

# Practical Solutions for Machine Learning Safety in Autonomous Vehicles

Sina Mohseni,<sup>1,2</sup> Mandar Pitale,<sup>2</sup> Vasu Singh,<sup>3</sup> Zhangyang Wang<sup>1</sup>

<sup>1</sup> Texas A&M University, College Station, TX

<sup>2</sup> NVIDIA, Santa Clara, CA

<sup>3</sup> NVIDIA, Munich, Germany

{sina.mohseni, atlaswang}@tamu.edu, {mpitale, vasus}@nvidia.com

## Abstract

Autonomous vehicles rely on machine learning to solve challenging tasks in perception and motion planning. However, automotive software safety standards have not fully evolved to address the challenges of machine learning safety such as interpretability, verification, and performance limitations. In this paper, we review and organize practical *machine learning safety* techniques that can complement *engineering safety* for machine learning based software in autonomous vehicles. Our organization maps safety strategies to state-of-the-art machine learning techniques in order to enhance dependability and safety of machine learning algorithms. We also discuss security limitations and user experience aspects of machine learning components in autonomous vehicles.

## 1 Introduction

Advances in machine learning (ML) have been one of the biggest innovations of the last decade. Nowadays, ML models are used extensively in different industrial fields like autonomous vehicles, medical diagnosis, and robotics to perform various tasks such as speech recognition, object detection, and motion planning. Among different ML models, Deep Neural Networks (DNNs) (LeCun, Bengio, and Hinton 2015) are well-known and widely used for their powerful representation learning in high-dimensional data. For instance, in the field of autonomous driving, various DNN object detection and image segmentation algorithms have been used as perception units to process camera (e.g., PilotNet (Bojarski et al. 2016), Fast RCNN (Wang, Shrivastava, and Gupta 2017)) and Lidar (e.g., VoxelNet (Zhou and Tuzel 2018)) data.

The development of safety critical systems relies on stringent safety methodologies, designs, and analyses to prevent hazards at the time of failure. In the automotive field, ISO26262 and ISO/PAS 21448 are two main safety standards used to address safety of electrical and electronic components. These standards mandate methodologies for system, hardware, and software development. Specifically for software development, the process ensures traceability across requirements, architectural and unit design, code, and verification. In cases of autonomous software with high complexity, iterative hazard analysis and risk assessment of

autonomous software is required to formally represent the operational design domain.

On the other hand, ML models have numerous inherent safety drawbacks including accuracy on their training set and robustness in operational domain in open-world limitations. For example, ML models are fragile to domain shift (Ganin and Lempitsky 2014), data corruption and natural perturbations (Hendrycks and Dietterich 2019). Also, prediction probability scores in DNNs do not provide a true representation of model uncertainty. Moreover, from a security perspective, it has been shown that DNNs are susceptible to adversarial attacks that make small perturbations to the input sample (indistinguishable by human eye) but can fool a DNN (Goodfellow, Shlens, and Szegedy 2014). Due to the lack of verification techniques for DNNs, validation of ML models often relies on simple accuracy measures on different large test sets to cover the targeted operation design domain. While an important metric to gauge the success of the algorithm, it is definitely insufficient to measure performance in safety-critical applications as the real-world examples may differ from the test set.

With the realization that ML models will be increasingly used in safety critical systems, we need to investigate gaps these models expose in existing engineering safety standards. Several examples of these gaps have been discussed in machine learning safety including interpretability and traceability into code, formal verification, and design specification (Salay, Queiroz, and Czarnecki 2017).

In this paper, we review challenges and opportunities in algorithmic techniques for ML safety to complement existing software safety standards for autonomous vehicles. Section 2 briefly reviews the two main automotive industry safety standards and identifies their five fundamental gaps with machine learning algorithms. Section 3 describes practical algorithmic safety techniques by surveying machine learning research on 1) error detectors and 2) model robustness. We briefly present three of our own implementations for safety-critical applications. Section 4 discusses open challenges, directions for future work, and concludes the paper.

## 2 Background

Industrial safety broadly refers to the management of all operations and events within an industry in order to protect its

employees and users by minimizing hazards, risks, and accidents. Given the importance of correct operation of Electric and Electronic (E/E) components, the IEC 61508 is a basic functional safety standard developed for electrical and electronic safety related systems with two fundamental principles: (a) Safety life cycle: an engineering process based on best practices to discover and eliminate design errors, (b) Failure analysis: a probabilistic approach to account for the safety impact of system failures. IEC 61508 has several derivations specific to different industries. For example, to achieve safety in the automotive industry, engineers are required to follow the ISO 26262 standard to minimize safety risks due to E/E faults to an acceptable level.

In this section, we briefly review the two main automotive safety standards followed by an organized list of fundamental ML limitations to meet safety requirements. We note that both automotive standards mandate a meticulous analysis of hazards and risks, followed by a detailed development process focusing on system requirements, documented architecture and design, well structured code, and a thorough verification strategy for unit, integration, and system-level testing.

### ISO 26262 Standard

ISO 26262 or the automotive E/E functional safety standard defines vehicle safety as the absence of unreasonable risk due to malfunctioning E/E components. It requires a Hazard Analysis and Risk Assessment (HARA) to determine vehicle level hazards. The potential hazards and risks guide the safety engineers towards safety goals that are then used to create functional safety requirements. These safety requirements then guide the system development process which is then decomposed into hardware and software development processes. Figure 1 presents an overview of this standard. Our focus in this paper is Part 6 of ISO 26262 that defines the V-model for the software development process.

The objective of the V-Model in Figure 1 is to ensure that software safety requirements are sufficiently fulfilled by the software architectural design and sufficiently tested by the software verification tests. Similarly, the software architectural design is verified and the software integration tests prove that the interactions between the architectural entities including the static and dynamic aspects are tested. At the lowest level of the V-model, the unit design specifies the design details of each unit (that is identified during software architectural design) such as input, output error handling, behavior of the unit so that it can be coded. Finally, the unit tests ensure that the unit is tested so that its requirements and design aspects are met and sufficient structural coverage of the unit is met.

ISO 26262 also identifies the methods for fault detection and avoidance to minimize the risk to an acceptable level. However, when applied to autonomous driving, ISO 26262 has several limitations. It cannot account for faults occurring due to the inability of the component to comprehend the environment - for example due to performance limitations or robustness issues - or foreseeable misuse of the system.

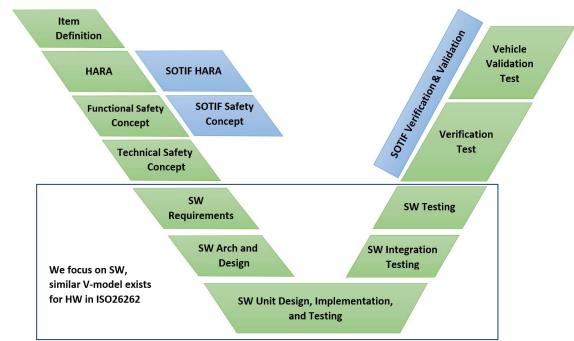


Figure 1: Comparison between the V-model in ISO 26262-6 (in Green) and ISO/PAS 21448 (in Blue) for software (SW) safety process.

### ISO/PAS 21448 Standard

ISO/PAS 21448 or SOTIF (Safety Of The Intended Functionality) standard describes an iterative development process consisting of design specification, development, and verification and validation phases. SOTIF standard recognizes performance limitations of software (including ML components) and expects that the scenarios/inputs that belong to unsafe-unknown (e.g., samples out of training distribution) and unsafe-known (e.g., samples out of operational design domain) situations shall be reduced to the extent such that the residual error risk is acceptable.

The SOTIF process augments the ISO 26262 with a SOTIF specific HARA and safety concept. The hazard analysis consists of identifying the hazards due to inadequate performance functionality, insufficient situational awareness, reasonably foreseeable misuse or shortcomings of the human machine interface. In comparison, the hazard and risk analysis of ISO 26262 is limited to hazards due to E/E failures. If the SOTIF analysis results in risk that is higher than acceptable, functional modifications are performed to reduce SOTIF risks. A verification and validation strategy is then developed to argue that the residual risk is below the acceptable level.

### Safety Gaps for ML Components

In recent years, there has been an increasing attention in identifying safety limitations of ML models. For example, Varshney et al. (Varshney 2016) discussed the definition of safety for ML models and compared it with four main *engineering safety* strategies (1) inherently safe design, (2) safety reserve, (3) safe fail, and (4) procedural safeguards) in industry. In a review of automotive software safety methods, Salay et al. (Salay, Queiroz, and Czarnecki 2017) presented an analysis of ISO-26262 part-6 methods with respect to safety of ML models. Their assessment of applicability of the software safety methods on machine learning algorithms (as software unit design) shows about 40% of software safety methods do not apply to ML models.

Safety and robustness of ML models problem has also been in the interest of ML scientists (Hernández-Orallo et al. 2019). Amodei et al. (Amodei et al. 2016) presented five

concrete research problems that could result in unintended and unsafe behavior of real-world AI systems. They focus and characterize safety problems of AI into defining and evaluating the objective functions. Also, Ortega et al. (Ortega and Maini 2018) introduced three areas of technical AI safety as specification (in design and emergent), robustness (for error prevention and recovery), and assurance (for monitoring and enforcement).

We categorize these open challenges in the terminology of automotive software safety standards and briefly review representative research (see first row of Table 1) to close these safety gaps.

- *Design Specification.* Documenting and reviewing the software specification is a crucial step in functional safety, however, the design specification of ML models is generally not adequate, as the models learn the patterns in data to discriminate or generate their distribution for new unseen input. Therefore, machine learning algorithms learn the target classes through their training data (and regularization constraints) and not using a formal specification. The lack of specificity could cause mismatch between “designer objectives” and “what the model actually learned” which could result in unintended functionality of the system. This data-driven optimization of variables to train machine learning models makes it impossible to define and pose specific safety constraints. To this end, Seshia et al. (Seshia et al. 2018) surveyed the landscape of formal specification for DNNs to lay an initial foundation for formalizing and reasoning about properties of DNNs. Another practical way to manage this design specification problem is to break machine learning components into smaller algorithms (with smaller tasks) to work in hierarchical structures. Related to this, Dreossi et al. (Dreossi et al. 2019) presented VerifAI toolkit for formal design and analysis of AI-based systems.
- *Implementation Transparency.* ISO26262 requires traceability from requirements to design. However, advanced ML models trained on high dimensional data are not transparent. The very large number of variables in the models makes them incomprehensible or so-called black-box for design review and inspection. In order to achieve traceability, significant research has been performed on interpretability methods for DNN to provide instance explanations of model prediction and DNN intermediate feature layers (Zeiler and Fergus 2014). In autonomous vehicles application, using the VisualBackProp technique (Bojarski et al. 2018) shows that a DNN algorithm trained to control a steering wheel would in fact learn patterns of lane, road edges, and parked vehicles to execute the targeted task. However, the completeness of interpretability methods to grant traceability is not proven yet (Adebayo et al. 2018) and in practice interpretability techniques are mainly used by designers to improve network structure and training process rather than support a safety assessment.
- *Testing and Verification.* Significant verification of work products are required for unit testing to meet the ISO 26262 standard. For example, coding guidelines for soft-

ware safety enforce that there is no dead or unreachable code. Depending upon the safety integrity level, complete statement, branch coverage, or modified condition and decision coverage are required to confirm adequacy of the unit tests. Coming to DNNs, formally verifying their correctness is challenging (provably NP-hard (Seshia, Sadigh, and Sastry 2016; 2016)) due to the high dimensionality of the data. Therefore, reaching a complete validation and testing bounded to the operational design domain is difficult. As a result, researchers proposed new techniques such as searching for unknown-unknowns (Bansal and Weld 2018) and predictor-verifier training (Dvijotham et al. 2018). Other techniques including neuron coverage and fuzz testing (Wang et al. 2018) in neural networks covers these aspects. Note that formal verification of shallow and linear models for low dimensional sensor data does not have the challenges of DNN verification.

- *Performance and Robustness.* SOTIF standard treats the ML models as a black box and suggests using methods to improve model performance and robustness. However, improving model performance and robustness in itself is a very challenging task. In learning problems, typically training a model ends with an error rate (due to false positive and false negative predictions) on the training set. *Empirical error* is prediction error rate of the learner function on its target distribution. *Generalization error* refers to the gap between model’s empirical error on its training and test sets. On top of these, *operational error* is referred to model’s error rate on open-world deployment that could be higher than test set error rate. Domain generalization refers to ability of model in learning generalizable data representations for open-world tasks. We will review more details and machine learning techniques to improve model robustness in Section 3.
- *Run-time Monitoring Function.* SOTIF and ISO 26262 standards suggest run-time monitoring functions as software error detection solutions. Monitoring functions in classical software are based on a rule-set to detect cases like transient hardware error, software crash, and exit from operation design domain. However, designing monitoring functions to predict ML failure (e.g., false positive and false negative error) is different in nature. ML models generate prediction probability for input instances but research shows prediction probability does not guarantee failure prediction (Hein, Andriushchenko, and Bitterwolf 2019). In fact, DNN and many other ML models could generate incorrect outputs with high confidence in cases of distribution shift and adversarial attacks. We will review more details of error detection techniques for ML models in Section 3.

### 3 Techniques for ML Safety

We introduce algorithmic techniques for ML safety to maintain dependability of machine learning algorithms for safe execution in open-world tasks. The techniques we review in this section are meant to complement classic engineering strategies in software safety. We also connect ML safety

techniques with appropriate engineering safety strategies (see Table 1) to help machine learning scientists and safety engineers find a common ground in the new topic of autonomous vehicle safety. We follow Varshney’s (Varshney 2016) four strategies to achieve machine learning safety to map AI Safety techniques with engineering strategies.

Considering the research debt left to assure different safety aspects of AI (see Section 2), we believe we are far from reaching (1) *Inherently Safe AI*. Therefore, we focus on practical machine learning solutions for the two following safety strategies:

- (2) *Safe Fail* refers to strategies to keep the vehicle on road in a safe state at the time of failure. This strategy can mitigate hazards at the time of fault by use of monitoring functions and graceful degradation plans such as notifying the driver to take vehicle’s control. We propose using run-time *error detection* techniques to detect erroneous output of machine learning algorithms (e.g. misclassification and misdetection) for the vehicle on the road.
- (3) *Safety Margins* in the context of machine learning is described as difference between model’s performance on training set and operational performance in open-world. We propose using *model robustness* techniques to improve resilience and hence Safety Margin of machine learning components.

We will also briefly review the importance of (4) *Procedural Safeguards* for non-experts end-users of autonomous vehicles (i.e. driver and passengers) later in the discussion about future work. To separate safety and security concerns, we consider external factors that intentionally exploit system vulnerabilities (e.g., sample manipulations in adversarial attack) as security concerns but not safety.

## Monitoring Function

Our first practical *ML safety* solution leverages from the range of techniques for machine learning misclassification error detection to achieve *Safe Fail* behavior. For example, when transient errors in hardware like sensors affect the functionality of the software like that for cruise control, an error detection unit (monitoring function) can detect the error and degrade the system by appropriate warnings and allowing the driver to take over. Similarly, various run-time monitoring functions and error detectors could be designed for machine learning components to predict model failure and trigger appropriate warnings. In the following we introduce three types of error detectors for machine learning and review their relations and limitations. Note that although the three following groups of error detectors overlap, we separate detectors by their targeted error types.

**Uncertainty Estimation** Uncertainty in probabilistic learners is an important factor to maintain fail-safety of the system. Even well trained and calibrated predictors that are robust to noise, corruption, and perturbations can benefit from uncertainty estimation to detect domain shift and out-of-distribution samples at run-time. Quantifying uncertainty can explain what a model does not know in terms of model confidence on its prediction (epistemic uncertainty or model

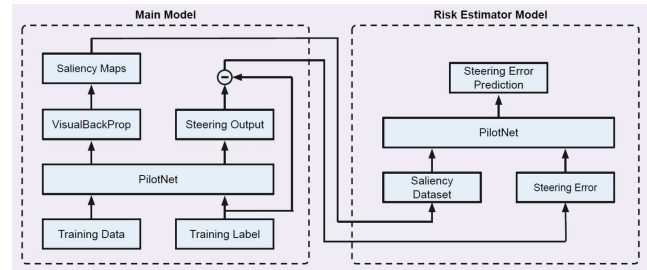


Figure 2: Model failure prediction based on saliency map for ConvNet regression models. We trained a student model as monitoring function for run-time failure predictor of a PilotNet model (Mohseni, Jagadeesh, and Wang 2019).

uncertainty) and uncertainty for unknown samples (aleatoric uncertainty or data uncertainty).

Regarding importance of uncertainty methods for safety critical applications, McAllister et al. (McAllister et al. 2017) proposed measuring uncertainties in ML models and propagating down in the decision-making pipeline as a key for safety of autonomous systems. However, quantifying predictive uncertainty in DNNs is a challenging task. Typically, DNN classification models generate normalised prediction scores which tend to be overconfident and regression DNN models do not give uncertainty representation in their output. Research in DNN presents solutions such as deep ensembles (Lakshminarayanan, Pritzel, and Blundell 2017) and Monte Carlo dropout (MC-dropout) (Gal and Ghahramani 2016) to estimate prediction uncertainty. Uncertainty estimation methods have been tested for various model error types (Kendall and Gal 2017) including adversarial attack detection (Smith and Gal 2018).

Although uncertainty estimation methods offer potential effective solutions for DNN failure prediction – in practice – they carry significant computation cost and latency which is not ideal for run-time failure prediction. For example, to design an error detector for a PilotNet algorithm, Michelmore et al. (Michelmore, Kwiatkowska, and Gal 2018) present an implementation of MC-dropout uncertainty estimation which needs 128 stochastic forward passes to estimate uncertainty of the model. Therefore, in resource limited setups and for reasons like computation simplicity, researchers work on alternative error detection solutions which we review in the next two subsection.

**In-distribution Error Detectors** Misclassification of in domain samples often happens due to weak representation learning. In recent years, advanced neural networks, regularization techniques, and large training datasets significantly improved DNNs representation learning and therefore model performance and robustness. However, run-time prediction error detectors are still needed to maintain safety of the system in case of model failure. Selective classification (also known as classification with reject option) is a technique to cautiously provide prediction for high-confidence samples and abstain when in doubt. This method for confident prediction can significantly improve model




Figure 3: Out-of-distribution sample detection base on prediction confidence. We proposed an out-of-distribution sample detector to train reject classes to learn outlier features in a self-supervised step. (Mohseni et al. 2020)

performance at the cost of test coverage. Geifman and El-Yaniv (Geifman and El-Yaniv 2017) present a simple and effective implementation of selective classification for DNNs. They introduce a reject function which guarantees control over the true risk based on DNN softmax output. They later introduce SelectiveNet (Geifman and El-Yaniv 2019), a three-headed network to jointly train classification and rejection functions on the normal domain. Along similar lines, Guo et al. (Guo et al. 2017) presented temperature scaling as a post-processing calibration technique to adjust the model probability estimates being off due to over fitting. In an application to autonomous vehicles, Hecker et al. (Hecker, Dai, and Van Gool 2018) added and trained a failure head to the network in order to learn to predict the occurrence of model failures.

In our recent paper, we presented an error detector for regression models in autonomous vehicle applications (Mohseni, Jagadeesh, and Wang 2019). We proposed a new design which trains a student model (*failure predictor*) for predicting the teacher model’s (*main model*) error at run-time. Figure 2 shows how the student model learns teacher model’s prediction loss on a validation set in order to predict its failures on the test set. We also train the student model using the saliency maps from the main model to improve failure prediction performance. We evaluate our failure predictor model based on prediction error and the driving safety gained by this system.

**Out-of-distribution Error Detectors** Out-of-distribution (OOD) sample or outlier refers to inputs that are outside the normal training distribution. OOD error is referred to ML model misclassification error for OOD samples. Examples of OOD samples in autonomous vehicles include unique, unusual or unknown road signs, road marks, or rare object or scenario which either was not included in the training set or the model was not able to learn (e.g. due to class imbalance) during the training process. OOD error is an inherent problem with ReLU family activation functions that they produce arbitrary high confidence as inputs get further from the training distribution (Hein, Andriushchenko, and Bitterwolf 2019). However, various techniques known as OOD detector, novelty detector, and outlier detector have been pro-

posed to detect OOD samples. Examples of OOD detector techniques include revising network architecture for learning the prediction-confidence (DeVries and Taylor 2018), employing ensemble of leaving-out classifiers (Vyas et al. 2018), and self-supervised representation learning (Golan and El-Yaniv 2018) approaches for outlier detection. On the other hand, a fast and low cost approach in OOD detection is to use class probabilities as a measure for OOD detection (Hendrycks and Gimpel 2016). In this regard, new techniques have been proposed to calibrate DNN decision-boundaries for robust OOD detection (Lee et al. 2017).

In our recent work, we present a fast and memory efficient OOD error detection technique by embedding and training reject options into any DNN discriminative model with minimal architectural changes (Mohseni et al. 2020). This is shown in Figure 3. Our fundamental idea is to exploit the high-level feature learning capacity of DNNs to jointly learn generalizable outlier features as well as in-distribution features for normal classification both in one network. Figure 3 shows how we train additional reject functions in the last layer of our neural network using a two steps of supervised (with labeled in-distribution training set) and self-supervised (with free unlabeled OOD natural samples) training. Our evaluation results show the proposed self-supervised learning of OOD features can very well generalize to reject other unseen distributions.

### Algorithm Robustness

The second practical solution for *ML safety* leverages robustness techniques to improve *Safety Margins* of ML models in autonomous vehicles. Robustness techniques in ML research improve resiliency of the algorithms to unseen samples, natural corruptions and perturbations, adversarial example, and domain shifts. ML literature present multiple techniques such as dataset augmentation, noise injection, and multi-task learning to regularize DNNs to learning generalizable features (Goodfellow, Bengio, and Courville 2016). Other techniques, including transfer learning (Hendrycks, Lee, and Mazeika 2019), have shown to improve model robustness by transferring universal representation from a pre-trained model to the new domain. Further, Zhang and LeCun (Zhang and LeCun 2017) explored using unlabeled free data to regularize model training for its robustness. In the following, we review two major safety-related types of machine learning robustness techniques for open-world tasks. We will also briefly review robustness and detection techniques for adversarial examples in the discussion section.

**Robustness to Domain Shift** Domain shift (also known as distribution shift and dataset shift) describes variations in the input data distribution in comparison to the training set. Distribution shift reduces the operational performance compared to the test set performance by breaking the i.i.d assumption between training and testing data. In this regard, *domain generalization* is a crucial aspect of machine learning algorithms for open-world applications such as autonomous vehicles which data is captured from uncontrolled and fast changing environment. Domain generaliza-

Table 1: Table of practical machine learning techniques to improve the safety of ML algorithms. Left column presents engineering safety strategies and right column maps machine learning techniques with three representative research papers for each row.

Safety Strategy	Practical AI Safety Opportunities
Inherently Safe Design	Design Specification ( (Amodei et al. 2016; Leike et al. 2017; Seshia et al. 2018))
	Implementation Transparency ( (Bojarski et al. 2018; Zeiler and Fergus 2014; Adebayo et al. 2018))
	Formal Verification ((Seshia, Sadigh, and Sastry 2016; Dvijotham et al. 2018; Wang et al. 2018))
Safe Fail	Uncertainty Estimation ((Lakshminarayanan, Pritzel, and Blundell 2017; Gal and Ghahramani 2016))
	In-distribution error detection ((Geifman and El-Yaniv 2017; 2019; Guo et al. 2017))
	Out-of-distribution error detection ( (Mohseni et al. 2020; Golan and El-Yaniv 2018; Vyas et al. 2018))
Safety Margin	Domain Generalization ((Zhang et al. 2019; Ganin and Lempitsky 2014; Lee, Eum, and Kwon 2019))
	Perturbation and Corruptions Robustness ((Geirhos et al. 2019; Huang et al. 2018; Hendrycks, Lee, and Mazeika 2019))

tion could be achieved from many different ways. One approach is adversarial domain adaptation (Ganin and Lempitsky 2014) that leverages from massive amount of unlabeled data captured from the targeted domain. For example, Zhang et al. (Zhang et al. 2019) implemented learning-based approaches to synthesize foreground objects and background contexts for new training samples. Multi-task learning is another technique to improve model robustness by simultaneously learning two (or more) tasks. For instance, Tang et al. (Tang et al. 2019) presents a pose-aware multi-task vehicle re-identification techniques to overcome view-point dependency of the objects. They created and used large-scale highly randomized synthetic dataset with automatically annotated vehicle attributes for training. From a different approach, Lee et al. (Lee, Eum, and Kwon 2019) propose using ensemble of models to capture and learn different poses and viewing angles of objects to improve overall robustness. Also, to improve robustness of object detection models to occlusions and deformations, Wang et al. (Wang, Shrivastava, and Gupta 2017) used an adversarial network to generate hard positive examples.

In a recent work (Wu et al. 2019), we proposed a new technique for improving model robustness to domain shift in unmanned aerial vehicle (UAV). We cast an object detection problem as a cross-domain object detection problem with multiple fine-grained domains. We then train our object detection model to extract invariant features shared by many different “non-ideal” variations (e.g., weather conditions, camera angle, light conditions) of the target domains. In order to do so, we add a nuisance disentangled feature transform block on the input and nuisance prediction branches (one for each non-ideal condition) in a modular fashion and jointly train the final network in an adversarial setting. Our implementation on Faster-RCNN backbone (Wang, Shrivastava, and Gupta 2017) shows superior results compared to vanilla baseline on improving model robustness to weather, altitude, and view variations in UAV images.

**Robustness to Corruptions and Perturbations** Natural data perturbation and corruption commonly exist in open-world settings. Benchmarking DNN robustness to corruption and perturbations (Hendrycks and Dietterich 2019)

shows machine learning models exhibit unexpected prediction errors on simple perturbations. Achieving model robustness to natural corruptions (e.g., due to camera lens, snow, rain, fog in image data) and perturbations (e.g., sensor transients error, electromagnetic interference on sensors) requires techniques to improve model robustness above their *clean datasets*. Previously, classical data augmentations were used to gain robustness to simple image variations like rotation and scaling (Goodfellow, Bengio, and Courville 2016). Other techniques such as using adaptive algorithm for choosing the augmentation transformations (Fawzi et al. 2016) and random patch erasing (Zhong et al. 2017) also shown to be effective for both robustness and representation learning. Recently, advanced augmentations such as style transfer (Geirhos et al. 2019) have shown to improve model robustness to texture bias. Another line of research proposes larger networks (Huang et al. 2018) to improve DNN robustness through multiscale and redundant feature learning.

Adversarial perturbation (Goodfellow, Shlens, and Szegedy 2014), on the other hand, are small but worst-case perturbations intentionally created by an attacker so that the perturbed sample results in the model misclassify the sample with high confidence. We distinguish safety hazard due to natural perturbations from security hazard due to adversarial perturbations as the latter intentionally exploit system vulnerabilities to cause harm. Security concerns related to adversarial perturbation are briefly mentioned in the discussion about future work.

## 4 Conclusion and Future Work

In this work, we presented a review and categorization of classical software safety methods with fundamental limitations in machine learning algorithms. The impetus of this work was to leverage from both engineering safety strategies and state-of-the-art machine learning techniques to enhance dependability and safety of machine learning components in autonomous systems. In this regard, maintaining the safety of autonomous vehicles requires a multidisciplinary effort across multiple fields including human-computer interactions, machine learning, software engineering, hardware engineering (Koopman and Wagner 2017). We briefly review

and discuss other dimensions of safety and security of ML-infused systems that could benefit from research communities attention.

**Security Risks of Adversarial Attacks:** An adversarial example is a clean image delicately perturbed (by an adversary) with a small distortion so that it visually looks identical to the original clean image but is misclassified by the ML model. Despite its popularity, adversarial attacks are not considered primarily a safety concern rather a security limitation (Carlini and Wagner 2017). Two main defense approaches for adversarial attacks are *detection* and *robustness*. For example, Smith and Gal (Smith and Gal 2018) present a case of MC-drop out uncertainty estimation technique for detecting adversarial examples. Further, to improve model robustness and resiliency to adversary perturbation Papernot et al. proposed an effective defense based on model distillation (Papernot et al. 2016). Still, the problem of adversarial examples is an unsolved problem when considering that attackers always react to current defense techniques by designing stronger attacks to undermine security of machine learning components (Carlini and Wagner 2017).

**Procedural Safeguards for ML Safety:** Beyond functional safety of the system, procedural safeguards help operators and product end-users (e.g. driver in autonomous vehicle) avoid unintentional misuse of system due to lack of instructions and unawareness (Varshney 2016). User experience (UX) design and algorithmic transparency are two approaches to improve safety of the operation in autonomous vehicles. In such cases, end-users can benefit from explainable UX designs which provides useful and comprehensible information about model reasoning (Gunning 2017) and prediction uncertainty (Michelmoré, Kwiatkowska, and Gal 2018). For example, UX design could leverage real-time visualization of model uncertainty for vehicle detection and path planning to help improve driver’s understanding of vehicle safety on the road.

In our future work, we plan to review recent procedural safeguards designs and studies for autonomous vehicle applications. Including key techniques and factors in calibrating user trust in human-AI systems.

## References

Adebayo, J.; Gilmer, J.; Muelly, M.; Goodfellow, I.; Hardt, M.; and Kim, B. 2018. Sanity checks for saliency maps. In *NIPS*.

Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; and Mané, D. 2016. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*.

Bansal, G., and Weld, D. S. 2018. A coverage-based utility model for identifying unknown unknowns. In *Thirty-Second AAAI Conference on AI*.

Bojarski, M.; Del Testa, D.; Dworakowski, D.; Firner, B.; Flepp, B.; Goyal, P.; Jackel, L. D.; Monfort, M.; Muller, U.; Zhang, J.; et al. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*.

Bojarski, M.; Choromanska, A.; Choromanski, K.; Firner, B.; Ackel, L. J.; Muller, U.; Yeres, P.; and Zieba, K. 2018. Visualbackprop: Efficient visualization of cnns for autonomous driving. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 1–8. IEEE.

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

DeVries, T., and Taylor, G. W. 2018. Learning confidence for out-of-distribution detection in neural networks. *ICLR*.

Dreossi, T.; Fremont, D. J.; Ghosh, S.; Kim, E.; Ravanbakhsh, H.; Vazquez-Chanlatte, M.; and Seshia, S. A. 2019. Verifai: A toolkit for the formal design and analysis of artificial intelligence-based systems. In *International Conference on Computer Aided Verification*, 432–442. Springer.

Dvijotham, K.; Goyal, S.; Stanforth, R.; Arandjelovic, R.; O’Donoghue, B.; Uesato, J.; and Kohli, P. 2018. Training verified learners with learned verifiers. *arXiv preprint arXiv:1805.10265*.

Fawzi, A.; Samulowitz, H.; Turaga, D.; and Frossard, P. 2016. Adaptive data augmentation for image classification. In *2016 IEEE International Conference on Image Processing (ICIP)*, 3688–3692. Ieee.

Gal, Y., and Ghahramani, Z. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 1050–1059.

Ganin, Y., and Lempitsky, V. 2014. Unsupervised domain adaptation by backpropagation. *arXiv preprint arXiv:1409.7495*.

Geifman, Y., and El-Yaniv, R. 2017. Selective classification for deep neural networks. In *NIPS*, 4878–4887.

Geifman, Y., and El-Yaniv, R. 2019. Selectivenet: A deep neural network with an integrated reject option. *arXiv preprint arXiv:1901.09192*.

Geirhos, R.; Rubisch, P.; Michaelis, C.; Bethge, M.; Wichmann, F. A.; and Brendel, W. 2019. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *ICLR*.

Golan, I., and El-Yaniv, R. 2018. Deep anomaly detection using geometric transformations. In *NIPS*, 9758–9769.

Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep learning*. MIT press.

Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.

Gunning, D. 2017. Explainable artificial intelligence (xai). *Defense Advanced Research Projects Agency (DARPA), nd Web 2*.

Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *ICML*.

Hecker, S.; Dai, D.; and Van Gool, L. 2018. Failure prediction for autonomous driving. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, 1792–1799. IEEE.

- Hein, M.; Andriushchenko, M.; and Bitterwolf, J. 2019. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *CVPR*.
- Hendrycks, D., and Dietterich, T. 2019. Benchmarking neural network robustness to common corruptions and perturbations. *ICLR*.
- Hendrycks, D., and Gimpel, K. 2016. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*.
- Hendrycks, D.; Lee, K.; and Mazeika, M. 2019. Using pre-training can improve model robustness and uncertainty. *arXiv preprint arXiv:1901.09960*.
- Hernández-Orallo, J.; Martínez-Plumed, F.; Avin, S.; and hÉigeartaigh, S. Ó. 2019. Surveying safety-relevant ai characteristics. In *SafeAI@ AAI*.
- Huang, G.; Chen, D.; Li, T.; Wu, F.; van der Maaten, L.; and Weinberger, K. Q. 2018. Multi-scale dense networks for resource efficient image classification. *ICLR*.
- Kendall, A., and Gal, Y. 2017. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*, 5574–5584.
- Koopman, P., and Wagner, M. 2017. Autonomous vehicle safety: An interdisciplinary challenge. *IEEE Intelligent Transportation Systems Magazine* 9(1):90–96.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NIPS*.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *nature* 521(7553):436–444.
- Lee, K.; Lee, H.; Lee, K.; and Shin, J. 2017. Training confidence-calibrated classifiers for detecting out-of-distribution samples. *arXiv preprint arXiv:1711.09325*.
- Lee, H.; Eum, S.; and Kwon, H. 2019. Me r-cnn: Multi-expert r-cnn for object detection. *IEEE Transactions on Image Processing*.
- Leike, J.; Martic, M.; Krakovna, V.; Ortega, P. A.; Everitt, T.; Lefrancq, A.; Orseau, L.; and Legg, S. 2017. Ai safety gridworlds. *arXiv preprint arXiv:1711.09883*.
- McAllister, R.; Gal, Y.; Kendall, A.; Van Der Wilk, M.; Shah, A.; Cipolla, R.; and Weller, A. V. 2017. Concrete problems for autonomous vehicle safety: Advantages of bayesian deep learning. *IJCAI*.
- Michelmore, R.; Kwiatkowska, M.; and Gal, Y. 2018. Evaluating uncertainty quantification in end-to-end autonomous driving control. *arXiv preprint arXiv:1811.06817*.
- Mohseni, S.; Pitale, M.; Yadawa, J.; and Wang, Z. 2020. Self-supervised learning for generalizable out-of-distribution detection. In *AAAI*.
- Mohseni, S.; Jagadeesh, A.; and Wang, Z. 2019. Predicting model failure using saliency maps in autonomous driving systems. *ICML Workshop on Uncertainty & Robustness in Deep Learning*.
- Ortega, P. A., and Maini, V. 2018. Building safe artificial intelligence: specification, robustness, and assurance. <https://medium.com/deepmindsafetyresearch/building-safe-artificial-intelligence-52f5f75058f1>.
- Papernot, N.; McDaniel, P.; Wu, X.; Jha, S.; and Swami, A. 2016. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*, 582–597. IEEE.
- Salay, R.; Queiroz, R.; and Czarnecki, K. 2017. An analysis of iso 26262: Using machine learning safely in automotive software. *arXiv preprint arXiv:1709.02435*.
- Seshia, S. A.; Desai, A.; Dreossi, T.; Fremont, D. J.; Ghosh, S.; Kim, E.; Shivakumar, S.; Vazquez-Chanlatte, M.; and Yue, X. 2018. Formal specification for deep neural networks. In *International Symposium on Automated Technology for Verification and Analysis*, 20–34. Springer.
- Seshia, S. A.; Sadigh, D.; and Sastry, S. S. 2016. Towards verified artificial intelligence. *arXiv preprint arXiv:1606.08514*.
- Smith, L., and Gal, Y. 2018. Understanding measures of uncertainty for adversarial example detection. *arXiv preprint arXiv:1803.08533*.
- Tang, Z.; Naphade, M.; Birchfield, S.; Tremblay, J.; Hodge, W.; Kumar, R.; Wang, S.; and Yang, X. 2019. Pamtri: Pose-aware multi-task learning for vehicle re-identification using highly randomized synthetic data. In *ICCV*.
- Varshney, K. R. 2016. Engineering safety in machine learning. In *2016 Information Theory and Applications Workshop (ITA)*, 1–5. IEEE.
- Vyas, A.; Jammalamadaka, N.; Zhu, X.; Das, D.; Kaul, B.; and Willke, T. L. 2018. Out-of-distribution detection using an ensemble of self supervised leave-out classifiers. In *ECCV*, 550–564.
- Wang, S.; Pei, K.; Whitehouse, J.; Yang, J.; and Jana, S. 2018. Efficient formal safety analysis of neural networks. In *NIPS*, 6367–6377.
- Wang, X.; Shrivastava, A.; and Gupta, A. 2017. A-fast-rcnn: Hard positive generation via adversary for object detection. In *CVPR*.
- Wu, Z.; Suresh, K.; Narayanan, P.; Xu, H.; Kwon, H.; and Wang, Z. 2019. Delving into robust object detection from unmanned aerial vehicles: A deep nuisance disentanglement approach. In *ICCV*.
- Zeiler, M. D., and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *ECCV*, 818–833. Springer.
- Zhang, X., and LeCun, Y. 2017. Universum prescription: Regularization using unlabeled data. In *AAAI*.
- Zhang, X.; Wang, Z.; Liu, D.; and Ling, Q. 2019. Dada: Deep adversarial data augmentation for extremely low data regime classification. In *ICASSP*. IEEE.
- Zhong, Z.; Zheng, L.; Kang, G.; Li, S.; and Yang, Y. 2017. Random erasing data augmentation. *arXiv preprint arXiv:1708.04896*.
- Zhou, Y., and Tuzel, O. 2018. Voxnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*.