

MIT Open Access Articles

Adversarial Training is Not Ready for Robot Learning

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Lechner, Mathias, Hasani, Ramin, Grosu, Radu, Rus, Daniela and Henzinger, Thomas A. 2021. "Adversarial Training is Not Ready for Robot Learning." 2021 IEEE International Conference on Robotics and Automation (ICRA).

As Published: 10.1109/ICRA48506.2021.9561036

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/144044>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Adversarial Training is Not Ready for Robot Learning

Mathias Lechner¹, Ramin Hasani², Radu Grosu³, Daniela Rus², Thomas A. Henzinger¹

Abstract—Adversarial training is an effective method to train deep learning models that are resilient to norm-bounded perturbations, with the cost of nominal performance drop. While adversarial training appears to enhance the robustness and safety of a deep model deployed in open-world decision-critical applications, counterintuitively, it induces undesired behaviors in robot learning settings. In this paper, we show theoretically and experimentally that neural controllers obtained via adversarial training are subjected to three types of defects, namely transient, systematic, and conditional errors. We first generalize adversarial training to a safety-domain optimization scheme allowing for more generic specifications. We then prove that such a learning process tends to cause certain error profiles. We support our theoretical results by a thorough experimental safety analysis in a robot-learning task. Our results suggest that adversarial training is not yet ready for robot learning.

I. INTRODUCTION

In this work, we discover that improving the safety of vision-based robot learning systems by adversarial training results in undesired side-effects of the robot’s real-world behavior. Training deep neural networks while accounting for adversarial examples robustifies the model to these visually imperceptible perturbations. This process trades nominal performance gained by standard empirical risk minimization (ERM) learning techniques, with worst-case performance under norm-bounded input perturbations [1], [2], [3].

Adversarial training has been mainly studied in image classification settings [4], [5], [6], [7], [8], [9], which exclusively focused on how much adversarial training trades nominal for robust test accuracy. While these metrics resemble the performance in static image classification tasks, robotic control tasks are inherently continuous and highly dynamic. Consequently, pure accuracy might not reflect the underlying performance of a robotic system accurately. For instance, for a closed control-loop, stability might be the highest priority, whereas faithfulness could be of high priority for vehicle routing algorithms.

In this work, we study how the nominal performance drop introduced by adversarial training methods is distributed over the real-world behavior in vision-based robot learning tasks. First, we propose safety-domain training, a generalization of adversarial training, which allows us to incorporate more general forms of safety specifications as secondary training objectives. We introduce a theoretical framework for characterizing error behaviors of learned controllers for robotic tasks. We then prove how safety-domain training changes the learned agent’s error-profile depending on the enforced safety

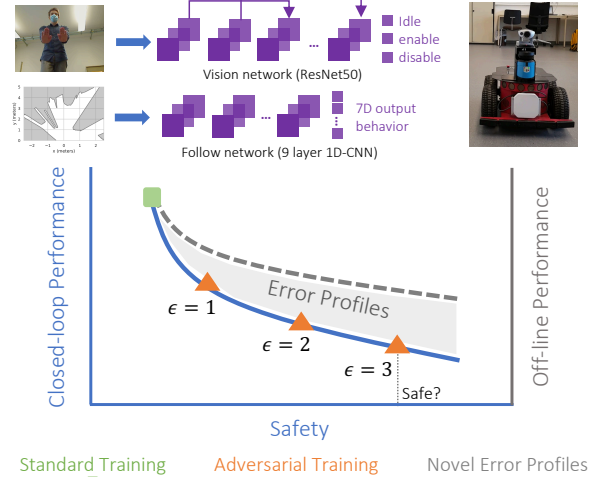


Fig. 1: Adversarial training gives rise to new error profiles that significantly decrease a neural controller’s performance.

specification. We test our training algorithm and confirm the consequences of our theory on an experimental case-study of an autonomous carrier robot with a variety of safety and robustness specifications. Based on our observations, we claim that:

Adversarial training is not yet ready for robot learning. More precisely, our experiments demonstrate that models trained by standard empirical risk minimization yield the best robotic performance in real-world scenarios. Counterintuitively, the best-performing agents are also vulnerable to making the robot crash under adversarial patterns. Conversely, while the models learned via our safety-domain training are provable immune to such worst-case behavior, they significantly perform worse in real-world scenarios. In particular, we observed that the strictness of the specifications enforced using adversarial training is the most dominant factor in determining the expected real-world robotic performance. Our empirical evaluations confirm that adversarial training of neural controllers requires rethinking before reliably using them in robot learning schemes [10], [11].

Summary of Contributions.

- 1) Formulate a generalization of adversarial training algorithms called safety-domain training with guarantees.
- 2) Theoretical framework consisting of erroneous behavior characterizations and induced error-profile by our safety-domain training algorithm.
- 3) Experimental confirmation of our framework on an image- and LiDAR-based autonomous carrier robot in various environments and assessment scenarios.

¹Institute of Science and Technology Austria (IST Austria)

³Massachusetts Institute of Technology (MIT), Cambridge, MA, USA

²Technische Universität Wien (TU Wien), 1040 Vienna, Austria

II. RELATED WORKS

Adversarial Training. Adversarial training has led to significant improvements of deep models’ resiliency to imperceptible perturbations. This was shown both empirically [2], [12], [13], [14] and with certification [15], [16], [17], [18], [19], [20], [21]. An emerging line of work suggests that the representations learned by adversarially trained models resemble visual features as perceived by humans more accurately compared to standard networks [22], [23], [24], [25], [26], [27]. In contrast, a large body of work tried to characterize the trade-off between a model’s robustness and accuracy when trained by adversarial learning schemes [28], [29], [30], [31]. Some gradient issues, such as gradient obfuscation [32], [33], during training, seemed to play a role in the mediocre performance of the models. Nevertheless, adversarially trained networks also showed to maintain their robustness properties [34] as well as their accuracy [35], [9] in transfer learning settings.

This work shows that despite the vast success of adversarially trained models in obtaining robustness properties on vision-based classification tasks, they can introduce novel error profiles in robot learning schemes. Our work aims to identify and report these profiles to enable the practical use of adversarial training in safety-critical applications.

Adversarial Training for Safe Robot Learning. Related approaches can be grouped into three categories; i) adversarial learning as a data augmentation technique; ii) Hand-crafted perturbation distribution; and iii) Task-specific models.

- (i) **Adversarial learning as a data augmentation technique** – A couple of recent works characterized generative adversarial networks (GANs) [36] as a data augmentation method to enhance neural controllers’ transferability. For example, [37] used GAN-based training for robotic visuomotor control, and [38] explored GANs to determine robust metric localization by using appearance transfer (e.g., day to night transformation of input images). These methods fundamentally differ from safety-related adversarial training frameworks [5] that we explore in this paper, as they refer to methods useful for data augmentation.
- (ii) **Hand-crafted perturbation distribution** – invariant sets, i.e., hand-crafted changes in underlying data distribution such as change of a gripper’s appearance and objects’ color used in task-relevant adversarial imitation learning [39]. These approaches require a simulator capable of generating domain-specific perturbations and are mainly designed for training performant agents. Our paper discusses a more general setting where we do not require domain-specific attributes.
- (iii) **Task-specific models** – Adversarial training in task-specific domains such as motion planning [40], [41], [42] and localization [43] has been used for enhancing robustness. Moreover, in reinforcement learning (RL) environments, adversarial training benefited agents in competitive scenarios such as active perception [44], interaction-aware multi-agent tracking, and behavior

prediction [45] and identifying weaknesses of a learned policy [46], [47]. These works do not evaluate existing general methods but propose tailored solutions for the specific task under-test. Our paper focuses on the broad vision-based robot learning problems that use contemporary adversarial training for enhancing safety.

III. ERROR PROFILES IN ROBOT LEARNING

Training a neural network f_θ in supervised learning, considers finding the best approximation function $f_\theta : x \mapsto y$ parameterized by θ , that maps the input data x to labels y . In a robotic learning setting, the data (x, y) is taken from a data distribution over a finite subset \mathbb{D} of the functional relation $\mathbb{R}^{N_x} \times \mathbb{R}^{N_y}$, i.e., sensor values and motor commands have finite precision. Limited training data, noise in the learning process [48], and inadequate causal modeling [49], [50], [51], [52] prevent the network from achieving a perfect mapping of the ground truth dependency between x and y . Consequently, these imperfections lead to errors during test time. In robot learning settings, we characterize these errors made by a neural controller by three categories: Systematic errors, transient errors, and conditional errors. Our objective is to show that these error profiles occur and lead to mediocre performance even when the neural controller is trained by safety-domain training methods such as adversarial training [5]. Let us first formally define these error profiles:

Definition 1 (Transient error): Given a loss function \mathcal{L} , a neural network f_θ , a threshold $\eta > 0$ and neighborhood $\epsilon > 0$, we call a point (x', y') transient error if

$$\mathcal{L}(y', f_\theta(x')) > \eta \text{ and } \mathcal{L}(\tilde{y}, f_\theta(\tilde{x})) < \eta, \quad (1)$$

for all (\tilde{x}, \tilde{y}) where $0 < \|\tilde{x} - x'\| \leq \epsilon$.

Definition 2 (Systematic error): Given a loss function \mathcal{L} , a neural network f_θ , a threshold $\eta > 0$, we say f_θ suffers from a systematic error if

$$\mathcal{L}(y, f_\theta(x)) > \eta, \quad (2)$$

for all $(x, y) \in \mathbb{D}$.

Definition 3 (Conditional error): Given a loss function \mathcal{L} , a neural network f_θ and a threshold $\eta > 0$, then we call a domain $\mathcal{D} = \{(x, y) | (x, y) \in \mathbb{D}\}$ a conditional error if

$$\mathbb{E}_{(x, y) \in \mathcal{D}} [\mathcal{L}(y, f_\theta(x))] > \eta, \quad (3)$$

and

$$\mathbb{E}_{(x, y) \in \mathcal{D}^c} [\mathcal{L}(y, f_\theta(x))] < \eta. \quad (4)$$

Fig. 2 illustrates these error types schematically for a single-dimensional learning problem. Transient errors might induce divergent behavior at the evaluation time (e.g., see the model’s approximation at the evaluation step x_{n-i}) in the second chart). Conditional errors can result in local mismatches that affect an agent’s local behavior to reach unsafe states. Systematic errors lead to distribution and baseline shifts for the entirety of the sample data. Next, we define a safety-domain robot learning scheme equipped with adversarial training and explore if this framework can give rise to the error types identified above.




Fig. 2: Different types of errors that can occur when fitting a network $f(x) = y$

IV. LEARNING WITH WORST-CASES IN MIND

This section defines a generalization of adversarial training by relaxing the ε -neighborhood for arbitrary domains. We call this approach *Safety-Domain Training*. We then explain how to solve the inner optimization loop of safety-domain training, either by empirical or certified safety methods and illustrate the resulting method in Algorithm 1.

Standard training of a neural network f_θ concerns optimizing θ to minimize the empirical risk as follows:

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(y_i, f_\theta(x_i)), \quad (5)$$

where \mathcal{L} is a loss function and $\{(x_i, y_i) | i = 1, 2, \dots, n\}$ the training samples. The optimization performed by mini-batch stochastic gradient descent (SGD) yields the best performing networks.

Adversarial training extends the objective in Eq. 5 by replacing each input x_i by an adversarial input \tilde{x}_i within the ε -neighborhood of x_i . Formally:

Definition 4 (Adversarial training [5]): Let f_θ be a neural network, $\{(x_i, y_i) | i = 1, 2, \dots, n\}$ the training samples, \mathcal{L} the loss function, and $\varepsilon > 0$ the adversarial attack radius. Then adversarial training optimizes the criterion

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \max_{\tilde{x}: \|x_i - \tilde{x}\| < \varepsilon} \mathcal{L}(y_i, f_\theta(\tilde{x})) \quad (6)$$

We generalize adversarial training to a more generic safety-domain training. In particular, we replace the ε -

neighborhoods of the training samples by arbitrary domains, i.e., labelled and connected sets.

Definition 5 (Safety-domain training): Let f_θ be a neural network, $\{(x_i, y_i) | i = 1, 2, \dots, n\}$ the training samples, \mathcal{L} the loss function, and $\{(D_i, z_i) | i = 1, 2, \dots, k\}$ the safety domains. Then safety-domain training optimizes the criterion

$$\min_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathcal{L}(y_i, f_\theta(x_i)) + \lambda \frac{1}{k} \sum_{i=1}^k \max_{\tilde{x} \in D_i} \mathcal{L}(z_i, f_\theta(\tilde{x})) \right], \quad (7)$$

where the hyperparameter λ specifies the tradeoff between optimizing the empirical training risk and the worst-case risk on the safety-domains.

Safety-domain training generalizes adversarial training and can be shown intuitively as follows:

Proposition 1: [Safety-domain training is a generalization of adversarial training] The criterion Eq (7) is a generalization of the adversarial training objective in Eq (6). In particular, this is true as we define $D_i := \{\tilde{x} : \|\tilde{x} - x_i\| < \varepsilon\}$ with $z_i := y_i$ and the training samples $\{y_i\}$.

Empirical vs. certified safety In practice, we have two options on how we solve the inner maximization step of safety-domain training. The first option is to perform several steps of projected gradient descent [2]. While this approach is computationally efficient and straightforward to implement, it provides no true worst-case guarantees as SGD does not ensure convergence to the global optimum. In practice, empirical approaches are often used for adversarial training of classifiers to account for computational complexity.

A more rigorous approach, albeit expensive, is to compute an upper-bound of the loss of each safety domain and minimize the upper bound via stochastic gradient descent [15], [53]. While computing an upper-bound of a network's output is difficult and may overestimate the true maximum, it provides certified guarantees on the worst-case loss. The interval bound propagation method falls into the category in which the upper bound of the loss is computed by an interval arithmetic abstraction of the network [54], [55]. The

Algorithm 1 Safety-domain training with guarantees

Input: Training data $\{(x_i, y_i) | i = 1 \dots n\}$, Safety domains $\{(z_i, D_i) | i = 1 \dots k\}$

Parameters: safety threshold δ , batch sizes b_t, b_s

Learning rate α , minimum training epochs i_{\min} .

safety_bound = ∞

while $i < i_{\min}$ **and** safety_bound $> \delta$ **do**

$i = i + 1$

$(\tilde{x}, \tilde{y}) = \text{sample_batch}(b_t, \{(x_i, y_i) | i = 1 \dots n\})$

$(\tilde{z}, \tilde{D}) = \text{sample_batch}(b_s, \{(z_i, D_i) | i = 1 \dots k\})$

$\nabla = \frac{\partial}{\partial \theta} \frac{1}{b_t} \sum_{i=1}^{b_t} \mathcal{L}(\tilde{y}_i, f_\theta(\tilde{x}_i))$

$\nabla = \nabla + \frac{\partial}{\partial \theta} \lambda \frac{1}{b_s} \sum_{i=1}^{b_s} \max_{x \in \tilde{D}_i} \mathcal{L}(\tilde{z}_i, f_\theta(x))$

$\theta = \theta - \alpha \nabla$

 safety_bound = $\max_{i=1 \dots k} \max_{x \in D_i} \mathcal{L}(z_i, f_\theta(x))$

end while

return θ

main difficulty of such certified approaches is to scale the training to large networks.

Algorithm 1 represents our framework for training an adversarially robust neural controller with safety guarantees.

V. ADVERSARIAL TRAINING NOT READY FOR ROBOT LEARNING

In this section, we show theoretically that adversarial training and even its generalization result in unexplored error profiles that lead to unsafe behavior. To construct our theory, we first describe a set of required assumptions.

A. Assumptions

To build the theory, we should rule out ill-posed edge-cases, thus relying on the following assumptions: 1) *Bounded generalization* The training loss provides a lower bound on the generalized loss [56], i.e., the expected loss over the data distribution. 2) *Bounded sets* Safety domains are bounded subsets of the data domain \mathbb{D} . 3) *Non-conflicting data* Safety-domains and ground truth data are non-conflicting, i.e., for every sample (x_i, y_i) and safety domain (D_j, z_j) it holds that if $x_i \in D_j \implies y_i = z_j$

Theorem 1: [Safety-domain training tend to cause conditional errors] Let f be a neural network, \mathcal{L} be a loss function, $\{(x_i, y_i) | i = 1, 2, \dots, m\}$ be the iid training data, and θ be the weights obtained by Algorithm 1 with the safety-domains $\{(D_i, z_i) | i = 1, 2, \dots, k\}$ and the threshold δ . Moreover, we assume δ to be a lower bound of the total training loss in Equation (7), i.e., ruling out cases where fitting f to the underlying data distribution is trivial. Let (x, y) be an arbitrary sample from the underlying data distribution, then

$$\mathbb{E}[\mathcal{L}(y, f_\theta(x)) | x \in \bigcup_{i=1}^k D_i] \leq \mathbb{E}[\mathcal{L}(y, f_\theta(x)) | x \notin \bigcup_{i=1}^k D_i] \quad (8)$$

Proof: For $x \in D_i$ we have

$$\begin{aligned} \mathcal{L}(y, f_\theta(x)) &\leq \max_{\tilde{x} \in D_i} \mathcal{L}(z_i, f_\theta(\tilde{x})) \\ &\leq \max_j \max_{\tilde{x} \in D_j} \mathcal{L}(z_j, f_\theta(\tilde{x})) \\ &= \delta. \end{aligned}$$

Thus

$$\mathbb{E}[\mathcal{L}(y, f_\theta(x)) | x \in \bigcup_{i=1}^k D_i] \leq \delta.$$

For $x \notin \bigcup_{i=1}^k D_i$, we know from our generalization bound

$$\begin{aligned} \mathbb{E}[\mathcal{L}(y, f_\theta(x)) | x \notin \bigcup_{i=1}^k D_i] &\geq \frac{1}{n} \sum_{i=1}^n \mathcal{L}(y_i, f_\theta(x_i)) \\ &\geq \delta, \end{aligned}$$

which shows the claim. \blacksquare

In the rest of this section, we discuss two important implications of Theorem 1 when adding additional assumptions about the networks and training process. In particular, we assume 1) *Locality* For any two training samples

$(x_1, y_1), (x_2, y_2)$ and a trained f_θ it holds that

$$\begin{aligned} \|x_1 - x_2\| > K &\implies \\ \mathbb{E}[\mathcal{L}(y_1, f_\theta(x_1))] &\approx \mathbb{E}[\mathcal{L}(y_1, f_\theta(x_1)) | \mathcal{L}(y_2, f_\theta(x_2))], \end{aligned}$$

i.e., the network's performance of a sample is not influenced by far apart samples. 2) *Training stability* For any training sample (x, y) and two trained networks $f_{\theta_1}, f_{\theta_2}$ it holds that

$$\mathbb{E}[\mathcal{L}(y, f_{\theta_1}(x))] \approx \mathbb{E}[\mathcal{L}(y, f_{\theta_2}(x))],$$

i.e., re-training a network does not introduce errors into the network.

Implication 1: The conditional errors introduced by the safety-domain training tends to occur near the boundaries of the safety-domains Let f_θ be a network obtained by safety-domain training and $f_{\theta'}$ be a network obtained by standard training. Moreover, let (x, y) be a sample with (x, y) with $\mathbb{E}[\mathcal{L}(y, f_\theta(x))] > \mathbb{E}[\mathcal{L}(y, f_{\theta'}(x))]$, then

$$\begin{aligned} (x, y) \in \{(x, y) | (x, y) \notin D_i \text{ and } \|x - \tilde{x}\| \geq K \\ \text{for all } (\tilde{x}, \tilde{y}) \in D_i \text{ and } i = 1, \dots, k\}, \end{aligned}$$

for some $K > 0$.

The claim that $(x, y) \notin D_i$ follows simply from applying Theorem 1. The second part, i.e., $\|x - \tilde{x}\| \geq K$, can be shown by deriving a contradiction when assuming the opposite is true. In particular, if $\|x - \tilde{x}\| < K$, our locality assumption implies that $\mathbb{E}[\mathcal{L}(y, f_\theta(x))]$ is independent from any $(\tilde{x}, \tilde{y}) \in D_i$. Due to this independence and our second training stability assumption, we know that the expected loss at (x, y) does not change when we retrain the network without any safety-domains. However, as safety-domain training with empty safety-domains corresponds to standard training, this contradicts our assumption that $\mathbb{E}[\mathcal{L}(y, f_\theta(x))] > \mathbb{E}[\mathcal{L}(y, f_{\theta'}(x))]$.

Proposition 2: [Transient and systematic errors are special cases of conditional errors] Let (x', y') be a transient error according to definition 1, then it implies a special case of a conditional error with $\mathcal{D} = (x', y')$ and $\mathcal{D}^c = (x, y) | (x, y) \in \mathbb{D} : 0 < \|\tilde{x} - x'\| = \epsilon\}$, i.e., a domain with only a single element and whose underlying data domain is defined locally. Moreover, if f_θ suffers from a systematic error as defined in 2, then it implies a special case of a conditional error with $\mathcal{D} = \mathbb{D}$, i.e., conditional error domain is the entire data domain.

Proposition 2 shows that adversarial training, i.e., when the safety-domains are small and sampled across the entire data domain, can lead to transient and systematic errors.

VI. EXPERIMENTAL EVALUATION: A CASE-STUDY

The objective of our experimental evaluation is to I) validate our claims empirically and II) study the different error-types on a more fine-grained level than the theory allows. In particular, we will study an end-to-end robotic control problem by learning from demonstration with deep models. Besides optimizing for high accuracy, we enforce secondary robustness and safety specifications on the networks. We impose the specification with various strictness




Fig. 3: State machine of the high-level controller. Transitions between states are triggered by a ResNet50 image classifier. In the active state the second neural network translates the LiDAR inputs to motor outputs.

levels to study how the specifications affect the network’s error characteristics on a fine-grained scale.

In our case study, we develop a controller for a vision + Lidar-based mobile robot. A human operator enables and disables the mobile robot via visual gestures. Once activated, the robot navigates such that it always faces the human operator at a distance of roughly one meter. The control software consists of two neural networks and a state-machine with two states. State transitions are triggered by a neural network (the vision network) processing the camera inputs. Fig. 3 shows an illustration of the state-machine and its transition profiles. The robot’s active behavior is realized via a second neural network (the follow network) that continuously translates a 2D-LiDAR scan of the environment into motor commands. The correct behavior is entirely determined by the networks’ performance, making the controller well suited for our empirical study. A video demonstration of the controller in action can be found at <https://youtu.be/xrgnSh1mk38>.

Our implementation approach is in contrast to traditional approaches for implementing such a robotic controller, which rely on hand-designed rules applied to infrared sensors [57] camera inputs [58], [59], or local localization protocols [60]. Perhaps the closest work to ours is the setup described in [61], which uses a stereo camera setup and machine learning using Support Vector Machines (SVMs).

Our physical robot is equipped with a Sick LMS1xx 2D-LiDAR rangefinder, a Logitech RGB camera, and a 4-wheeled differential drive. Consequently, our application allows for an additional layer of safety compared to pure vision-based approaches.

For our vision network we use a ResNet50 [62] pre-trained on ImageNet [63]. The fine-tuning task concerns classifying 1825 training images in three categories, i.e., idle, enable, disable gestures, as illustrated in Fig. 3. We avoid overfitting of the network by fine-tuning only the last few layers. We train the vision network by adversarial training with the fast-gradient-sign method [6] and three different values for ϵ , i.e., l_∞ neighborhoods with $\epsilon \in \{0, 1, 2\}$, see Fig. 6 for an example. Note that adversarial training with $\epsilon = 0$ is equivalent to a standard empirical risk minimization training. The training and validation accuracy is reported in Table I.

Our command-following network is a 9-layer 1D convolutional neural network mapping the 541-dimensional inputs to




Fig. 4: **Left:** Visualization of a safety-domain. No LiDAR signal in green area should be classified as a “forward” decision. **Right:** The network trained with standard ERM can be attacked to output a “forward” despite the LiDAR signal indicating a large object 10cm in front of the robot.

7 possible output categories, i.e., stay, forward, left forward, right forward, backward, left backward, and right backward. As the Follow network directly controls the motors, it potentially crashes into the human operator or an obstacle causing physical damage. To avoid such worst-case outcomes, we enforced safety-specifications on the network. In particular, we want to avoid the forward movement of the robot in case an object is in front of it. In Table II, we define four levels of safety-domains that characterize our safety requirements with increasing strictness.

Safety level 0 corresponds to standard empirical risk minimization. The other safety levels are enforced via safety-domain training using interval bound propagation (IBP), i.e., certified safety compared to the empirical safety of the vision network. A visualization of a safety-domain from the level 1 specification is shown in Fig. 4 on the left. We collected a total of 2705 training and 570 validation samples, as illustrated in Fig. 5. the training and validation accuracy is shown in Table III.

For both the vision and follow network, we evaluate each specification in seven standardized scenarios. The scenarios differ in complexity, e.g., operator commands, obstacles, environment, and lighting conditions. For the vision network, we report the number of misinterpreted gestures, i.e., an enabling or disabling of the controller without the operator’s command. Consequently, some types of errors are masked out (for instance, “enabling when the controller is already enabled”).

We report a holistic metric for the Follow network if the robot maneuvered correctly for the entire scenario. The results for the adversarially fine-tuned vision network are

TABLE I: Training and validation accuracy of the vision network trained with different adversarial perturbation radii. Training accuracy represents adversarial accuracy and validation accuracy represents clean accuracy.

Level	Training acc. (adversarial)	Validation acc. (clean)
0	99.7 %	98.4%
1	52.0%	92.8%
2	32.5%	71.9%

TABLE II: Specification of the safety domains D_i for the different safety levels.

Level	Description of safety domains D_i
0	$D_i = \emptyset$
1	$D_i = \left\{ x \mid 0 \leq x_j \leq 0.2 \text{ for } j \in \{i-1, i, i+1\} \text{ and } 0 \leq x_j \leq 3 \text{ for } j \notin \{i-1, i, i+1\} \right\}$
2	$D_i = \left\{ x \mid 0 \leq x_i \leq 0.2 \text{ and } 0 \leq x_j \leq 3 \text{ for } j \neq i \right\}$
3	$D_i = \left\{ x \mid 0 \leq x_i \leq 0.2 \text{ and } 0 \leq x_j \leq 4 \text{ for } j \neq i \right\}$

Setup Details: For each level there are 240 domains, i.e., $i = 150 \dots 390$. The corresponding labels z_i are defined as any non-forward moving category, i.e., $z_i \in \{\text{stay, backward, left backward, right backward}\}$. The domains with increase safety level represent super-set of the lower safety level, e.g. the conditions considered at level 1 are a strict subset of the level 2 safety. Level 1 safety only considers cases where at least three consecutive LiDAR rays are less than 20 cm, whereas one ray is enough for level 2 and 3. Level 3 differs from level 2 in terms of the upper bounds on the other rays.




Fig. 5: Three training samples of the active-mode following network. The network has to detect feet patterns in laser-rangefinder scans from seven different categories.

shown in Table V. While the network trained with $\epsilon = 1$ performed as well as the model trained by standard ERM, the performance significantly dropped when increasing the adversarial attack budget. Given that for human observers, adversarial perturbations with $\epsilon = 2$ are imperceptible, our results indicate that current training methods are unable to enforce non-trivial adversarial robustness on an image classifier in a robotic learning context.

Moreover, the results confirm that the errors occurred uniformly across different scenarios. Therefore, we conclude that imposing adversarial robustness specification causes transient errors confirming our theory.

The results for the safety-domain imposed follow network is reported in Table IV. Only the network trained with standard ERM could successfully handle all scenarios. Interestingly, Fig. 4 (right image) shows that this network is highly vulnerable to adversarial misclassifications and would output

TABLE III: Training and validation accuracy of the follow network trained when enforcing different safety-levels.

Level	Training accuracy	Validation accuracy
0	98.8%	84.7%
1	99.7%	76.8%
2	97.1%	73.4%
3	57.3%	53.2%




Fig. 6: Visualization of an adversarial attack on our vision network. Adding an adversarial mask flips the decision from a stop command to an activation command. The adversarial mask changes each pixel by only ± 2 steps on the 8-bit color channels, making the attack imperceptible to humans.

TABLE IV: Evaluation of the LiDAR follow networks

#	Scenario Description	Standard training	Safety Level 1	Safety Level 2	Safety Level 3
1	Plain	✓	✓	✓	Fail
2	Around boxes	✓	Fail	Fail	Fail
3	Out of corner	✓	✓	✓	Fail
4	Through gate	✓	✓	Fail	Fail
5	Around table	✓	✓	✓	Fail
6	Garage parking	✓	✓	✓	Fail
7	Narrow hallway	✓	Fail	Fail	Fail
Total		7/7	5/7	4/7	0/7

Note: Evaluation of the LiDAR follow networks with various safety specification enforced on seven standardized test scenarios. Successful navigation of a scenario is marked by a ✓. Fail indicates unsuccessful tests.

a forward decision if the large object is directly in front of the robot. While the networks with safety-level one and above are immune to such attacks, they perform significantly worse on the seven test-scenarios. With increasing specification level, the performance monotonously decreases until the network trained with the most rigorous safety specification cannot handle any scenario at all.

In contrast to the adversarial experiment of the Vision network, the defects made by the certified networks are conditional errors. In particular, if a network with specification level 1 could not solve a scenario, then a network with 2 and 3 could not either. Moreover, the failure of the level 1 and level 2 networks happened only during forward locomotion, i.e., close to the border of the safety-domains. This observation also supports our claims that errors induced by safety-domain training occur conditionally.

VII. CONCLUSION

Adversarial training and its generalization, safety-domain training can, in principle, learn robust and safe deep learning models. However, in this work, we showed that these methods induce unexplored error profiles in robotic tasks. We proposed a framework for characterizing these errors and predicting their presence based on the type of safety specification enforced during the training.

We empirically validated our claims and demonstrated that the type and strictness of the enforced specification govern the real-world performance of learned controllers for robotic

TABLE V: Evaluation of the image recognition networks.

#	Scenario description	Adversarial training radius		
		$\epsilon = 0$	$\epsilon = 1$	$\epsilon = 2$
1	Forward-backward	1	-	5
2	With surgical mask	-	1	3
3	Against direct sunlight	-	-	1
4	Staying idle	1	-	4
5	Summon out of garage	-	1	1
6	Artificial lighting (idle)	-	-	3
7	Artificial lighting (follow)	-	-	1
Total		2	2	18

Note: Evaluation of the image recognition networks trained with and without adversarial training on seven standardized test scenarios. Numbers indicate number of misclassified gestures that triggered a change in operation mode, i.e., errors without an effect are not counted. Dash represent zero misinterpretations.

environments. Our results concluded that adversarial training requires rethinking before being deployed in robot learning [64].

ACKNOWLEDGMENT

M.L. and T.A.H. are supported in part by the Austrian Science Fund (FWF) under grant Z211-N23 (Wittgenstein Award). R.H. and D.R. are partially supported by Boeing. R.G. are partially supported by Horizon-2020 ECSEL Project grant no. 783163 (iDev40).

REFERENCES

[1] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial machine learning at scale,” *arXiv preprint arXiv:1611.01236*, 2016.

[2] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*, 2017.

[3] C. Xie, Y. Wu, L. v. d. Maaten, A. L. Yuille, and K. He, “Feature denoising for improving adversarial robustness,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 501–509.

[4] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šmrdić, P. Laskov, G. Giacinto, and F. Roli, “Evasion attacks against machine learning at test time,” in *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 2013, pp. 387–402.

[5] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” *arXiv preprint arXiv:1312.6199*, 2013.

[6] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*, 2014.

[7] N. Carlini and D. Wagner, “Adversarial examples are not easily detected: Bypassing ten detection methods,” in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, 2017, pp. 3–14.

[8] D. Stutz, M. Hein, and B. Schiele, “Disentangling adversarial robustness and generalization,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6976–6987.

[9] H. Salman, A. Ilyas, L. Engstrom, A. Kapoor, and A. Madry, “Do adversarially robust imagenet models transfer better?” *arXiv preprint arXiv:2007.08489*, 2020.

[10] M. Lechner, R. Hasani, D. Rus, and R. Grosu, “Gershgorin loss stabilizes the recurrent neural network compartment of an end-to-end robot learning scheme,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 5446–5452.

[11] A. Brunnbauer, L. Berducci, A. Brandstätter, M. Lechner, R. Hasani, D. Rus, and R. Grosu, “Model-based versus model-free deep reinforcement learning for autonomous racing cars,” *arXiv preprint arXiv:2103.04909*, 2021.

[12] T. Miyato, S.-i. Maeda, M. Koyama, and S. Ishii, “Virtual adversarial training: a regularization method for supervised and semi-supervised learning,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 8, pp. 1979–1993, 2018.

[13] Y. Balaji, T. Goldstein, and J. Hoffman, “Instance adaptive adversarial training: Improved accuracy tradeoffs in neural nets,” *arXiv preprint arXiv:1910.08051*, 2019.

[14] H. Zhang, Y. Yu, J. Jiao, E. Xing, L. El Ghaoui, and M. Jordan, “Theoretically principled trade-off between robustness and accuracy,” in *International Conference on Machine Learning*, 2019, pp. 7472–7482.

[15] M. Lecuyer, V. Atlidakis, R. Geambasu, D. Hsu, and S. Jana, “Certified robustness to adversarial examples with differential privacy,” in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 656–672.

[16] L. Weng, H. Zhang, H. Chen, Z. Song, C.-J. Hsieh, L. Daniel, D. Boning, and I. Dhillon, “Towards fast computation of certified robustness for relu networks,” in *International Conference on Machine Learning*, 2018, pp. 5276–5285.

[17] E. Wong and Z. Kolter, “Provable defenses against adversarial examples via the convex outer adversarial polytope,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 5286–5295.

[18] A. Raghunathan, J. Steinhardt, and P. Liang, “Certified defenses against adversarial examples,” in *International Conference on Learning Representations*, 2018.

[19] J. Cohen, E. Rosenfeld, and Z. Kolter, “Certified adversarial robustness via randomized smoothing,” in *International Conference on Machine Learning*, 2019, pp. 1310–1320.

[20] H. Salman, J. Li, I. Razenshteyn, P. Zhang, H. Zhang, S. Bubeck, and G. Yang, “Provably robust deep learning via adversarially trained smoothed classifiers,” in *Advances in Neural Information Processing Systems*, 2019, pp. 11 292–11 303.

[21] G. Yang, T. Duan, E. Hu, H. Salman, I. Razenshteyn, and J. Li, “Randomized smoothing of all shapes and sizes,” *arXiv preprint arXiv:2002.08118*, 2020.

[22] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, “Adversarial examples are not bugs, they are features,” in *Advances in Neural Information Processing Systems*, 2019, pp. 125–136.

[23] L. Engstrom, A. Ilyas, S. Santurkar, D. Tsipras, B. Tran, and A. Madry, “Learning perceptually-aligned representations via adversarial robustness,” *arXiv preprint arXiv:1906.00945*, vol. 2, no. 3, p. 5, 2019.

[24] S. Santurkar, A. Ilyas, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, “Image synthesis with a single (robust) classifier,” in *Advances in Neural Information Processing Systems*, 2019, pp. 1262–1273.

[25] Z. Allen-Zhu and Y. Li, “Feature purification: How adversarial training performs robust deep learning,” *arXiv preprint arXiv:2005.10190*, 2020.

[26] B. Kim, J. Seo, and T. Jeon, “Bridging adversarial robustness and gradient interpretability,” *arXiv preprint arXiv:1903.11626*, 2019.

[27] S. Kaur, J. Cohen, and Z. C. Lipton, “Are perceptually-aligned gradients a general property of robust classifiers?” *arXiv preprint arXiv:1910.08640*, 2019.

[28] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, “Robustness may be at odds with accuracy,” in *International Conference on Learning Representations*, 2018.

[29] S. Bubeck, Y. T. Lee, E. Price, and I. Razenshteyn, “Adversarial examples from computational constraints,” in *International Conference on Machine Learning*, 2019, pp. 831–840.

[30] D. Su, H. Zhang, H. Chen, J. Yi, P.-Y. Chen, and Y. Gao, “Is robustness the cost of accuracy?—a comprehensive study on the robustness of 18 deep image classification models,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 631–648.

[31] A. Raghunathan, S. M. Xie, F. Yang, J. C. Duchi, and P. Liang, “Adversarial training can hurt generalization,” *arXiv preprint arXiv:1906.06032*, 2019.

[32] A. Athalye, N. Carlini, and D. Wagner, “Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples,” *arXiv preprint arXiv:1802.00420*, 2018.

[33] J. Uesato, B. O’Donoghue, A. v. d. Oord, and P. Kohli, “Adversarial risk and the dangers of evaluating against weak attacks,” *arXiv preprint arXiv:1802.05666*, 2018.

[34] A. Shafahi, P. Saadatpanah, C. Zhu, A. Ghiasi, C. Studer, D. Jacobs, and T. Goldstein, “Adversarially robust transfer learning,” in *International Conference on Learning Representations*, 2019.

- [35] F. Utrera, E. Kravitz, N. B. Erichson, R. Khanna, and M. W. Mahoney, "Adversarially-trained deep nets transfer better," *arXiv preprint arXiv:2007.05869*, 2020.
- [36] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [37] X. Chen, A. Ghadrizadeh, M. Björkman, and P. Jensfelt, "Adversarial feature training for generalizable robotic visuomotor control," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1142–1148.
- [38] H. Porav, W. Maddern, and P. Newman, "Adversarial training for adverse conditions: Robust metric localisation using appearance transfer," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1011–1018.
- [39] K. Zolna, S. Reed, A. Novikov, S. G. Colmenarej, D. Budden, S. Cabi, M. Denil, N. de Freitas, and Z. Wang, "Task-relevant adversarial imitation learning," *arXiv preprint arXiv:1910.01077*, 2019.
- [40] L. Janson, T. Hu, and M. Pavone, "Safe motion planning in unknown environments: Optimality benchmarks and tractable policies," *arXiv preprint arXiv:1804.05804*, 2018.
- [41] G. Shi, L. Zhou, and P. Tokekar, "Robust multiple-path orienteering problem: Securing against adversarial attacks," *arXiv preprint arXiv:2003.13896*, 2020.
- [42] C. Innes and S. Ramamoorthy, "Elaborating on learned demonstrations with temporal logic specifications," *arXiv preprint arXiv:2002.00784*, 2020.
- [43] Y. Yang and G. Huang, "Map-based localization under adversarial attacks," in *Robotics Research*. Springer, 2020, pp. 775–790.
- [44] M. Shen and J. P. How, "Active perception in adversarial scenarios using maximum entropy deep reinforcement learning," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3384–3390.
- [45] J. Li, H. Ma, and M. Tomizuka, "Interaction-aware multi-agent tracking and probabilistic behavior prediction via adversarial learning," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6658–6664.
- [46] X. Pan, D. Seita, Y. Gao, and J. Canny, "Risk averse robust adversarial reinforcement learning," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8522–8528.
- [47] S. Kuutti, S. Fallah, and R. Bowden, "Training adversarial agents to exploit weaknesses in deep control policies," *arXiv preprint arXiv:2002.12078*, 2020.
- [48] R. Hasani, G. Wang, and R. Grosu, "A machine learning suite for machine components' health-monitoring," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 9472–9477.
- [49] B. Schölkopf, "Causality for machine learning," *arXiv preprint arXiv:1911.10500*, 2019.
- [50] R. Hasani, M. Lechner, A. Amini, D. Rus, and R. Grosu, "Liquid time-constant networks," *arXiv preprint arXiv:2006.04439*, 2020.
- [51] M. Lechner, R. Hasani, A. Amini, T. A. Henzinger, D. Rus, and R. Grosu, "Neural circuit policies enabling auditable autonomy," *Nature Machine Intelligence*, vol. 2, no. 10, pp. 642–652, 2020.
- [52] M. Lechner and R. Hasani, "Learning long-term dependencies in irregularly-sampled time series," *arXiv preprint arXiv:2006.04418*, 2020.
- [53] S. Gruenbacher, R. Hasani, M. Lechner, J. Cyranka, S. A. Smolka, and R. Grosu, "On the verification of neural odes with stochastic guarantees," *arXiv preprint arXiv:2012.08863*, 2020.
- [54] S. Goyal, K. Dvijotham, R. Stanforth, R. Bunel, C. Qin, J. Uesato, R. Arandjelovic, T. Mann, and P. Kohli, "On the effectiveness of interval bound propagation for training verifiably robust models," *arXiv preprint arXiv:1810.12715*, 2018.
- [55] P.-S. Huang, R. Stanforth, J. Welbl, C. Dyer, D. Yogatama, S. Goyal, K. Dvijotham, and P. Kohli, "Achieving verified robustness to symbol substitutions via interval bound propagation," *arXiv preprint arXiv:1909.01492*, 2019.
- [56] O. Bousquet, U. Luxburg, and G. Rätsch, "Advanced lectures on machine learning, ml summer schools 2003," 01 2004.
- [57] S. Afghani and M. I. Javed, "Follow me robot using infrared beacons," *Academic Research International*, vol. 4, no. 3, p. 63, 2013.
- [58] T. Yoshimi, M. Nishiyama, T. Sonoura, H. Nakamoto, S. Tokura, H. Sato, F. Ozaki, N. Matsuhira, and H. Mizoguchi, "Development of a person following robot with vision based target detection," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 5286–5291.
- [59] M. Munaro, F. Basso, and E. Menegatti, "Tracking people within groups with rgb-d data," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 2101–2107.
- [60] B. V. Pradeep, E. Rahul, and R. R. Bhavani, "Follow me robot using bluetooth-based position estimation," in *2017 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*. IEEE, 2017, pp. 584–589.
- [61] J. Satake and J. Miura, "Robust stereo-based person detection and tracking for a person following robot," in *ICRA Workshop on People Detection and Tracking*, 2009, pp. 1–10.
- [62] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [63] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., "Imagenet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [64] M. Lechner, R. Hasani, M. Zimmer, T. A. Henzinger, and R. Grosu, "Designing worm-inspired neural networks for interpretable robotic control," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 87–94.