1) it can be largely mitigated by zero-padding the input image before applying the FFT-based convolution, and 2) even if ignored entirely, the effect of a few additional circular terms in the covariance estimation is minimal.

Algorithm 1 Learning GMRF Parameters

1: procedure LEARNGMRF($\{\mathbf{x}^{(i)}, y^{(i)}\}_{i=1}^{m}, \delta$) 2: Draw *n* vectors $\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(n)} \sim \mathcal{N}(0, \mathbf{I})$ 3: Compute gradient estimates $\mathbf{g}^{(1)}, \dots, \mathbf{g}^{(mn)}$ where $\mathbf{g}^{(ij)} = \mathbf{u}^{(j)}(L(\mathbf{x}^{(i)} + \delta \mathbf{u}^{(j)}, y^{(i)}) - L(\mathbf{x}^{(i)}, y^{(i)}))/\delta$ 4: while Not converged do 5: Compute objective: $f(\theta) = \operatorname{tr}(\mathbf{S}\Lambda(\theta)) - \operatorname{logdet}(\Lambda(\theta))$ $= \operatorname{sum}(\mathbf{G} \times \operatorname{conv2d}(\mathbf{K}(\theta), \mathbf{G})) + \operatorname{sum}(\operatorname{log}(\operatorname{FFT}(\operatorname{Pad}(\mathbf{K}(\theta))))))$ 6: Perform Newton update: $\theta := \theta - \nabla^2 f(\theta)^{-1} \nabla f(\theta)$ (using automatic differentiation) 7: return θ

the attack success rate. Under the aforementioned GMRF framework, we can interpret black-box gradient estimation as a Gaussian inference problem. Specifically, we have assumed that the gradient at an input x follows the normal distribution with the prescribed inverse covariance

$$\mathbf{g} \sim \mathcal{N}(0, \Lambda^{-1})$$

We next construct a linear observation model by observing the fact that the loss function value at some point \mathbf{x}' can be viewed as a noisy observation of a linear function of the gradient

$$L(\mathbf{x}') \approx L(\mathbf{x}) + \mathbf{g}^{\top}(\mathbf{x}' - \mathbf{x})$$

where we drop the y in $L(\mathbf{x}, y)$ for notational simplicity. In order to estimate the posterior mean for \mathbf{g} given loss function values, we perturb the input using scaled version of supplied vectors $\{\mathbf{z}^{(1)}\}_{i=1}^{m}$, i.e., $\delta_1\{\mathbf{z}^{(1)}\}_{i=1}^{m}$ so as to obtain the perturbed points $\{\mathbf{x}'^{(i)}\}_{i=1}^{m}$. Thus, given a set of sample points $\mathbf{x}'^{(1)}, \ldots, \mathbf{x}'^{(m)}$ (which here refer to different *perturbations* of a single underlying point \mathbf{x}), corresponding loss function values $L(\mathbf{x}'^{(1)}), \ldots, L(\mathbf{x}'^{(m)})$, we have the following characterization of the distribution

$$\mathbf{L}|\mathbf{g} \sim \mathcal{N}(\mathbf{X}\mathbf{g}, \sigma^2 \mathbf{I})$$

where

$$\mathbf{L} = \left[L(\mathbf{x}^{\prime(1)}) - L(\mathbf{x}), \cdots, L(\mathbf{x}^{\prime(m)}) - L(\mathbf{x}) \right], \mathbf{X} = \left[(\mathbf{x}^{\prime(1)} - \mathbf{x}^{\prime})^{\top}, \cdots, (\mathbf{x}^{\prime(m)} - \mathbf{x}^{\prime})^{\top} \right].$$
(9)

With the above formalism in place, using standard Gaussian manipulation rules, the posterior of interest g|L, is given by

$$\mathbf{g}|\mathbf{L} \sim \mathcal{N}\left(\left(\boldsymbol{\Lambda} + \mathbf{X}^{\top}\mathbf{X}/\sigma^{2}\right)^{-1}\mathbf{X}^{\top}\mathbf{L}_{1}/\sigma^{2}, \left(\boldsymbol{\Lambda} + \mathbf{X}^{\top}\mathbf{X}/\sigma^{2}\right)^{-1}\right).$$
(10)

To compute this estimate, we need to solve a linear equation in the matrix $\Lambda + \mathbf{X}^{\top}\mathbf{X}/\sigma^2$. This is a convolution plus a low rank matrix (the $\mathbf{X}^{\top}\mathbf{X}$ term is rank m, and we typically have $m \ll n$ because we have relatively few samples and a high dimensional input). We can not solve for this matrix exactly using the FFT, but we can still solve for it efficiently (requiring only an $m \times m$ inverse) using the matrix inversion lemma, specifically using Woodburry's matrix inversion lemma

$$\left(\Lambda + \mathbf{X}^{\top} \mathbf{X} / \sigma^2\right)^{-1} = \Lambda^{-1} - \Lambda^{-1} \mathbf{X}^{\top} \left(\sigma^2 \mathbf{I} + \mathbf{X} \Lambda^{-1} \mathbf{X}^{\top}\right)^{-1} \mathbf{X} \Lambda^{-1}.$$
 (11)

Since the term needs to be computed explicitly for at least the inner inverse, we explicitly maintain the term $\mathbf{U} = \Lambda^{-1} \mathbf{X}^{\top}$. Note that in the sequential sampling setting (where we sequentially sample $\mathbf{x}^{(i)}$ points one at a time), this matrix can be maintained over all samples, and the factorization updated sequentially.

Finally, we use this posterior inference within a simple FGSM-like attack method for ℓ_{∞} attacks, which first computes the mean estimate of the gradient $\hat{\mathbf{g}}$, then constructs an adversarial example in the direction of its sign. One free parameter of the method is how we choose the perturbations to create $\mathbf{x'}^{(1:m)}$. While we could use simple Gaussian perturbations, our covariance structure strongly motivates an alternative choice: using the top FFT basis vectors, which correspond to the dominant eigenvalue of the Covariance, and thus given the directions of highest expected variance. While FFT vectors have been used previously in adversarial attacks [28], our formulation thus further provides a motivation for their adoption. The full method is shown in Algorithm 2.

Algorithm 2 GMRF based black-box FGSM

1:	procedure	BB-FGSM(σ_{i}	$(\delta, \epsilon, \theta)$
----	-----------	-----------------------	------------------------------

- 2: Construct perturbed points: $\mathbf{x}^{\prime(1:m)} := \mathbf{x} + \delta \cdot \text{Top FFT basis vectors}$
- 3: Form L and X using equation 9
- 4: Using FFT, compute: $\hat{\mathbf{g}} := \left(\Lambda(\theta) + \mathbf{X}^{\top}\mathbf{X}/\sigma^2\right)^{-1}\mathbf{X}^{\top}\mathbf{L}/\sigma^2$
- 5: **return** $\mathbf{x} + \epsilon \cdot \operatorname{sign}(\hat{\mathbf{g}})$

A note on the threat model. One important point to note is that, unlike typical black box models, the BB-FGSM algorithm implicitly uses some information from *multiple* images to form it's attack (because we estimate the GMRF parameters θ on separate images, using additional queries to the classifier). However, this is a relatively minor difference in practice: for a given classifier we estimate θ once offline, and because θ has so few parameters (typically just 3–4), we need a very small number of total queries to the classifier to estimate it (e.g., we use 10 additional images with 10 queries each to estimate the parameters for ImageNet classifier). After learning these parameters offline, they can be used indefinitely to run attacks against the classifier, essentially amortizing away any computational cost of estimating the parameters.

In order to achieve query efficiency our framework incorporates gradient correlation into gradient estimation through the GMRF framework which we then further exploit by the procedure described above. In order to estimate the adversarial direction efficiently in Algorithm ??, a key role is played by the vectors $\{\mathbf{z}^{(1)}\}_{i=1}^{m}$ which perturb the input. One particular choice being sampling the directions from a normal distribution. However, it is worth noting that the inverse covariance distribution of the gradients by construction is a convolution operator and hence is diagonalized by the FFT basis. Low frequency components of the FFT basis vectors enhanced the attack accuracy significantly, an example of which is depicted in Figure 7 for black-box attacks on VGG16-bn classifier for ImageNet with the ℓ_{∞} perturbation constrained to $\epsilon = 0.05$. The adversarial input for x is generated using FGSM as follows:

$$\mathbf{x}_{adv} = \mathbf{x} + \epsilon \operatorname{sign}\left(\widehat{\mathbf{g}}\right). \tag{12}$$

The entire procedure consisting of GMRF inference and gradient inference is presented in Algorithm 2. We defer the discussion of computing FFT basis vectors to the appendix.

5 Experiments

In this section, we evaluate our proposed gradient estimation and resulting black box attack on the MNIST [14] and ImageNet [6] data sets, considering untargeted ℓ_{∞} attacks in particular. We evaluate the performance of our approach against several alternative gradient-based black-box attacks, in terms of attack success rate² with a given query budget and average query count (we don't compare to non-gradient-based methods, as the overall evaluation here is to emphasize the importance of gradient estimation in the low-query setting, not necessarily to demonstrate the absolute most success black-box attacks). More complete experimental details are presented in the appendix. The code is available at https://github.com/anitksahu/GMRF.

5.1 Experiments on MNIST

We first compare the performance of the proposed method with that of white-box FGSM [7] across different values of ℓ_{∞} bounds ranging from 0.05 to 0.3 in increments of 0.05 over query budgets from 20 to 200. To further illustrate the importance of incorporating a non-identity gradient covariance, we also provide experimental results for a version of our proposed algorithm which takes the gradient covariance to be an identity matrix.

As shown in Figure 4, our proposed attack method exhibits better attack accuracy than white-box FGSM in around 75 queries and the gap in the performance is further magnified with increasing

²The attack success rate is defined as the ratio of the number of images successfully misclassified to the total number of input images. Among the set of input images, we discard images which were misclassified by the classifier. The average query count is computed on the queries made for successful attacks.



Figure 2: Performance comparison of the proposed attack, the proposed attack with identity gradient covariance and white-box FGSM in terms of attack accuracy with query budgets between 20 and 200 for MNIST on LeNet.



Figure 3: Performance comparison of the proposed attack (GMRF), the proposed attack with identity gradient covariance and white-box FGSM in terms of attack accuracy with query budgets between 20 and 200 for MNIST on LeNet.

number of queries. On the other hand, the version of our algorithm with identity gradient covariance consistently under performs with respect to the white-box FGSM attack. This further illustrates the effectiveness of our proposed algorithm and the importance of modelling the gradient covariance. We defer additional discussions concerning cosine similarities of the adversarial directions of the proposed GMRF based attack framework and white-box FGSM to the appendix.

The superior performance of our proposed framework as compared to white-box FGSM as demonstrated in Figure 4 can be attributed to the following reason. First, incorporating the gradient non-identity covariance structure into the gradient estimation scheme, allows our perturbation to be able to use structural gradient information from other images too. On the other hand, white box FGSM treats gradient of every image to be independent of gradients of other images. This is further illustrated by our experimental findings based on the version of our algorithm considering the gradient covariance to be an identity matrix.

5.2 Experiments on ImageNet

We next compare the performance of the proposed method on ImageNet with that of NES [10], BANDITS-TD [11], a version of NES using FFT basis vectors and PARSIMONIOUS [17], which are



Figure 4: VGG16-bn: Attack Figure 5: ResNet50: Attack suc- Figure 6: Inception-v3: Attack success rate as a function of cess rate as a function of query success rate as a function of query budget

budget

query budget

Table 1: Summary of ℓ_{∞} attacks with $\epsilon = 0.05$ ImageNet attacks using NES, BANDITS-TD, PARSIMONIOUS and the proposed method with a query budget of 1000 per image

Attack	VGG16-bn		Resnet50		Inception-v3	
	Success	Avg. Queries	Success	Avg. Queries	Success	Avg. Queries
NES	72.52%	452.86	55.65%	463.63	45.03%	448.40
NES-FFT	84.43%	395.10	70.53%	405.86	53.59%	417.80
Parsimonious	99.06%	165.47	97.33%	178.61	79.89%	253.95
Bandits _{TD}	85.79%	141.16	77.49%	190.47	68.40%	204.06
Proposed	95.17%	56.79	92.01%	84.06	75.03%	170.09

the current state of the art in ℓ_∞ gradient-based black-box attacks. We consider three pre-trained classifiers namely, ResNet50 [9], Inception-v3 [25] and VGG16-bn [24]. We use the pre-trained models provided by PyTorch for attacking these classifiers. We evaluation on a randomly sampled set of 1000 images from ImageNet validation set.

The ℓ_{∞} perturbation bound is set to $\epsilon = 0.05$ (with images noramlized to [0,1]. We use the implementation and hyperparameters provided by Ilyas et al. [11] for NES and BANDITS-TD. Similarly for PARSIMONIOUS, we use the implementation and hyperparameters given by [17]. The specifics of the GMRF model, the values of the parameters and the associated hyperparameters (which were obtained by grid search) for the proposed algorithm for the three classifiers are relegated to the Appendix.

Figures 4 - 6 show the evolution of attack accuracy with different querying budgets (the amortized number of queries needed to estimate GMRF parameters are also added to the totals for the GMRF method, but this adds an average of only 0.2 queries per example). In the low query budget setting, with less than 450 query budget, our algorithm outperforms PARSIMONIOUS, though it exhibits slightly inferior performance with higher queries. As shown in Table 1, our proposed algorithm, in spite of being a single-step attack, outperforms BANDITS-TD, NES-FFT and NES by achieving higher attack success rate. Despite the higher success rate, the proposed method uses much fewer queries on an average as compared to BANDITS-TD, NES-FFT and NES. Thus, the proposed method strictly dominates BANDITS-TD, NES-FFT and NES on every metric.

The proficiency of our proposed scheme in the low query budget regime can be attributed to the utilization of the GMRF framework and the usage of FFT basis vectors. In particular, the FFT basis vectors being the eigen vectors of the covariance matrix provide for systematic dimensionality reduction.

Conclusion 6

In this paper, we proposed an efficient method for modeling the covariance of gradient estimated in deep networks using Gaussian Markov random fields. Although the method naively would be computationally impractical, we show that by employing an FFT approach, we can efficiently learn and perform inference over the Gaussian covariance structure. We then use this approach to highlight an efficient gradient-based black-box adversarial attack, which, owing to capturing this correlational structure, typically requires fewer samples than existing gradient-based techniques.

Although the highlight of this work involved using our gradient estimation procedure within a black box adversarial attack, the proposed methods also have applicability in any setting where one wishes to obtain an estimate of neural network input gradients, using limited queries to the model. We hope to explore further applications and algorithmic advances using these techniques is future work.

7 Broader Impacts

The sensitivity of deep learning classifiers to adversarial attacks raises serious concerns about their practical application in many real-world, safety-critical domains. Black box attacks, in particular, highlight the potential of these attacks to apply to practically deployed systems, where we do not have access to the full classifier definition. At the same time, the methods proposed here are not easily applied immediately to many deployed classifiers in the real world, as they require multiple classifications of the same input with different perturbations, and also require access to the loss. Thus, while this is a step toward demonstrating the feasibility of real-world attacks, there are still obstacles to overcome if adversaries wished to deploy them in real-world settings.

Nonetheless, the shift towards more black-box versions of adversarial attacks continues to pose a major challenge for deployed systems, which should be constantly evaluated. Our hope is that the presence of effective attacks makes it clear the limitations of such systems, and ultimately encourages the development of more robust models to deploy in the real world.

References

- [1] A. Al-Dujaili and U.-M. O'Reilly. There are no bit parts for sign bits in black-box attacks. *arXiv preprint arXiv:1902.06894*, 2019.
- [2] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein. Square attack: a query-efficient black-box adversarial attack via random search. *arXiv preprint arXiv:1912.00049*, 2019.
- [3] W. Brendel, J. Rauber, and M. Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. *arXiv preprint arXiv:1712.04248*, 2017.
- [4] J. Chen, M. I. Jordan, and M. J. Wainwright. Hopskipjumpattack: A query-efficient decisionbased attack. In 2020 IEEE Symposium on Security and Privacy (SP), 2019.
- [5] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, AISec '17, pages 15–26, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-5202-4. doi: 10.1145/3128572.3140448. URL http://doi.acm.org/10.1145/3128572.3140448.
- [6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [7] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572, 2014.
- [8] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- [9] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778, 2015.
- [10] A. Ilyas, L. Engstrom, A. Athalye, and J. Lin. Black-box adversarial attacks with limited queries and information. In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2137–2146, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL http://proceedings.mlr.press/v80/ilyas18a.html.
- [11] A. Ilyas, L. Engstrom, and A. Madry. Prior convictions: Black-box adversarial attacks with bandits and priors. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=BkMiWhR5K7.

- [12] J. Kiefer, J. Wolfowitz, et al. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3):462–466, 1952.
- [13] A. Kurakin, I. Goodfellow, and S. Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016.
- [14] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.
- [15] A. T. Lin, Y. Dukler, W. Li, and G. Montúfar. Wasserstein diffusion tikhonov regularization. arXiv preprint arXiv:1909.06860, 2019.
- [16] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=rJzIBfZAb.
- [17] S. Moon, G. An, and H. O. Song. Parsimonious black-box adversarial attacks via efficient combinatorial optimization. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the* 36th International Conference on Machine Learning, volume 97 of Proceedings of Machine Learning Research, pages 4636–4645, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL http://proceedings.mlr.press/v97/moon19a.html.
- [18] N. Narodytska and S. Kasiviswanathan. Simple black-box adversarial attacks on deep neural networks. In 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 1310–1318. IEEE, 2017.
- [19] B. Nelson, B. I. Rubinstein, L. Huang, A. D. Joseph, S. J. Lee, S. Rao, and J. Tygar. Query strategies for evading convex-inducing classifiers. *Journal of Machine Learning Research*, 13 (May):1293–1332, 2012.
- [20] Y. Nesterov and V. Spokoiny. Random gradient-free minimization of convex functions. Foundations of Computational Mathematics, 17(2):527–566, 2017.
- [21] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, pages 506–519. ACM, 2017.
- [22] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- [23] S. N. Shukla, A. K. Sahu, D. Willmott, and J. Z. Kolter. Black-box adversarial attacks with bayesian optimization. arXiv preprint arXiv:1909.13857, 2019.
- [24] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition, 2014. cite arxiv:1409.1556.
- [25] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2818–2826, 2015.
- [26] C. Tu, P. Ting, P. Chen, S. Liu, H. Zhang, J. Yi, C. Hsieh, and S. Cheng. Autozoom: Autoencoderbased zeroth order optimization method for attacking black-box neural networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January* 27 - February 1, 2019., pages 742–749, 2019. URL https://aaai.org/ojs/index.php/ AAAI/article/view/3852.
- [27] W. Xu, Y. Qi, and D. Evans. Automatically evading classifiers. In Proceedings of the 2016 Network and Distributed Systems Symposium, pages 21–24, 2016.
- [28] D. Yin, R. G. Lopes, J. Shlens, E. D. Cubuk, and J. Gilmer. A fourier perspective on model robustness in computer vision. In *Advances in Neural Information Processing Systems*, pages 13255–13265, 2019.



Figure 7: FFT basis vectors vs normal distribution for VGG16-bn.

A Appendix

A.1 Computation of FFT basis

We use the fact that the covariance and the inverse covariance matrix because of being convolutional operators are diagonalized by the FFT basis. Let us assume the image is of size $c \times h \times w$, where c,h and w denote the number of channels, height and width of the image. We define a tensor S of zeros of size $c \times h \times w \times 2$, where the last dimension is to account for both the real and complex components. In order to generate the lowest frequency basis vector, we set the the first element of the tensor of the first channel, i.e, $S_{0,0,0,0} = 1$ and take the inverse FFT. This gives us the lowest cosine basis vector. We do the same for the other channels, by just setting the corresponding component to 1 and taking the inverse FFT. Setting, $S_{0,0,0,1} = 1$ and then taking the inverse FFT yields the lowest frequent sine component. In order to generate the low frequency components, we start from the beginning of a row and proceed along diagonally by incrementing the row and column index by one. At each entry of the tensor, we repeat it for every channel once at a time.

A.2 Comparison of effectiveness of directions

In order to estimate the adversarial direction efficiently in Algorithm ??, a key role is played by the vectors $\{\mathbf{z}^{(1)}\}_{i=1}^{m}$ which perturb the input. One particular choice being sampling the directions from a normal distribution. However, it is worth noting that the inverse covariance distribution of the gradients by construction is a convolution operator and hence is diagonalized by the FFT basis. Low frequency components of the FFT basis vectors enhanced the attack accuracy significantly, an example of which is depicted in Figure 7 for black-box attacks on VGG16-bn classifier for ImageNet with the ℓ_{∞} perturbation constrained to $\epsilon = 0.05$.

A.3 MNIST Experiments

The GMRF model used for MNIST is given by $\Lambda_{i,i} = \alpha, \Lambda_{i,i+1} = \Lambda_{i+1,i} = \Lambda_{i,i-1} = \Lambda_{i-1,i} = \beta, \Lambda_{i+1,i+1} = \Lambda_{i-1,i-1} = \Lambda_{i-1,i+1} = \Lambda_{i+1,i-1} = \gamma$, where $\Lambda_{i,j}$ denotes the (i, j)-th element of Λ . For estimating the GMRF parameters, we use the last 5 images of the MNIST test set and perturb each of them with 5 vectors drawn from a normal distribution. For the attack, we use low frequency basis vectors of the FFT basis. The following table gives the values of the different hyperparameters used in the attack. Except for the GMRF parameters, all the other parameters were determined using

δ	0.1
$\mid \alpha \mid$	21094408
β	-5116365
γ	284558.1562
σ	10^{-3}
$ \delta_1 $	0.15

Table	2:	MNIST	Experiment	Settings
raore		1111 110 1	Enperment	Dettingo

grid search.

A.4 VGG16 Experiments

The GMRF model used for VGG16 for Imagenet is given by $\Lambda_{0,i,i} = \alpha$, $\Lambda_{0,i,i+1} = \Lambda_{0,i+1,i} = \Lambda_{0,i,i-1} = \Lambda_{0,i-1,i} = \beta$, $\Lambda_{0,i+1,i+1} = \Lambda_{0,i-1,i+1} = \Lambda_{0,i-1,i-1} = \Lambda_{0,i+1,i-1} = \kappa$, $\Lambda_{1,i,i} = \Lambda_{-1,i,i} = \gamma$, where in $\Lambda_{k,i,j}$, k denotes the channel. We also tried out GMRF models of lower and higher degree of association and we selected the one performing the best. For estimating the GMRF parameters, we use the last 10 images of the ImageNet validation set and perturb each of them with 10 vectors drawn from a normal distribution. For the attack, we use low frequency cosine basis vectors of the FFT basis. The following table gives the values of the different hyperparameters used in the attack. Except for the GMRF parameters, all the other parameters were determined using grid search.

δ	1.0
$ \alpha $	633.44
β	-24.05
γ	-232.04
κ	-2.00
σ	1.0
δ_1	0.04

Table 3: ImageNet VGG-16 Experiment Settings

A.5 ResNet50 Experiments

The GMRF model used for VGG16 for Imagenet is given by $\Lambda_{0,i,i} = \alpha, \Lambda_{0,i,i+1} = \Lambda_{0,i+1,i} = \Lambda_{0,i,i-1} = \Lambda_{0,i-1,i} = \beta, \Lambda_{0,i+1,i+1} = \Lambda_{0,i-1,i+1} = \Lambda_{0,i-1,i-1} = \Lambda_{0,i+1,i-1} = \kappa, \Lambda_{0,i,i+2} = \Lambda_{0,i,i-2} = \Lambda_{0,i-2,i} = \Lambda_{0,i+2,i} = \Lambda_{0,i+1,i+2} = \Lambda_{0,i-1,i+2} = \Lambda_{0,i+2,i+1} = \Lambda_{0,i+2,i-1} = \Lambda_{0,i-2,i+1} = \lambda_{0,i-2,i+1} = \nu, \Lambda_{1,i,i} = \Lambda_{-1,i,i} = \gamma$, where in $\Lambda_{k,i,j}$, k denotes the channel. We also tried out GMRF models of lower and higher degree of association and we selected the one performing the best. For estimating the GMRF parameters, we use the last 10 images of the ImageNet validation set and perturb each of them with 10 vectors drawn from a normal distribution. For the attack, we use low frequency cosine basis vectors of the FFT basis. The following table gives the values of the different hyperparameters used in the attack. Except for the

Table 4:	ImageNet	ResNet50	Experiment	Settings

δ	1.0
$ \alpha $	2631.93
β	-263.33
$ \gamma $	-837.16
κ	6.78
ν	28.09
σ	0.5
δ_1	0.0375

GMRF parameters, all the other parameters were determined using grid search.

A.6 Inception v3 Experiments



Figure 8: Cosine Similarity MNIST: 40 queries

Figure 9: Cosine Similarity MNIST: 200 queries

denotes the channel. We also tried out GMRF models of lower and higher degree of association and we selected the one performing the best. For estimating the GMRF parameters, we use the last 10 images of the ImageNet validation set and perturb each of them with 10 vectors drawn from a normal distribution. For the attack, we use low frequency cosine basis vectors of the FFT basis. The following table gives the values of the different hyperparameters used in the attack. Except for the

δ	1.0
$\mid \alpha \mid$	8964.89
β	-2960.87
γ	-841.13
κ	1155.66
ν	286.03
σ	0.1
δ_1	0.045

Table 5: ImageNet Inception v3 Experiment Settings

GMRF parameters, all the other parameters were determined using grid search.

A.7 Gradient Estimation Performance

We explore the cosine similarities of the adversarial directions of the proposed GMRF based attack framework and white-box FGSM. For the 40 query budget setting, the cosine similarity is centered around 0.3, while for the 200 query budget setting, the cosine similarity is centered around 0.2 as depicted in figures 8 and 9 respectively. Had our gradient estimates completely aligned with that of the true gradient as in white-box FGSM our performance would have been upper bounded by the performance of white-box FGSM. In essence, our scheme finds directions for adversarial perturbations, which in itself does not maximize the loss but is able to find a direction which leads to misclassification of the examples and outperform white-box FGSM. We further illustrate the performance of our gradient estimation scheme through experiments on the ImageNet dataset using the VGG-16bn architecture. We use two metrics namely, mean squared error of the normalized estimated gradient with respect to the normalized true gradient and the cosine similarity between the estimated gradient and the true gradient. In order to perform our analysis, we use 500 data samples from the test set to estimate the gradient using the GMRF framework. First, we estimate the GMRF parameters as previously described in Algorithm 1 and then perform the MAP estimation for the gradient.

The gradient estimation performance for ImageNet using VGG16-bn is depicted in Figure 8. In order to perform our analysis, we sample 500 data samples from the ImageNet validation set to estimate the gradient using the GMRF framework. First, we estimate the GMRF parameters as previously described in Algorithm 1 and then perform the MAP estimation for the gradient. We specifically consider the query budget to be 200. Out of the 400 correctly classified images, white-box FGSM and our proposed algorithm attain attack accuracies of 0.9268 and 0.8794 respectively. While, the gradient estimation performance in terms of MSE is impressive, the cosine similarity shows that

the estimated gradient does not quite coincide directionally with the true gradient. The difference in the directions explains the inferior performance of our proposed scheme in this regime. It is worth noting that the dimension of the input data for VGG16-bn is 150528. From classical results in zeroth-order optimization it is well known that in a *d*-dimensional space, O(d) queries are required to obtain a nearly bias-free gradient estimate. Our framework uses only 200 queries to estimate the gradient which resides in a 150528 dimensional space. In spite of the possible erroneous directional characteristic of our estimated gradient, it still manages to achieve a 0.87 success rate and outperforms BANDITS-TD and PARSIMONIOUS in the 200 query budget regime.



Figure 10: VGG-16bn, ImageNet, 200 query budget

A.8 Autocorrelation

We provide further evidence for considering a non-identity covariance for modelling the gradient covariance. In figure 11, we plot the autocorrelation of the true gradients from 100 images sampled from the MNIST test set for kernel size of 11×11 for the LeNet model. In figure 12, we plot the autocorrelation of the third channel of the true gradients from 50 images sampled from the ImageNet validation set for kernel size of 11×11 for the VGG16-bn model. The autocorrelation plots show the correlation across dimensions and across images to be substantial so as to provide even more evidence and reason for the gradient model we have considered in this paper.



Figure 11: Autocorrelation of the gradients: MNIST, LeNet, 11×11 kernel



Figure 12: Autocorrelation of the gradients: ImageNet, VGG16-bn, 11×11 kernel