TEST-TIME ADAPTATION AND ADVERSARIAL ROBUSTNESS

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper studies test-time adaptation in the context of adversarial robustness. We first formulate an adversarial threat model for test-time adaptation, where the defender may have a unique advantage as the adversarial game becomes a maximin game, instead of a minimax game as in the classic adversarial robustness threat model. We then study whether the maximin threat model admits more "good solutions" than the minimax threat model, and is thus *strictly weaker*. On the positive side, we show that, if one is allowed to access the training data, then Domain Adversarial Neural Networks (DANN), an algorithm designed for unsupervised domain adaptation, can provide nontrivial robustness in the test-time maximin threat model against strong transfer attacks and adaptive fixed point attacks. This is somewhat surprising since DANN is not designed specifically for adversarial robustness (e.g., against norm-based attacks), and provides no robustness in the minimax model. On the negative side, we show that recent data-oblivious test-time adaptations, in contrast to DANN, can be easily attacked. We take a step to discuss moving towards adversarially robust test-time adaptation and examine its various implications.

1 INTRODUCTION

There is a surge of interest to study test-time adaptation to help generalization to unseen domains (e.g., recent work by Sun et al. (2020); Wang et al. (2020); Nado et al. (2020)). At the high level, a generic test-time adaptation can be modeled as an algorithm Γ which accepts an (optional) labeled training dataset D, an (optional) model F trained on D (usually used as a starting point), and an unlabeled test feature set U, outputs a model $\widetilde{F} = \Gamma(F, D, U)$, in order to achieve high test accuracy on U. For large test set U, test-time adaptation can be viewed as a form of *transductive learning* (Joachims (1999); Vapnik (1998)) (i.e., using D, U to train a model to predict specific instances in U), which is argued to be easier than more traditional inductive learning.

This paper studies test-time adaptation in the context of adversarial robustness (i.e., there is an active agent who tries to fool the test-time adaptation by perturbing the input so that $\widetilde{F}$ gives wrong predictions). There are several motivations in pursuing this direction. First, this question is of practical interest: Many practical ML pipelines run in a *batch mode*[1], where they first collect a set of unlabelled data points, and then send them to a model (e.g. Nado et al. (2020)). In such cases, data in the batch may have been adversarially perturbed, and it is a natural question *whether we can leverage the large batch size and test-time adaptation to enhance adversarial robustness.* Second, from a purely theoretical point of view, since test-time adaptation is a form of transductive learning, it is intriguing to ask *whether transductive adversarial learning can be easier*, given that traditional adversarial robustness is formulated in the *inductive* learning setting (e.g. Madry et al. (2018)). To this end, a recent work by Goldwasser et al. (2020) shows that, with transductive learning, one can achieve nontrivial guarantees for classes of bounded VC dimension with *arbitrary* train and test distributions. The current work complements their paper in the setting of deep learning.

To study these questions, we first formalize a threat model, which we call *(test-time) maximin threat model*, for the adversarial robustness of test-time adaptation. Recall that the classic adversarial robustness game is a minimax game $\min_F \mathbb{E}_V[\max_{\widetilde{V}} L(F, \widetilde{V})]$, where V is random sampled data,

[1]For example, Instagram collects a large batch of photos before sending them to a model to tag people.

$\widetilde{V}$ is the perturbed data generated from V by the adversary, and $L(F, \widetilde{V})$ is the loss of the model F on $\widetilde{V}$. By contrast, in the maximin threat model, we allow V to be sampled from a different domain, and the game is maximin: $\mathbb{E}_V[\max_U \min_{\widetilde{F}} L(\widetilde{F}, \widetilde{V})]$ (where U is the perturbed features of V, subject to the attack type, and $\widetilde{V}$ is the labeled perturbed data, see Definition 2). By the maximin inequality, it follows that this threat model is *no harder* than the minimax model (to allow source and target domains to be different, we need to generalize the classic minimax model, see Definition 3).

We then move on to the focus of this work: Whether the maximin threat model is "strictly weaker" than the minimax threat model. We note that any good defender solution (a robust model) in the minimax game induces a good defender solution in the maximin game (an adaptation algorithm that outputs that robust model), thus intuitively, the good defender solutions of the minimax model is a subset of the good defender solutions of the maximin threat model. We ask *whether such a containment is proper*: That is, whether there *exists* a defender solution that is good in the maximin threat model (i.e., provides robustness), but is bad in the minimax threat model (provides no robustness). The existence of such a defender will demonstrate that the maximin threat model admits more good solutions. Besides theoretical interest, this question is also of practical importance since these "new" solutions may possess *desirable properties that good solutions in the minimax threat model may lack*. For example, one such property is that the defender solution is *attack agnostic* (Goodfellow (2018) (pp.30)): That is, the solution is not to directly optimize the performance measure for a particular type of perturbation[2].

On the positive side, we consider ℓ_∞ attacks (more results on ℓ_2 attacks are reported in the appendix), and empirically show that DANN (Ganin et al. (2017)), an algorithm designed for *unsupervised domain adaptation* (UDA), is a candidate that exhibits such a separation between the maximin and minimax models. This is somewhat surprising as DANN is also agnostic to ℓ_∞ attacks. Specifically, we generate an adapted model $\widetilde{F} = \mathsf{DANN}(\mathsf{RI}, D, U)$ to classify U, where RI is random initialization, D is the labeled training set, and U is generated by either transfer attacks using Projected Gradient Descent (PGD) attacks (Madry et al. (2018)), or adaptive fixed point attacks. In the homogeneous case where the train and test distributions are the same, on CIFAR10 for example, this achieves about 50% test accuracy, comparable with robustness achieved by an adversarially trained CIFAR10 model (which is trained specifically for ℓ_∞ attacks) in the minimax model. In the inhomogeneous case, on CIFAR10→CIFAR10c-fog for example, DANN provides about 52% accuracy, significantly larger than the 17% accuracy using an adversarially trained CIFAR10 model (note that because the defender does not see clean labeled CIFAR10c-fog data, one cannot hope to use an adversarially trained CIFAR10c-fog model). Not surprisingly, the accuracy becomes very low when we evaluate $\widetilde{F}$ in the minimax model, confirming that it is specifically for the maximin model.

On the negative side, we explore the maximin robustness of the recent data-oblivious adaptation algorithms (namely, the adaptation algorithms do not use D, but just the pretrained model F and unlabeled test set U). Specifically, we consider Test-Time Training (TTT) by Sun et al. (2020)[3]. We show that TTT can be easily attacked using simple transfer attacks. While this is not surprising as authors of Sun et al. (2020) have cautioned that TTT is not designed for adversarial robustness, the situation is in sharp contrast to our results with DANN.

The rest of the paper is organized as follows: Section 2 presents necessary setup. In Section 3 we define the maximin threat model, and adversarial semi-supervised minimax threat model (a generalization of the classic minimax threat model). In Section 4 we show that DANN provides, somewhat surprisingly, adversarial robustness in the maximin model. Finally, Section 5 explores the maximin robustness of oblivious test-time adaptation, and concludes the paper with future directions.

2 PRELIMINARIES

Let F be a model, for a data point $(x, y) \in \mathcal{X} \times \mathcal{Y}$, a loss function $\ell(F; x, y)$ give the loss of F on x given the true label y. Let V be a set of labeled data points. We use the notation $L(F, V) =$

[2] Another consideration, which is beyond the scope of this paper, is the computational feasibility of finding a good solution, given the hardness of minimax optimization Katz et al. (2017); Daskalakis et al. (2020).

[3] While TTT does not use training data D at the test time, it has a special self-training component, and the joint architecture is a Y-structure. A more domain agnostic approach is discussed in Wang et al. (2020).

$\frac{1}{|V|}\sum_{(x,y)\in V}\ell(F;x,y)$ to denote the empirical loss of F on V. For example, if we use binary loss $\ell^{0,1}(F;x,y)=\mathbb{1}[F(x)\neq y]$, this gives the test error of F on V. We use the notation $V|_X$ to denote the projection of V to its features, that is $\{(x_i,y_i)\}_{i=1}^m \mapsto \{x_1,\ldots,x_m\}$.

Threat model for classic adversarial robustness. To formulate the threat model for test-time adaptation, we first present a threat model for the classic adversarial robustness. Although the classic adversarial robustness can be written down succinctly as a minimax objective, namely $\min_F \mathbb{E}_{(x,y)\sim(X,Y)}\left[\max_{x'\in N(x)}[\ell(F;x',y)]\right]$ ($N(x)$ is a neighborhood function of x, determined by the attack type), a threat model formulation will help us develop more nuanced models.

Definition 1 (Threat model for classic adversarial robustness). *Attacker and defender agree on a particular attack type. Attacker is an algorithm $\mathcal{A}$, and defender is a supervised learning algorithm $\mathcal{T}$.*

Before game starts
- *A (labeled) training set D is sampled i.i.d. from from (X,Y).*

Training time
- ***(Defender)** Train a model F on D as $F=\mathcal{T}(D)$.*

Test time
- *A (labeled) natural test set V is sampled i.i.d. from (X,Y).*
- ***(Attacker)** On input F, D, and V, $\mathcal{A}$ perturbs each point $(x,y)\in V$ to (x',y) (subject to the agreed attack type), giving $\widetilde{V}=\mathcal{A}(F,D,V)$.*

Evaluation:
Evaluate the test loss of F on $\widetilde{V}$, $L(F,\widetilde{V})$. Attacker's goal is to maximize the test loss, while the defender's goal is to minimize the test loss.

We stress that the i.i.d sampling of V is important (which is also present in the expectation in the minimax objective): Otherwise an attacker can pick any point that fools F and repeat it arbitrarily many times. (we refer readers to Goodfellow (2019) for more discussions along this line).

Notations for models and attacks. In this paper we mainly use the PGD attacks (Projected Gradient Descent attacks) with norm-based perturbations Madry et al. (2018). For example, given a model F, we use the notation $\mathsf{PGD}(F,V)$ to denote PGD attacks against F, on data V (the attack type is specified in the context). We adopt the following notations:

Notation	Meaning
T	A target model trained on the labeled target data V.
AdvT	An adversarially trained target model using the labeled target data V.
S	A source model trained on the labeled source data D.
AdvS	An adversarially trained source model using the labeled source data D.
$\mathsf{PGD}(\cdot,\cdot)$	PGD Attacks on a model and data. For example, $\mathsf{PGD}(\mathsf{AdvT},V)$ means to use PGD attacks on the model AdvT and data V.

Test-time defenses and BPDA. Various previous work have investigated *test-time defenses* where a pretrained model is fixed and there is a "preprocessing procedure" to preprocess an input before sending it to the model. Several such defenses were described and attacked in Athalye et al. (2018), by the BPDA technique (Backward Pass Differentiable Approximation). While syntactically one can fit test-time defenses into our framework, the test-time adaptation concept discussed in this paper is significantly different, which usually allows one to completely modify a model. As we will see later in the paper, we use an unsupervised domain adaptation algorithm, such as DANN, for test-time adaptation. In these cases, it is unclear how to apply BPDA.

3 Threat Models

3.1 Test-time maximin threat model

The intuition behind the test-time maximin threat model is as follows: After we receive the adversarially perturbed data U to classify (at test time), the defender trains a model based on U, and we evaluate the test accuracy only for U. (i.e., for different test set U we may have different models and

different test accuracy.) This perspective of training a model using labeled data D and unlabeled data U, and only test on U, is essentially the *transductive learning* by Vapnik (1998) (however, we consider it in an adversarial setting). Formally, we have the following definition:

Definition 2 (**Maximin threat model for test-time adaptation).** *Fix an adversarial perturbation type. We have a source domain (X^s, Y^s), and a target domain (X^t, Y^t). Attacker is an algorithm $\mathcal{A}$, and defender is a pair of algorithms $(\mathcal{T}, \Gamma)$, where $\mathcal{T}$ is a supervised learning algorithm, and Γ is a test-time adaptation algorithm.*

Before game starts
- *A (labeled) training set D is sampled i.i.d. from (X^s, Y^s).*

Training time
- **(Defender, optional)** *Train a model F trained on D as $F = \mathcal{T}(D)$.*

Test time
- *A (labeled) natural test set V is sampled i.i.d from (X^t, Y^t), V is sent to the attacker.*
- **(Attacker)** *Produce an unlabeled dataset U as follows:*
  1. *On input Γ, F, D, and V, $\mathcal{A}$ perturbs each point $(x, y) \in V$ to (x', y) (subject to the agreed attack type), giving $\widetilde{V} = \mathcal{A}(\Gamma, F, D, V)$.*
  2. *Send $U = \widetilde{V}|_X$ (that is, the feature vectors of $\widetilde{V}$) to the defender.*
- **(Defender)** *Produce an adapted model as $\widetilde{F} = \Gamma(F, D, U)$.*

Evaluation: *Evaluate the test loss of $\widetilde{F}$ on $\widetilde{V}$, $L(\widetilde{F}, \widetilde{V})$.*

On the adversary: *While the adversary knows the adaptation algorithm, Γ may use some private randomness, such as random initialization. Since defender proceeds after the attacker to apply Γ, these private randomness is assumed not to be known by the adversary.*

On notations: *If there is no need for a pretrained model, we use the notation $\Gamma(\perp, \cdot, \cdot)$.*

Remarks on the threat model. We note the following aspects:

Modeling differences. **(1)** The first obvious distinction is that *source domain may differ from the target domain.* **(2)** This is called a *test-time* threat model since $\widetilde{F}$ is trained based on $U = \widetilde{V}|_X$, and the test error is only evaluated on $\widetilde{V}$ (the perturbed test set). This is in sharp contrast with the classic minimax threat model, where the model F is trained on D, which is independent from V and so $\widetilde{V}$. **(3)** This threat model enables us to study whether *a large test set (i.e. large $|U|$) can benefit a defender for adversarial robustness*. This is in sharp contrast with the classic minimax threat model, where the attacker is granted much power, and can pick points in a one-by-one, or "online", fashion (i.e., sample a point from the nature, perturb it, and send to the defender). **(4)** This is called a *maximin* threat model since the adversary *must move first* to submit the challenge set U, and the defender then moves based on it. In fact, this can be written as a maximin game $\mathbb{E}_V \left[\max_U \min_{\widetilde{F}} \text{test_err}(\widetilde{F}, \widetilde{V}) \right]$, different from the minimax game in the classic threat model.

Same security goal. However, in the case where these two domains coincide, the maximin threat model is still nontrivial. In fact, from the security perspective, both threat models are about the *same* end-to-end security goal: *Correct predictions on adversarially perturbed data.*

Attacking a model vs. Attacking an algorithm. In the classic threat model, a model F is fixed after training, and at the test time the adversary attacks the model. However, in the maximin threat model, the adversary must attack an adaptation *algorithm* Γ.

Restrictions of the threat model. Astute readers may realize that we have intentionally left the definition of the test-time adaptation Γ to have a lot of freedom. For example, Γ can leverage the labeled training data D, which may be practically not available or computationally infeasible for test-time computations. We consider the following possibilities of restrictions:

Homogeneity. If the source and target domains equal, we call it the *homogeneous maximin model*. This setting is directly related to the classic threat model (Definition 1), where we have a single domain (X, Y), but an adversary can perturb the test instances. From a security perspective, maximin threat model gives a *relaxed* threat model, but *the same end-to-end security guarantees*. (i.e., correct predictions against adversarially perturbed input)

Data-obliviousness. In this case, the adaptation algorithm Γ cannot leverage the labeled source data D. In other words, the agent can only leverage the pretrained model F and the unlabeled data set U.

3.2 ADVERSARIAL SEMI-SUPERVISED MINIMAX THREAT MODEL

An unsatisfying aspect of the theory thus far is that the classic minimax threat model is only defined in the homogeneous case, but the maximin threat model can be inhomogeneous, which makes the direct comparison difficult. To bridge the gap, we provide a generalization of the classic minimax threat model that lies between the classic minimax threat model (Definition 1) and the test-time maximin threat model (Definition 2).

Definition 3 (**Adversarial semi-supervised minimax threat model).** *Fix an adversarial perturbation type. We have a source domain (X^s, Y^s), and a target domain (X^t, Y^t). The attacker is a pair of algorithms $(\mathcal{A}_0, \mathcal{A}_1)$, and the defender is a pair of algorithms $(\mathcal{T}, \Gamma)$, where $\mathcal{T}$ is a supervised learning algorithm, and Γ is a semi-supervised learning algorithm.*

Before game starts
- *A training set D is sampled i.i.d. from (X^s, Y^s).*
- *A semi-supervision set V is sampled i.i.d from (X^t, Y^t), V is sent to the attacker.*

Training time
- **(Defender, optional)** *Train a model F trained on D as $F = \mathcal{T}(D)$.*
- **(Attacker)** *Produce an unlabeled dataset U as follows:*
  1. *On input Γ, F, D, and V, $\mathcal{A}_0$ perturbs each point $(x, y) \in V$ to (x', y) (subject to the agreed attack type), giving $\widetilde{V} = \mathcal{A}_0(\Gamma, F, D, V)$.*
  2. *Send $U = \widetilde{V}|_X$ (that is, the feature vectors of $\widetilde{V}$) to the defender.*
- **(Defender)** *Produce a final model as $\Gamma : \widetilde{F} = \Gamma(F, D, U)$.*

Test time
- *A (labeled) natural test set V' is sampled i.i.d. from (X^t, Y^t).*
- **(Attacker)** *On input $\widetilde{F}$, D, and V', $\mathcal{A}_1$ perturbs each point $(x, y) \in V'$ to (x', y) (subject to the agreed attack type), giving $\widetilde{V'} = \mathcal{A}_1(\widetilde{F}, D, V')$.*

Evaluation: *Evaluate the test loss of $\widetilde{F}$ on $\widetilde{V'}$, $L(\widetilde{F}, \widetilde{V'})$. The goal of the adversary $(\mathcal{A}_0, \mathcal{A}_1)$ is to (jointly) maximize $L(\widetilde{F}, \widetilde{V'})$.*

Remarks on the threat model. We note the following aspects:

Modeling differences. **(1)** This threat model is *semi-supervised*, because the defender (learner) receives an unlabeled data set U from the target domain. Note that the procedure to produce $\widetilde{F}$ at the training time of this threat model is exactly the same as the procedure to produce $\widetilde{F}$ in the test-time maximin threat model. The key difference is that, once $\widetilde{F}$ is trained, we evaluate it on $\widetilde{V'}$, which is the adversarially perturbed data on *independently sampled* V'. This is thus inductive learning, and a minimax game. **(2)** This threat model is *adversarial*, because the attacker can adversarially perturb the clean semi-supervision set V to produce U.

Classic minimax model as a special case. The classic minimax threat model is a *special case* of this threat model, by putting source and target domains equal, and choosing a trivial Γ: $\Gamma(F, D, U) = F = \mathcal{T}(D)$. We list several other specializations of this threat model in Appendix B.3. Therefore, without loss of generality, for the rest of the paper, by *minimax model* we mean the *adversarial semi-supervised minimax threat model*.

Γ***: One algorithm, two interpretations.*** We note that for an algorithm $\Gamma : (F, D, U) \mapsto \widetilde{F}$, one can have now two interpretations: In the maximin threat model, we interpret it as a adaptation algorithm, because we are only going to evaluate $\widetilde{F}$ on $\widetilde{V}$; in the minimax threat model, we interpret it as a semi-supervised learning algorithm, because we are going to evaluate $\widetilde{F}$ on unseen points $\widetilde{V'}$.

4 MAXIMIN VS. MINIMAX: A DANN PERSPECTIVE

We now move on to the focus of this work:

Is the maximin threat model "strictly weaker" than the minimax threat model?

Theoretical understanding. We first show that the maximin threat model is no harder than the minimax threat model.

Proposition 1 (**Homogeneous maximin vs. Classic minimax threat model).** *Let $k \geq 1$ be a natural number, and $\mathcal{F}$ be the hypothesis class. For a given V, the domain of $\widetilde{V}$ is a well-defined function of V (e.g., ℓ_∞ ball around V). We have that:*

$$\mathop{\mathbb{E}}_{V\sim(X,Y)^k}\left[\max_{U}\min_{\widetilde{F}\in\mathcal{F}}\{L(\widetilde{F},\widetilde{V})\}\right] \leq \min_{\widetilde{F}\in\mathcal{F}}\mathop{\mathbb{E}}_{V\sim(X,Y)^k}\left[\max_{\widetilde{V}}\{L(\widetilde{F},\widetilde{V})\}\right]$$

Note that in the proof (A.1) the homogeneity condition is not essential (one can set (X, Y) as the target domain and the proof carries through). Further, the proof holds verbatim if we extend the family of models in the minimax game to $\mathcal{F}(D, U)$. This thus gives the following corollary,

Corollary 1. *Maximin threat model is no harder than the minimax threat model (Definition 3).*

We also note the following:

Proposition 2 (**Good minimax solution is also a good maximin solution).** *Suppose $\mathcal{T}^*$ is a supervised learning algorithm which trains a model $F^* = \mathcal{T}^*(D)$, where its adversarial gain in the adversarial semi-supervised minimax model is bounded by κ (i.e. $\mathbb{E}_{V'}[\max_{\widetilde{V}'} L(F^*, \widetilde{V})] \leq \kappa$.) Then in the maximin threat model, the adversarial gain of the strategy $(\mathcal{T}^*, \Gamma^*)$, where $\Gamma^*(F^*, D, U) = F^* = \mathcal{T}^*(D)$, is also upper bounded by κ.*

4.1 Are good minimax solutions strictly contained in good maximin solutions?

Proposition 1 and Proposition 2 lead to the following natural question:

Is there a Γ such that:
- *As a test-time adaptation algorithm in the maximin model, it provides robustness, but*
- *As a learning algorithm in the minimax model, the model it produces has no robustness.*

The existence of such a Γ provides a separation between the minimax and the maximin threat models, and indicates that the maximin threat model admits more good solutions. Despite theoretical interest, this question is of practical importance since these "new" solutions may possess *desirable properties that good solutions in the minimax threat model may lack*. In this section we consider one such desirable property that the defender solution is *attack agnostic* (Goodfellow (2018) (pp.30)): That is, The defender strategy is not to directly optimize the performance measure. (e.g., we know the attack is ℓ_∞-norm attacks, and we directly train for it).

High-level experiment design. While we are not able to prove the existence of a defender solution that separates maximin and minimax threat models, we design and implement experiments to provide evidence that such a strategy may actually exist. The experiment design is as follows:

(1) We use the DANN algorithm with random initialization (RI), as a test-time adaptation; that is, DANN(RI, $\cdot, \cdot$). There are several motivations of choosing DANN: **(A)** DANN fits our framework as it is designed for unsupervised domain adaptation. **(B)** DANN is however *not* designed for adversarial robustness. Thus is it a very interesting question whether DANN can provide test-time adversarial robustness against attacks (e.g. norm-based attacks) that *it is not specifically designed for*. **(C)** DANN algorithm leverages source data D, which could benefit the maximin robustness.

(2) We generate adversarially perturbed $\widetilde{V}$, and check whether DANN can provide robustness.

(3) Note that DANN(RI, $\cdot, \cdot$) can be viewed as a *semi-supervised learning algorithm in the adversarial semi-supervised minimax threat model*. Therefore, we check whether the adapted model $\widetilde{F}$ = DANN(RI, D, U) is robust in the minimax model.

(4) If (2) and (3) show that DANN(RI, $\cdot, \cdot$) is *significantly more robust* in the test-time maximin threat model than in the minimax model, then the experiment goal is achieved.

Algorithm 1 FIXED POINT ATTACK $\mathsf{FPA}[k, F]$

Require: A training dataset D, a natural test set V, a pretrained model F used by Γ if Γ relies on a pretrained model, and an integer parameter $k \geq 0$.
1: Attack F on V to get adversarially perturbed examples V_0. Set $U_0 = V_0|_X$.
2: Set $F_0 = \perp$ if Γ does not use a pretrained model, otherwise set $F_0 = F$.
3: **for** $i = 1, 2, \ldots, k$ **do**
4: Compute $F_i = \Gamma(F_0, D, U_{i-1})$.
5: Attack F_i on V to get adversarially perturbed V_i. Set $U_i = V_i|_X$.
6: **end for**
7: **return** U_k.

Attacks. We use ℓ_∞ attacks in this section (note that DANN is not designed for ℓ_∞ attacks). We report more results about ℓ_2 attacks in Appendix D.3. We consider two categories of attacks:

Transfer Attacks. Transfer attacks are a common class of attacks, where we transfer attacks on one model to the target model (in our context, produced by a maximin defender strategy). In this paper we will mainly apply PGD attacks on adversarially trained models to generate transfer attacks.

Fixed Point Attacks. We also consider how an adversary can leverage the knowledge of the adaptation algorithm Γ, and generate *adaptive* attacks. We propose a new family of adaptive attacks, which we call the *fixed point attacks (*FPA*)*. The intuition behind FPA is that since our goal is to find U on which $\Gamma(F, D, U)$ fails (thus U is a "fixed point"), we thus start with a U_0, and iteratively compute $F_i = \Gamma(F, D, U_{i-1})$, generate U_i by attacking F_i, and hope that with large number of iterations we can converge to a good $U = U_k$. Formally, FPA (Algorithm 1) accepts k, F as parameters, and instantiate U_0 by attacking F, and then repeatedly attacking the models produced by Γ. Note that for the case where Γ does not rely on pretraining, there are different ways to pick F. For example we can choose $F = \mathsf{AdvT}$, and then $U_0 = \mathsf{PGD}(\mathsf{AdvT}, V)|_X$.

Datasets. For the homogeneous case, we consider MNIST and CIFAR10 (i.e., for both source and target domains). The homogeneous case represents the security problem considered in the classic threat model. For the inhomogeneous case, we consider MNIST→MNIST-M (Ganin et al., 2017), and CIFAR10→CIFAR10c-fog (Hendrycks & Dietterich, 2019). For CIFAR10c-fog, all 5 levels combined and perform and 80-20 train test split, to obtain the training set of size 40000 and test dataset of size 10000 (for space reasons, we combine all corruption levels for experiments in this section, results where we separately study different levels are reported in Section D.2).

Models. For MNIST, we use the original DANN architecture from Ganin et al. (2017) with slight modifications (e.g. adding batch normalization and dropout layers). For CIFAR10, we use pre-activation ResNets from He et al. (2016) for the prediction branch of DANN, and for the domain prediction branch we use the architecture from Sun et al. (2020) (convolution layer). This architecture is slightly different from the the typical DANN architecture used for MNIST, which assumes vectorized feature and fully connected domain prediction branch.

Homogeneous ($S = T$)			**Inhomogeneous** ($S \neq T$)		
Data	**Defender**	**Accuracy**	**Data**	**Defender**	**Accuracy**
MNIST	DANN	**98.79%**	MNIST→MNIST-M	DANN	**62.73%**
MNIST	AdvS	98.14%	MNIST→MNIST-M	AdvS	35.95%
CIFAR10	DANN	**49.47%**	CIFAR10→CIFAR10c-fog	DANN	**52.28%**
CIFAR10	AdvS	46.39%	CIFAR10→CIFAR10c-fog	AdvS	17.22%

Table 1: Adversarial robustness under transfer attacks with ℓ_∞ attacks. For MNIST we use perturbation budget 0.3. For MNIST-M we use perturbation budget $8/255$ (we found robustness on MNIST-M degrades fast on MNIST-M, even for adversarially trained models. We choose this number to illustrate relative performance). For all CIFAR experiments we use perturbation budget $8/255$. In all cases, DANN demonstrated better accuracy (albeit in a weaker threat model).

Experiment results. (A) Table 1 summarizes the results under transfer attacks. In the homogeneous case, the adversarial accuracy DANN provided in the maximin model is comparable to the adversarial accuracy an adversarially trained model provided in the minimax model. And the adversarial accuracy becomes significantly higher in the inhomogeneous case (compared to using an adversar-

k	MNIST	MNIST→ MNIST-M	CIFAR10	CIFAR10→ CIFAR10c-fog
1	98.57%	56.37%	63.65%	67.99%
2	98.85%	55.62%	83.62%	56.06%
3	99.20%	55.54%	83.52%	50.12%
4	99.42%	52.87%	88.01%	49.52%
5	98.23%	52.31%	88.18%	55.09%
6	98.71%	53.49%	88.52%	57.25%
7	99.35%	52.57%	88.24%	55.85%
8	98.69%	52.16%	87.91%	54.89%

Table 2: Adversarial robustness under fixed point attacks with F = AdvT (i.e., the attacker initializes the initial model as the adversarially trained target model). DANN exhibits significant robustness even for large k.

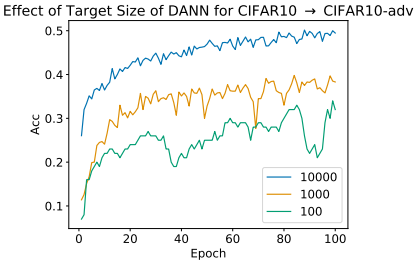

Figure 1: Effect of test size on CIFAR10. We plot the accuracy as we train more epochs with DANN during the adaptation, for different target size. Adversarial accuracy of smaller target size is consistently lower.

ially trained source model). **(B)** Table 2 summarizes the results under FPA attacks. Similar to the transfer attack case, DANN provides noticeable robustness. Note that since the defender moves after the attacker, he or she always applies adaptation "one more time". **(C)** Figure 1 gives the robustness results under different target size. As we can see, the robustness of DANN degrades as the target size decreases, confirming our intuition that large target size benefits test-time robustness. **(D)** Finally, we also evaluate the accuracy of the adapted DANN models in the minimax threat model. Not surprisingly, the accuracy becomes very low. This thus concludes our experiment that DANN provides robustness in the test-time maximin threat model, but not in minimax threat model.

5 Towards adversarially robust test-time adaptation

We briefly explore the adversarial robustness of the recent data-oblivious test-time adaptation algorithms. Specifically, we focus on the TTT algorithm by Sun et al. (2020). Recall that a test-time adaptation algorithm is data-oblivious, if it does not use the labeled training data D at the test time. TTT algorithm uses a specially trained (with self training) pretrained model, which we denote as PrTTT. Similar to our DANN experiment, we conduct experiments on CIFAR10→CIFAR10c-fog: **(1)** We found that attacking PrTTT can already break TTT algorithm. While this is not surprising as authors of Sun et al. (2020) have cautioned that TTT is not designed for adversarial robustness, this is in sharp contrast to our results with DANN. **(2)** We then use an adversarially trained model as the pretrained model for TTT. We found that this indeed increases the maximin-robustness of TTT. However, the robustness is roughly the same (about 1% difference) as directly using the pretrained model. This indicates that the robustness mainly comes from the adversarially trained model, and TTT algorithm provides no robustness. This is again in sharp contrast to the situation of DANN.

We list several implications as future directions:

• *Is transductive adversarial deep learning easier?* Both the current work and the work of Goldwasser et al. (2020) have given indication that adversarial robustness may be easier in the transductive learning setting. Can we provide more theoretical evidence, in particular, for deep learning?

• *The surprising power of distribution matching for robustness.* It is somewhat surprising that a distribution matching algorithm, such as DANN, can provide adversarial robustness (albeit in a weaker threat model), even when the target domain is already a corrupted domain such as CIFAR10c-fog. This means that for practical applications where we run an ML model in a batch mode on some collected unlabeled data, one can consider first apply UDA to produce an adapted model and then classify the batch, and this may provide adversarial robustness. To this end, can we certify such robustness? Or can we find better attacks?

• *Does robust test-time adaptation always need to access source data?* We note that while DANN provides maximin robustness, it needs to access the labeled training data D at test time, and while TTT does not need D, it provides no maximin robustness. Can we achieve the best of both of worlds and get adversarially robust *oblivious* test-time adaptation algorithms?

REFERENCES

Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, and Mario Marchand. Domain-adversarial neural networks. *stat*, 1050:15, 2014.

Anish Athalye, Nicholas Carlini, and David A. Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 274–283. PMLR, 2018.

Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.

John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman. Learning bounds for domain adaptation. In *Advances in neural information processing systems*, pp. 129–136, 2008.

Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C. Duchi, and Percy Liang. Unlabeled data improves adversarial robustness. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pp. 11190–11201, 2019.

Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. In *Advances in Neural Information Processing Systems*, pp. 3730–3739, 2017.

Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization, 2020.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. Domain-adversarial training of neural networks. In Gabriela Csurka (ed.), *Domain Adaptation in Computer Vision Applications*, Advances in Computer Vision and Pattern Recognition, pp. 189–209. Springer, 2017.

Shafi Goldwasser, Adam Tauman Kalai, Yael Tauman Kalai, and Omar Montasser. Beyond perturbations: Learning guarantees with arbitrary adversarial test examples. *CoRR*, abs/2007.05145, 2020. URL https://arxiv.org/abs/2007.05145.

Ian J. Goodfellow. Defense against the dark arts: An overview of adversarial example security research and future research directions. *CoRR*, abs/1806.04169, 2018. URL http://arxiv.org/abs/1806.04169.

Ian J. Goodfellow. A research agenda: Dynamic models to defend against correlated attacks. *CoRR*, abs/1903.06293, 2019. URL http://arxiv.org/abs/1903.06293.

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. *Proceedings of the International Conference on Learning Representations*, 2019.

Thorsten Joachims. Transductive inference for text classification using support vector machines. In Ivan Bratko and Saso Dzeroski (eds.), *Proceedings of the Sixteenth International Conference on Machine Learning (ICML 1999), Bled, Slovenia, June 27 - 30, 1999*, pp. 200–209. Morgan Kaufmann, 1999.

Guy Katz, Clark W. Barrett, David L. Dill, Kyle Julian, and Mykel J. Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In Rupak Majumdar and Viktor Kuncak (eds.), *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part I*, volume 10426 of *Lecture Notes in Computer Science*, pp. 97–117. Springer, 2017.

Daniel Kifer, Shai Ben-David, and Johannes Gehrke. Detecting change in data streams. In *VLDB*, volume 4, pp. 180–191. Toronto, Canada, 2004.

Mingsheng Long, Jianmin Wang, Guiguang Ding, Jiaguang Sun, and Philip S Yu. Transfer joint matching for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1410–1417, 2014.

Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. *arXiv preprint arXiv:1502.02791*, 2015.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.

Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Multiple source adaptation and the rényi divergence. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pp. 367–374, 2009.

Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(8):1979–1993, 2019.

Zachary Nado, Shreyas Padhy, D Sculley, Alexander D'Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020.

Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Rui Shu, Hung H. Bui, Hirokazu Narui, and Stefano Ermon. A DIRT-T approach to unsupervised domain adaptation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.

Aman Sinha, Hongseok Namkoong, and John C. Duchi. Certifying some distributional robustness with principled adversarial training. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.

Yu Sun, Xiaolong Wang, Liu Zhuang, John Miller, Moritz Hardt, and Alexei A. Efros. Test-time training with self-supervision for generalization under distribution shifts. In *ICML*, 2020.

Jonathan Uesato, Jean-Baptiste Alayrac, Po-Sen Huang, Alhussein Fawzi, Robert Stanforth, and Pushmeet Kohli. Are labels required for improving adversarial robustness? In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pp. 12192–12202, 2019.

Vladimir Vapnik. *Statistical learning theory*. Wiley, 1998. ISBN 978-0-471-03003-4.

Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno A. Olshausen, and Trevor Darrell. Fully test-time adaptation by entropy minimization. *CoRR*, abs/2006.10726, 2020.

Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. In *Advances in neural information processing systems*, pp. 8559–8570, 2018.

A PROOFS

A.1 PROOF OF PROPOSITION 1

Let $\mathcal{F}$ be the family of models $\widetilde{F}$ we can choose from. From the maximin inequality, we have that

$$\max_{U} \min_{\widetilde{F} \in \mathcal{F}} \{L(\widetilde{F}, \widetilde{V})\} \leq \min_{\widetilde{F} \in \mathcal{F}} \max_{\widetilde{V}} \{L(\widetilde{F}, \widetilde{V})\}$$

Note that for the minimax, the max over $\widetilde{V}$ is also constrained to perturb features (as we want to find adversarial examples). If we take expectation over V, we have then

$$\mathop{\mathbb{E}}_{V} \left[\max_{U} \min_{\widetilde{F} \in \mathcal{F}} \{L(\widetilde{F}, \widetilde{V})\} \right] \leq \mathop{\mathbb{E}}_{V} \left[\min_{\widetilde{F} \in \mathcal{F}} \max_{\widetilde{V}} \{L(\widetilde{F}, \widetilde{V})\} \right]$$

Note that

$$\mathop{\mathbb{E}}_{V} \left[\min_{\widetilde{F} \in \mathcal{F}} \max_{\widetilde{V}} \{L(\widetilde{F}, \widetilde{V})\} \right] \leq \min_{\widetilde{F} \in \mathcal{F}} \mathop{\mathbb{E}}_{V} \left[\max_{\widetilde{V}} \{L(\widetilde{F}, \widetilde{V})\} \right],$$

which completes the proof. □

A.2 PROOF OF PROPOSITION 2

In the maximin threat model, let us take the strategy $(\mathcal{T}^*, \Gamma^*)$ as defined. Therefore the maximin adversarial gain is

$$\mathop{\mathbb{E}}_{V} \left[\max_{\widetilde{V}} L(\Gamma^*(\mathcal{T}^*(D), D, \widetilde{V}|_X), \widetilde{V}) \right] = \mathop{\mathbb{E}}_{V} \left[\max_{\widetilde{V}} L(\mathcal{T}^*(D), \widetilde{V}) \right] = \mathop{\mathbb{E}}_{V} \left[\max_{\widetilde{V}} L(F^*, \widetilde{V}) \right]$$

which, by assumption, is bounded by κ. □

B MORE ON THREAT MODELS

B.1 ON THE MODELING

Why is the adversary constrained to attack each point in Definition 2 and Definition 3? This is set up as the rule of the threat model, because if not, an adversary can pick any point that fails the model and repeats it arbitrary number of times, which trivializes the model. In fact, such one-by-one attack is implicit in the original minimax objective, and is also the common practice of evaluating adversarial robustness (i.e., typical evaluation is to take a test set, attack the points one by one, and then evaluate adversarial accuracy). Our threat model formulation makes this hidden assumption explicit. We refer readers to Goodfellow (2019) for more discussion if one wants to break the i.i.d. sampling assumption.

Why in the minimax modeling, we do not assume $\mathcal{A}$ in Definition 1, or $\mathcal{A}_1$ in Definition 3 to know the training algorithms? For example, $\mathcal{T}$ is not in the input of $\mathcal{A}$ in Definition 1. This is the case because in the inner most expression $L(F, \widetilde{V})$ only depends on F, and F is fixed after the training. Therefore, bringing in the training algorithm gives no additional information (note that this is also the typical practice of evaluating adversarial robustness in the classic setting, where we just attack the model). By contrast, this is not the case for the maximin threat model because attacker moves first and the model is not fixed yet, so in that case, a strong white-box attacker may need to take into consideration the adaptation algorithm, not just the pretrained model.

Why do we need to have two algorithms $\mathcal{A}_0$, $\mathcal{A}_1$ in Definition 3? $\mathcal{A}_0$ attacks the semi-supervised learning algorithm, and thus has Γ as input. On the other hand, $\mathcal{A}_1$ attacks a model, so neither Γ nor $\mathcal{T}$ is in the input. Therefore, these two algorithms are different and are separated in the definition. Note that the joint goal of $\mathcal{A}_0, \mathcal{A}_1$ is to fail the final model on $\widetilde{V'}$, so $\mathcal{A}_0$ can be totally different from $\mathcal{A}_1$ in order to fail the training. In this respect, $\mathcal{A}_0$ is like data poisoning the unlabeled data set in the semi-supervised learning. Unlike data poisoning though, we have the constraint that it must attack each point in V.

There are two algorithms $\mathcal{A}_0, \mathcal{A}_1$ in Definition 3, wouldn't $\mathcal{A}_0$ help the defender because it leaks information about the target distribution. We have shown that the defender can choose to ignore the unlabeled data set, and this degenerates to the classic threat model in the homogeneous case. Otherwise, the semi-supervised learning setting may indeed help the defender, because the information about the target domain (or more unlabeled data if we are in the homogeneous case) is leaked to the defender. This is, however, exactly why in these cases that this threat model lies between the classic threat model and the maximin threat model. In the maximin threat model, the attacker leaks U to the defender, and the defender is only required to predict U. In the semi-supervised threat model, things becomes more challenging, because finally we want to evaluate the model on independent samples.

In the experiments, we essentially instantiated the semi-supervised threat model, for example, as $(\mathcal{A}_0 = \mathsf{FPA}, \mathcal{A}_1 = \mathsf{PGD})$. Is that too weak? A strong adversary of course may want to use a strong attack $\mathcal{A}_0$ to fail the training, and thus fail the final model, and choosing FPA may indeed be a weak choice. However, we note that our goal in Step (3) of the experiment design is to show that we can *break* the adapted model, and $(\mathsf{FPA}, \mathsf{PGD})$ already full-fills this purpose. As another note, in the maximin threat model, currently FPA is the most natural adaptive attack we can think of. Exploring stronger attacks is an intriguing future direction.

B.2 More extensions and restrictions for maximin threat model

Singleton. In the standard model, we allow U to be a batch which consists of many point. In the singleton restriction, U is restricted to have only one point (for example, the test-time training algorithm by Sun et al. (2020) works in this setting).

Definition 4 (Online version). *In the online maximin model, we assume that we will sequentially receive batches of test samples, one after another, and we allow Γ to leverage the previous samples received. We assume, for simplicity of this work however, that the batches are sampled independently from each other.*

For the online maximin model, data obliviousness means that the agent cannot directly leverage the previous unlabeled samples he or she received. However, the agent can save the previous model parameters and use that.

B.3 Specialization of the adversarial semi-supervised minimax threat model

Adversarial training considered in Miyato et al. (2019); Shu et al. (2018). The setting becomes a special case of our general minimax threat model by:

- Instantiating $\mathcal{A}_0$ as a trivial adversary that does not do any perturbation.That is the defender receives a set of clean set U of unlabeled features.
- The defender adaptation strategy does adversarial training on the unlabeled data, given by

$$\begin{aligned} &d(p(y|x), p(y|x + r^*, \theta)) \\ &\text{where } r^* = \operatorname*{argmax}_{r:\|r\|\leq\varepsilon} d(p(y|x), p(y|x + r, \theta)) \end{aligned}$$

Unsupervised Adversarial Training Uesato et al. (2019) This is similar to the above except that the source and target domains are equal.

C Related Work

Adversarial robustness. Adversarial robustness of deep learning has received significant attention in recent year. By far a standard approach is to do adversarial training directly for the attack type Madry et al. (2018); Sinha et al. (2018). Several recent works have studied adversarial training in a semi-supervised setting Carmon et al. (2019); Uesato et al. (2019). As we have discussed, these work can be viewed as instantiations in the adversarial semi-supervised minimax threat model. Our work can be viewed as one step further to study adversarial robustness in the transductive setting for deep learning.

Test-time adaptation and transductive learning. Transductive learning is a long line of research (for example, see the classic work of Transductive Support Vector Machine Joachims (1999)), where one attempts to leverage unlabeled data U to predict on specific instances of U Vapnik (1998). On the other hand, transductive learning has not received much attention in the deep learning setting. The recent proposals of test-time adaptation Sun et al. (2020); Wang et al. (2020); Nado et al. (2020) seems to be an reincarnation of this idea, with the additional twist of leveraging a pretrained model. Our work considers test-time adaptation in an adversarial deep learning setting, which to the best of our knowledge, is new.

Unsupervised Domain Adaptation. A classic approach for analyzing domain adaption is based on $\mathcal{H}$-divergence Kifer et al. (2004); Blitzer et al. (2008); Ben-David et al. (2010). That theoretical framework is the basis for a line of methods that uses adversarial training with neural networks to learn representations that are indistinguishable between source and target domain, in particular domain adversarial neural network (DANN) Ajakan et al. (2014); Ganin et al. (2017) and related techniques Pei et al. (2018); Zhao et al. (2018). Some other approach used different divergence notions, such as MMD Long et al. (2014; 2015), Wasserstein distance Courty et al. (2017); Shen et al. (2018), and Rényi divergence Mansour et al. (2009).

D MORE EXPERIMENTS

D.1 BASELINES: ACCURACY OF ADVERSARIALLY TRAINED MODELS.

Data	Attack distance	Training PGD steps ($\mathcal{T}$)	Test PGD steps ($\mathcal{A}$)	Nat. test acc	Adv. test acc
CIFAR10	8/255	7	20	78.31%	46.39%
CIFAR10-fog	8/255	7	20	71.39%	29.1%
MNIST	0.3	40	100	98.81%	98.14%
MNIST-M	8/255	40	100	96.50%	80.88%

Table 3: Accuracy of adversarially trained models in the minimax threat model. For the Nat. test acc column we do not use the adversary $\mathcal{A}$. Note that for the MNIST-M row it is the case that we directly do adversarial training on the clean labeled target data. This cannot be achieved in our threat model because defender only sees perturbed unlabeled target data. Nevertheless, this row gives the best adversarial accuracy one can hope for.

See Table 3 for the results.

D.2 TRANSFER ATTACKS ON DIFFERENT CORRUPTION LEVELS

Corruption level	AdvS on PGD(AdvT, CIFAR10c-fog)	AdvT on PGD(AdvT, CIFAR10c-fog)	DANN on PGD(AdvT, CIFAR10c-fog)	DANN on CIFAR10c-fog
1	37.99%	47.00%	70.73%	76.32%
2	23.07%	39.84%	63.05%	77.00%
2	13.72%	30.07%	59.01%	75.56%
4	8.45%	21.41%	52.93%	74.41%
5	3.37%	10.58%	40.58%	62.79%

Table 4: Transfer attacks (ℓ_∞ attack) with different corruption levels.

See Table 4 for the results.

D.3 TRANSFER ATTACKS WITH ℓ_2 ATTACKS

Homogeneous ($S = T$)			**Inhomogeneous** ($S \neq T$)		
Data	**Defender**	**Accuracy**	**Data**	**Defender**	**Accuracy**
MNIST	DANN	99.19%	MNIST→MNIST-M	DANN	68.37%
MNIST	AdvS	99.07%	MNIST→MNIST-M	AdvS	44.30%
CIFAR10	DANN	69.57%	CIFAR10→CIFAR10c-fog	DANN	69.98%
CIFAR10	AdvS	70.97%	CIFAR10→CIFAR10c-fog	AdvS	37.32%

Table 5: Transfer attacks (ℓ_2 attacks with $\epsilon = 80/255$) for all tasks.

See Table 5 for the results.

Plots of fixed point attacks See Figure 2 for the results.

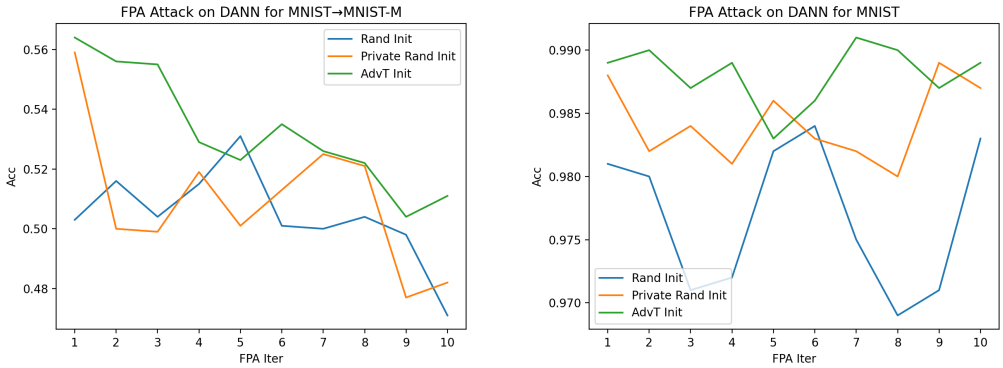

(a) MNIST, MNIST→MNIST-M.

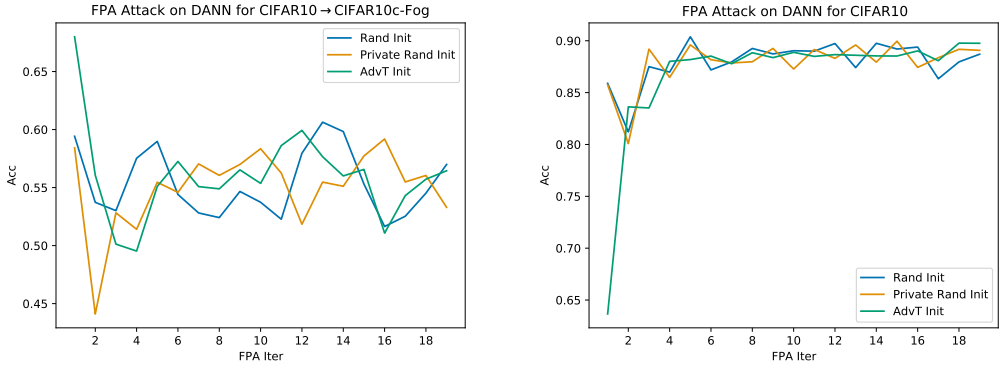

(b) CIFAR10, CIFAR10→CIFAR10-fog.

Figure 2: Plots for fixed point attacks.