

---

# Deep Probabilistic Accelerated Evaluation: A Robust Certifiable Rare-Event Simulation Methodology for Black-Box Safety-Critical Systems

---

**Mansur Arief**  
Carnegie Mellon University

**Zhiyuan Huang**  
Tongji University

**Guru K.S. Kumar**  
Carnegie Mellon University

**Yuanlu Bai**  
Columbia University

**Shengyi He**  
Columbia University

**Wenhao Ding**  
Carnegie Mellon University

**Henry Lam**  
Columbia University

**Ding Zhao**  
Carnegie Mellon University

## Abstract

Evaluating the reliability of intelligent physical systems against rare safety-critical events poses a huge testing burden for real-world applications. Simulation provides a useful platform to evaluate the extremal risks of these systems before their deployments. Importance Sampling (IS), while proven to be powerful for rare-event simulation, faces challenges in handling these learning-based systems due to their black-box nature that fundamentally undermines its efficiency guarantee, which can lead to under-estimation without diagnostically detected. We propose a framework called Deep Probabilistic Accelerated Evaluation (Deep-PrAE) to design statistically guaranteed IS, by converting black-box samplers that are versatile but could lack guarantees, into one with what we call a relaxed efficiency certificate that allows accurate estimation of bounds on the safety-critical event probability. We present the theory of Deep-PrAE that combines the dominating point concept with rare-event set learning via deep neural network classifiers, and demonstrate its effectiveness in numerical examples including the safety-testing of an intelligent driving algorithm.

## 1 Introduction

The unprecedented deployment of intelligent physical systems on many real-world applications comes with the need for safety validation and certification (Kalra and Paddock, 2016; Koopman and Wagner, 2017; Uesato et al., 2018). For systems that interact with humans and are potentially safety-critical - which can range from medical systems to self-driving cars and personal assistive robots - it is imperative to rigorously assess their risks before their full-scale deployments. The challenge, however, is that these risks are often associated precisely to how AI reacts in rare and catastrophic scenarios which, by their own nature, are not sufficiently observed.

The challenge of validating the safety of intelligent systems described above is, unfortunately, insusceptible to traditional test methods. In the self-driving context, for instance, the goal of validation is to ensure the AI-enabled system reduces human-level accident rate (in the order of 1.5 per  $10^8$  miles of driving), thus delivering enhanced safety promise to the public (Evan, 2016; Kalra and Paddock, 2016; NHTSA, 2016). Formal verification, which mathematically analyzes and verifies autonomous design, faces challenges when applied to black-box or complex models due to the lack of analytic tractability to formulate failure cases or consider all execution trajectories (Clarke et al., 2018). Automated scenario selection approaches generate test cases based on domain knowledge (Wegener and Bühler, 2004) or adaptive searching algorithms (such as adaptive stress testing; Koren et al., 2018), which is more implementable but falls short of rigor. Test matrix approaches, such as Euro NCAP (NHTSA, 2007), use prepopulated test cases extracted from crash databases, but they only contain historical human-driver information. The closest analog to the latter for self-driving ve-

hicles is “naturalistic tests”, which means placing them in real-world environments and gathering observations. This method, however, is economically prohibitive because of the rarity of the target conflict events (Zhao et al., 2017; Arief et al., 2018; Claybrook and Kildare, 2018; O’Kelly et al., 2018).

Because of all these limitations, simulation-based tests surface as a powerful approach to validate complex black-box designs (Corso et al., 2020). This approach operates by integrating the target intelligent algorithm into an interacting virtual simulation platform that models the surrounding environment. By running enough Monte Carlo sampling of this (stochastic) environment, one hopes to observe catastrophic conflict events and subsequently conduct statistical analyses. This approach is flexible and scalable, as it hinges on building a virtual environment instead of physical systems, and provides a probabilistic assessment on the occurrences and behaviors of safety-critical events (Koopman and Wagner, 2018).

Nonetheless, similar to the challenge encountered by naturalistic tests, because of their rarity, safety-critical events are seldom observed in the simulation experiments. In other words, it could take an enormous amount of Monte Carlo simulation runs to observe one “hit”, and this in turn manifests statistically as a large estimation variance per simulation run relative to the target probability of interest (i.e., the so-called *relative error*; Lécuyer et al., 2010). This problem, which is called rare-event simulation (Bucklew, 2013), is addressed conventionally under the umbrella of variance reduction, which includes a range of techniques from importance sampling (IS) (Juneja and Shahabuddin, 2006; Blanchet and Lam, 2012) to multi-level splitting (Glasserman et al., 1999; Villén-Altamirano and Villén-Altamirano, 1994). Typically, to ensure the relative error is dramatically reduced, one has to analyze the underlying model structures to gain understanding of the rare-event behaviors, and leverage this knowledge to design good Monte Carlo schemes (Juneja and Shahabuddin, 2006; Dean and Dupuis, 2009). For convenience, we call such relative error reduction guarantee an *efficiency certificate*.

Our main focus of this paper is on rare-event problems with the underlying model unknown or too complicated to support analytical tractability. In this case, traditional variance reduction approaches may fail to provide an efficiency certificate. Moreover, we will explain how some existing “black-box” variance reduction techniques, while versatile and powerful, could lead to dangerous *under-estimation* of a rare-event probability without detected diagnostically due to a lack of efficiency certificate. This motivates us to study a framework to convert these black-box methods into

one that has rigorous certificate. More precisely, our framework consists of three ingredients:

**Relaxed efficiency certificate:** We shift the estimation of target rare-event probability to an upper (and lower) bound, in a way that supports the integration of learning errors into variance reduction without giving up estimation correctness.

**Set-learning with one-sided error:** We design learning algorithms based on deep neural network classifier to create outer (or inner) approximations of rare-event sets. This classifier has a special property that, under a geometric property called orthogonal monotonicity, it exhibits zero false negative rates.

**Deep-learning-based IS:** With the deep-learning based rare-event set approximation, we search the so-called *dominating points* in rare-event analysis to create IS that achieves the relaxed efficiency certificate.

We call our framework consisting of the three ingredients above *Deep Probabilistic Accelerated Evaluation (Deep-PrAE)*, where “Accelerated Evaluation” follows terminologies in recent approaches for the safety-testing of autonomous vehicles (Zhao et al., 2016; Huang et al., 2018). In the set-learning step in Deep-PrAE, the samples fed into our deep classifier can be generated by any black-box algorithms including the cross-entropy (CE) method (De Boer et al., 2005; Rubinstein and Kroese, 2013) and particle approaches such as adaptive multi-level splitting (AMS) (Au and Beck, 2001; Cérou and Guyader, 2007; Webb et al., 2018). Deep-PrAE turns these samples into an IS with an efficiency certificate against undetected under-estimation. Our approach is robust in the sense that it provides a tight bound for the target rare-event probability if the underlying classifier is expressive enough, while it still provides a correct, though conservative, bound if the classifier is weak. To our best knowledge, such type of guarantees and robustness features is the first of its kind in the rare-event simulation literature, and we envision our work to lay the foundation for further improvements to design certified methods for evaluating more sophisticated intelligent designs.

## 2 Statistical Challenges in Black-Box Rare-Event Simulation

Our evaluation goal is the probabilistic assessment of a complex physical system invoking rare but catastrophic events in a stochastic environment. For concreteness, we write this rare-event probability  $\mu = P(X \in \mathcal{S}_\gamma)$ . Here  $X$  is a random vector in  $\mathbb{R}^d$  that denotes the environment, and is distributed according to  $p$ .  $\mathcal{S}_\gamma$  denotes a safety-critical set on the interaction between the physical system and the environment. The “rarity”

parameter  $\gamma \in \mathbb{R}$  is considered a large number, with the property that as  $\gamma \rightarrow \infty$ ,  $\mu \rightarrow 0$  (Think of, e.g.,  $\mathcal{S}_\gamma = \{x : f(x) \geq \gamma\}$  for some risk function  $f$  and exceedance threshold  $\gamma$ ). We will work with Gaussian  $p$  for the ease of analysis, but our framework is more general (i.e., applies to Gaussian mixtures and other light-tailed distributions). Here, we explain intuitively the main concepts and challenges in black-box rare-event simulation, leaving the details to Appendix A.

**Monte Carlo Efficiency.** Suppose we use a Monte Carlo estimator  $\hat{\mu}_n$  to estimate  $\mu$ , by running  $n$  simulation runs in total. Since  $\mu$  is tiny, the error of a meaningful estimation must be measured in relative term, i.e., we would like

$$P(|\hat{\mu}_n - \mu| > \epsilon\mu) \leq \delta \quad (1)$$

where  $\delta$  is some confidence level (e.g.,  $\delta = 5\%$ ) and  $0 < \epsilon < 1$ .

Suppose that  $\hat{\mu}_n$  is unbiased and is an average of  $n$  i.i.d. simulation runs, i.e.,  $\hat{\mu}_n = (1/n) \sum_{i=1}^n Z_i$  for some random unbiased output  $Z_i$ . We define the *relative error*  $RE = \text{Var}(Z_i)/\mu^2$  as the ratio of variance (per-run) and squared mean. Importantly, to attain (1), a sufficient condition is  $n \geq RE/(\delta\epsilon^2)$ . So, when  $RE$  is large, the required Monte Carlo size is also large.

**Challenges in Naive Monte Carlo.** Let  $Z_i = I(X_i \in \mathcal{S}_\gamma)$  where  $I(\cdot)$  denotes the indicator function, and  $X_i$  is an i.i.d. copy of  $X$ . Since  $Z_i$  follows a Bernoulli distribution,  $RE = (1 - \mu)/\mu$ . Thus, the required  $n$  scales linearly in  $1/\mu$  (when  $\mu$  is tiny). This demanding condition is a manifestation of the difficulty in hitting  $\mathcal{S}_\gamma$ . In the standard large deviations regime (Amir Dembo, 2010; Dupuis and Ellis, 2011) where  $\mu$  is exponentially small in  $\gamma$ , the required Monte Carlo size  $n$  would grow *exponentially* in  $\gamma$ .

**Variance Reduction.** The severe burden when using naive Monte Carlo motivates techniques to drive down  $RE$ . First we introduce the following notion:

**Definition 1.** We say an estimator  $\hat{\mu}_n$  satisfies an *efficiency certificate* to estimate  $\mu$  if it achieves (1) with  $n = \tilde{O}(\log(1/\mu))$ , for given  $0 < \epsilon, \delta < 1$ .

In the above,  $\tilde{O}(\cdot)$  denotes a polynomial growth in  $\cdot$ . If  $\hat{\mu}_n$  is constructed from  $n$  i.i.d. samples, then the efficiency certificate can be attained with  $RE = \tilde{O}(\log(1/\mu))$ . Note that in the large deviations regime, the sample size  $n$  used in a certifiable estimator is reduced from exponential in  $\gamma$  in naive Monte Carlo to *polynomial* in  $\gamma$ .

*Importance sampling (IS)* is a prominent technique to achieve efficiency certificate (Glynn and Iglehart, 1989). IS generates  $X$  from another distribution  $\tilde{p}$  (called IS distribution), and outputs  $\hat{\mu}_n =$

$(1/n) \sum_{i=1}^n L(X_i)I(X_i \in \mathcal{S}_\gamma)$  where  $L = dp/d\tilde{p}$  is the likelihood ratio, or the Radon-Nikodym derivative, between  $p$  and  $\tilde{p}$ . Via a change of measure, it is easy to see that  $\hat{\mu}_n$  is unbiased for  $\mu$ . The key is to control its  $RE$  by selecting a good  $\tilde{p}$ . This requires analyzing the behavior of the likelihood ratio  $L$  under the rare event, and in turn understanding the rare-event sample path dynamics (Juneja and Shahabuddin, 2006).

**Perils of Black-Box Variance Reduction Algorithms.** Unfortunately, in black-box settings where complete model knowledge and analytical tractability are unavailable, the classical IS methodology faces severe challenges. To explain this, we first need to understand how efficiency certificate can be obtained based on the concept of *dominating points*. From now on, we consider input  $X \in \mathbb{R}^d$  from a Gaussian distribution  $N(\lambda, \Sigma)$  where  $\Sigma$  is positive definite.

**Definition 2.** A set  $A_\gamma \subset \mathbb{R}^d$  is a *dominating set* for the set  $\mathcal{S}_\gamma \subset \mathbb{R}^d$  associated with the distribution  $N(\lambda, \Sigma)$  if for any  $x \in \mathcal{S}_\gamma$ , there exists at least one  $a \in A_\gamma$  such that  $(a - \lambda)^T \Sigma^{-1}(x - a) \geq 0$ . Moreover, this set is *minimal* in the sense that if any point in  $A_\gamma$  is removed, then the remaining set no longer satisfies the above condition. We call any point in  $A_\gamma$  a *dominating point*.

The dominating set comprises the “corner” cases where the rare event occurs (Sadowsky and Bucklew, 1990). In other words, each dominating point  $a$  encodes, in a local region, the most likely scenario should the rare event happen, and this typically corresponds to the highest-density point in this region. Locality here refers to the portion of the rare-event set that is on one side of the hyperplane cutting through  $a$  (see Figure 1(a)).

Intuitively, to increase the frequency of hitting the rare-event set (and subsequently to reduce variance), an IS would translate the distributional mean from  $\lambda$  to the global highest-density point in the rare-event set. The delicacy, however, is that this is *insufficient* to control the variance, due to the “overshoots” arising from sampling randomness. In order to properly control the overall variance, one needs to divide the rare-event set into local regions governed by dominating points, and using a mixture IS distribution that accounts for *all* of them. This approach gives a certifiable IS, described as follows:

**Theorem 1 (Certifiable IS).** Suppose  $\mathcal{S}_\gamma = \bigcup_j \mathcal{S}_\gamma^j$ , where each  $\mathcal{S}_\gamma^j$  is a “local” region corresponding to a dominating point  $a_j \in A_\gamma$  associated with the distribution  $N(\lambda, \Sigma)$ , with conditions stated precisely in Theorem 5 in the Appendix. Then the IS estimator constructed by  $n$  i.i.d. outputs drawn from the IS distribution  $\sum_j \alpha_j N(a_j, \Sigma)$  achieves an efficiency certificate in estimating  $\mu = P(X \in \mathcal{S}_\gamma)$ .

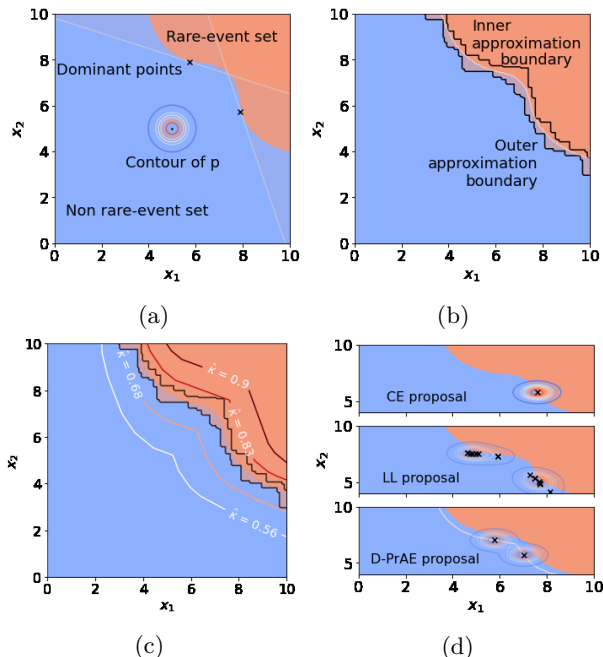


Figure 1: (a) An example of  $\mathcal{S}_\gamma$  with two dominating points (b) Outer- and inner- approximations of  $\mathcal{S}_\gamma$  (c)  $\hat{\kappa}$  tuning for Stage 1 Alg. 1 (d) IS proposals: too few dominating points for CE with too simple parametric class, too many for LL, and a balance for Deep-PrAE.

On the contrary, if the Gaussian (mixture) IS distribution misses any of the dominating points, then the resulting estimate may be utterly unreliable for two reasons. First, not only that efficiency certificate may fail to hold, but its RE can be arbitrarily large. Second, even more dangerously, this poor performance can be empirically hidden and leads to a systematic *under-estimation* of the rare-event probability without being detected. In other words, in a given experiment, we may observe a reasonable empirical relative error (i.e., sample variance over squared sample mean), yet the estimate is much lower than the correct value. These are revealed in the following example:

**Theorem 2** (Perils of under-estimation). *Suppose we estimate  $\mu = P(X \geq \gamma \text{ or } X \leq -k\gamma)$  where  $X \sim p = N(0, 1)$  and  $0 < k < 3$ . We choose  $\tilde{p} = N(\gamma, 1)$  as the IS distribution to obtain  $\hat{\mu}_n$ . Then 1) The relative error of  $\hat{\mu}_n$  grows exponentially in  $\gamma$ . 2) If  $n$  is polynomial in  $\gamma$ , we have  $P(|\hat{\mu}_n - \Phi(\gamma)| > \varepsilon \Phi(\gamma)) = O(\frac{\gamma}{n\varepsilon^2})$  for any  $\varepsilon > 0$  where  $\Phi(\gamma) = P(X \geq \gamma) < \mu$ , and the empirical relative error =  $O(n^2)$  with probability higher than  $1 - 1/2^n$ .*

The second conclusion in Theorem 2 implies that the estimator  $\hat{\mu}_n$ , built from an IS with a missed dominating point, systematically under-estimates the target  $\mu$ , yet with high probability its empirical relative error

grows only polynomially in  $\gamma$ , thus wrongly fooling the user that the estimator is efficient.

With this, we now explain why using black-box variance reduction algorithms can be dangerous - in the sense of not having an efficiency certificate and, more importantly, the risk of an unnoticed systematic under-estimation. In the literature, there are two lines of techniques that apply to black-box problems. The first line is the CE method, which uses optimization to search for a good parametrization over a parametric class of IS. The objective criteria include the cross-entropy (with an oracle-best zero-variance IS distribution; Rubinstein and Kroese 2013; De Boer et al. 2005) and estimation variance (Arouna, 2004). Without closed-form expressions, and also to combat the rare-event issue, one typically solves a sequence of empirical optimization problems, starting from a “less rare” problem (i.e., smaller  $\gamma$ ) and gradually increasing the rarity with updated empirical objectives using better IS samples. Achieving efficiency requires both a sufficiently expressive parametric IS class and parameter convergence (so that at the end all the dominating points are accounted for). The second line of methods is the multi-level splitting or subsimulation (Au and Beck, 2001; Cérou and Guyader, 2007), a particle method in lieu of IS, which relies on enough mixing of descendant particles. Full analyses on these methods to reach efficiency certificate appear challenging, and without one the estimators could be under-estimated, and without detected, as illustrated in Theorem 2. We discuss more details of CE and AMS in Appendix E. Note that there are other variants of CE and AMS. The former include enhanced CE such as Markov chain IS (Botev et al., 2013, 2016; Grace et al., 2014), neural network IS (Müller et al., 2019) and nonparametric CE (Rubinstein, 2005). The latter include RESTART which works similarly as subset simulation and splitting but performs a number of simulation retrials after entering regions with a higher importance function value (Villén-Altamirano, 2010). Similar to standard CE and AMS, these methods also face challenges in satisfying an efficiency certificate.

Lastly, we briefly review several other methods with guarantees similar to our efficiency certificate, but relies heavily on structural knowledge. The first one is large-deviations-based IS including sequential exponential tilting (Bucklew, 2004; Asmussen and Glynn, 2007; Siegmund, 1976) and mixture-based proposals (Chen et al., 2019). Another method, which is especially powerful for heavy tailed problems, is conditional Monte Carlo which reduces the variance by sampling conditional on some auxiliary random variables (Asmussen and Kroese, 2006).

Compared to existing methods as reviewed above, our

novelty is to tackle black-box problems while sustaining a correctness guarantee, via a new certificate and a careful integration of set-learning with the dominating point machinery.

### 3 The Deep Probabilistic Accelerated Evaluation Framework

We propose the Deep-PrAE framework to overcome the challenges faced by black-box variance reduction algorithms. This framework comprises two stages: First is to learn the rare-event set from a first-stage sample batch, by viewing set learning as a classification task. These first-stage samples can be drawn from any rare-event sampling methods including CE and AMS. Second is to apply an efficiency-certified IS on the rare-event probability over the learned set. Algorithm 1 shows our main procedure. The key to achieving an ultimate efficiency certificate lies in how we learn the rare-event set in Stage 1, which requires two properties:

**Small one-sided generalization error:** “One-sided” generalization error here means the learned set is either an outer or an inner approximation of the unknown true rare-event set, with probability 1. Converting this into a classification, this means the false negative (or positive) rate is exactly 0. “Small” here then refers to the other type of error being controlled.

**Decomposability:** The learned set is decomposable according to dominating points in the form of Theorem 1, so that an efficient mixture IS can apply.

The first property ensures that, even though the learned set can contain errors, the learned rare-event probability is either an upper or lower bound of the truth. This requirement is important as it is difficult to translate the impact of generalization errors into rare-event estimation errors. By Theorem 2, we know that any non-zero error implies the risk of missing out on important regions of the rare-event set, undetectably. The one-sided generalization error allows a shift of our target to valid upper and lower bounds that can be correctly estimated, which is the core novelty of Deep-PrAE. To this end, we introduce a new efficiency notion:

**Definition 3.** We say an estimator  $\hat{\mu}_n$  satisfies an upper-bound relaxed efficiency certificate to estimate  $\mu$  if  $P(\hat{\mu}_n - \mu < -\epsilon\mu) \leq \delta$  with  $n \geq \tilde{O}(\log(1/\mu))$ , for given  $0 < \epsilon, \delta < 1$ .

Compared with the efficiency certificate in (1), Definition 3 is relaxed to only requiring  $\hat{\mu}_n$  to be an upper bound of  $\mu$ , up to an error of  $\epsilon\mu$ . An analogous lower-bound relaxed efficiency certificate can be seen in Appendix D. From a risk quantification viewpoint, the upper bound for  $\mu$  is more crucial, and the lower

---

#### Algorithm 1: Deep-PrAE to estimate $\mu = P(X \in \mathcal{S}_\gamma)$ .

---

**Input:** Black-box evaluator  $I(\cdot \in \mathcal{S}_\gamma)$ , initial Stage 1 samples  $\{(\tilde{X}_i, Y_i)\}_{i=1, \dots, n_1}$ , Stage 2 sampling budget  $n_2$ , input distribution  $N(\lambda, \Sigma)$ .

**Output:** IS estimate  $\hat{\mu}_n$ .

1 **Stage 1 (Set Learning):**

2 Train classifier with positive decision region

$$\bar{\mathcal{S}}_\gamma^\kappa = \{x : \hat{g}(x) \geq \kappa\} \text{ using } \{(\tilde{X}_i, Y_i)\}_{i=1, \dots, n_1};$$

3 Replace  $\kappa$  by  $\hat{\kappa} = \max\{\kappa \in \mathbb{R} : (\bar{\mathcal{S}}_\gamma^\kappa)^c \subset \mathcal{H}(T_0)\}$ ;

4 **Stage 2 (Mixture IS based on Searched dominating points):**

5 Start with  $\hat{A}_\gamma = \emptyset$ ;

6 **While**  $\{x : \hat{g}(x) \geq \hat{\kappa}, (x_j^* - \lambda)^T \Sigma^{-1}(x - x_j^*) < 0, \forall x_j^* \in \hat{A}_\gamma\} \neq \emptyset$  **do**

7 Find a dominating point  $x^*$  by solving the optimization problem

$$x^* = \arg \min_x (x - \lambda)^T \Sigma^{-1}(x - \lambda)$$

$$\text{s.t. } \hat{g}(x) \geq \hat{\kappa},$$

$$(x_j^* - \lambda)^T \Sigma^{-1}(x - x_j^*) < 0, \forall x_j^* \in \hat{A}_\gamma$$

and update  $\hat{A}_\gamma \leftarrow \hat{A}_\gamma \cup \{x^*\}$ ;

8 **End**

9 Sample  $X_1, \dots, X_{n_2}$  from the mixture distribution  $\sum_{a \in \hat{A}_\gamma} (1/|\hat{A}_\gamma|) N(a, \Sigma)$ .

10 Compute the IS estimator

$$\hat{\mu}_n = (1/n_2) \sum_{i=1}^{n_2} L(X_i) I(X_i \in \bar{\mathcal{S}}_\gamma^{\hat{\kappa}}), \text{ where the likelihood ratio } L(X_i) =$$

$$\phi(X_i; \lambda, \Sigma) / (\sum_{a \in \hat{A}_\gamma} (1/|\hat{A}_\gamma|) \phi(X_i; a, \Sigma)) \text{ and}$$

$$\phi(\cdot; \alpha, \Sigma) \text{ denotes the density of } N(\alpha, \Sigma).$$


---

bound serves to assess an estimation gap. The following provides a handy certification:

**Proposition 1** (Achieving relaxed efficiency certificate). Suppose  $\hat{\mu}_n$  is upward biased, i.e.,  $\bar{\mu} := E[\hat{\mu}_n] \geq \mu$ . Moreover, suppose  $\hat{\mu}_n$  takes the form of an average of  $n$  i.i.d. simulation runs  $Z_i$ , with  $RE = \text{Var}(Z_i)/\bar{\mu}^2 = \tilde{O}(\log(1/\bar{\mu}))$ . Then  $\hat{\mu}_n$  possesses the upper-bound relaxed efficiency certificate.

Proposition 1 stipulates that a relaxed efficiency certificate can be attained by an upward biased estimator that has a logarithmic relative error with respect to the biased mean. Appendix C shows an extension of Proposition 1 to two-stage procedures, where the first stage determines the upward biased mean. This upward biased mean, in turn, can be obtained by learning an outer approximation for the rare-event set, giving:

**Corollary 1** (Set-learning + IS). Consider estimat-

ing  $\mu = P(X \in \mathcal{S}_\gamma)$ . Suppose we can learn a set  $\bar{\mathcal{S}}_\gamma$  with any number  $n_1$  of i.i.d. samples  $D_{n_1}$  (drawn from some distribution) such that  $\bar{\mathcal{S}}_\gamma \supset \mathcal{S}_\gamma$  with probability 1. Also suppose that there is an efficiency certificate for an IS estimator for  $\bar{\mu}(D_{n_1}) := P(X \in \bar{\mathcal{S}}_\gamma)$ . Then a two-stage estimator where a constant  $n_1$  number of samples  $D_{n_1}$  are first used to construct  $\bar{\mathcal{S}}_\gamma$ , and  $n_2 = \tilde{O}(\log(1/\bar{\mu}(D_{n_1})))$  samples are used for the IS in the second stage, achieves the upper-bound relaxed efficiency certificate.

To execute the procedure in Corollary 1, we need to learn an outer approximation of the rare-event set. To this end, consider set learning as a classification problem. Suppose we have collected  $n_1$  Stage 1 samples  $\{(\tilde{X}_i, Y_i)\}_{i=1, \dots, n_1}$ , where  $Y_i = 1$  if  $\tilde{X}_i$  is in the rare-event set  $\mathcal{S}_\gamma$ , and 0 otherwise. Here, it is beneficial to use Stage 1 samples that have sufficient presence in  $\mathcal{S}_\gamma$ , which can be achieved via any black-box variance reduction methods. We then consider the pairs  $\{(\tilde{X}_i, Y_i)\}$  where  $\tilde{X}_i$  is regarded as the feature and  $Y_i$  as the binary label, and construct a classifier, say  $\hat{g}(x) : \mathbb{R}^d \rightarrow [0, 1]$ , from some hypothesis class  $\mathcal{G}$  that (nominally) signifies  $P(Y = 1 | X = x)$ . The learned rare-event set  $\bar{\mathcal{S}}_\gamma$  is taken to be  $\{x : \hat{g}(x) \geq \kappa\}$  for some threshold  $\kappa \in \mathbb{R}$ .

The outer approximation requirement  $\bar{\mathcal{S}}_\gamma \supset \mathcal{S}_\gamma$  means that all true positive (i.e., 1) labels must be correctly classified, or in other words, the false negative (i.e., 0) rate is zero, i.e.,

$$P(X \in \bar{\mathcal{S}}_\gamma^c, Y = 1) = 0 \quad (2)$$

Typically, achieving such a zero ‘‘Type I’’ misclassification rate is impossible for any finite sample except in degenerate cases. However, this is achievable under a geometric premise on the rare-event set  $\mathcal{S}_\gamma$  that we call *orthogonal monotonicity*. To facilitate discussion, suppose from now on that the rare-event set is known to lie entirely in the positive quadrant  $\mathbb{R}_+^d$ , so in learning the set, we only consider sampling points in  $\mathbb{R}_+^d$  (analogous development can be extended to the entire space).

**Definition 4.** We call a set  $\mathcal{S} \subset \mathbb{R}_+^d$  orthogonally monotone if for any two points  $x, x' \in \mathbb{R}_+^d$ , we have  $x \leq x'$  (where the inequality is defined coordinate-wise) and  $x \in \mathcal{S}$  implies  $x' \in \mathcal{S}$  too.

Definition 4 means that any point that is more ‘‘extreme’’ than a point in the rare-event set must also lie inside the same set. This is an intuitive assumption that appears to hold in some safety-critical rare-event settings (see Section 4). Note that, even with such a monotonicity property, the boundary of the rare-event set can still be very complex. The key is that, with orthogonal monotonicity, we can now produce a classification procedure that satisfies 2. In fact, the simplest

approach is to use what we call an orthogonally monotone hull:

**Definition 5.** For a set of points  $D = \{x_1, \dots, x_n\} \subset \mathbb{R}_+^d$ , we define the orthogonally monotone hull of  $D$  (with respect to the origin) as  $\mathcal{H}(D) = \cup_i \mathcal{R}(x_i)$ , where  $\mathcal{R}(x_i)$  is the rectangle that contains both  $x_i$  and the origin as two of its corners.

In other words, the orthogonally monotone hull consists of the union of all the rectangles each wrapping each point  $x_i$  and the origin 0. Now, denote  $T_0 = \{\tilde{X}_i : Y_i = 0\}$  as the non-rare-event sampled points. Evidently, if  $\mathcal{S}_\gamma$  is orthogonally monotone, then  $\mathcal{H}(T_0) \subset \mathcal{S}_\gamma^c$  (where complement is with respect to  $\mathbb{R}_+^d$ ), or equivalently,  $\mathcal{H}(T_0)^c \supset \mathcal{S}_\gamma$ , i.e.,  $\mathcal{H}(T_0)^c$  is an outer approximation of the rare-event set  $\mathcal{S}_\gamma$ . Figure 1(b) shows this outer approximation (and also the inner counterpart). Moreover,  $\mathcal{H}(T_0)^c$  is the smallest region (in terms of set volume) such that 2 holds, because any smaller region could exclude a point that has label 1 with positive probability.

**Lazy-Learner IS.** We now consider an estimator for  $\mu$  where, given the  $n_1$  samples in Stage 1, we build the mixture IS depicted in Theorem 1 to estimate  $P(X \in \mathcal{H}(T_0)^c)$  in Stage 2. Since  $\mathcal{H}(T_0)^c$  takes the form  $(\cup_{i: Y_i=0} \mathcal{R}(\tilde{X}_i))^c$ , it has a finite number of dominating points, which can be found by a sequential algorithm (similar to the one that we will discuss momentarily). We call this the ‘‘lazy-learner’’ approach. Its problem, however, is that  $\mathcal{H}(T_0)^c$  tends to have a very rough boundary. This generates a large number of dominating points, many of which are unnecessary in that they do not correspond to any ‘‘true’’ dominating points in the original rare-event set  $\mathcal{S}_\gamma$  (see the middle of Figure 1(d)). This in turn leads to a large number of mixture components that degrades the IS efficiency, as the RE bound in Theorem 1 scales linearly with the number of mixture components.

**Deep-Learning-Based IS.** Our main approach is a deep-learning alternative that resolves the statistical degradation of the lazy learner. We train a neural network classifier, say  $\hat{g}$ , using all the Stage 1 samples  $\{(\tilde{X}_i, Y_i)\}$ , and obtain an approximate non-rare-event region  $(\bar{\mathcal{S}}_\gamma^\kappa)^c = \{x : \hat{g}(x) < \kappa\}$ , where  $\kappa$  is say 1/2. Then we adjust  $\kappa$  minimally away from 1/2, say to  $\hat{\kappa}$ , so that  $(\bar{\mathcal{S}}_\gamma^{\hat{\kappa}})^c \subset \mathcal{H}(T_0)$ , i.e.,  $\hat{\kappa} = \max\{\kappa \in \mathbb{R} : (\bar{\mathcal{S}}_\gamma^\kappa)^c \subset \mathcal{H}(T_0)\}$ . Then  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}} \supset \mathcal{H}(T_0)^c \supset \mathcal{S}_\gamma$ , so that  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  is an outer approximation for  $\mathcal{S}_\gamma$  (see Figure 1(c), where  $\hat{\kappa} = 0.68$ ). Stage 1 in Algorithm 1 shows this procedure. With this, we can run mixture IS to estimate  $P(X \in \bar{\mathcal{S}}_\gamma^{\hat{\kappa}})$  in Stage 2.

The execution of this algorithm requires the set  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}} = \{x : \hat{g}(x) \geq \hat{\kappa}\}$  to be in a form susceptible to The-

orem 1 and the search of all its dominating points. When  $\hat{g}$  is a ReLU-activated neural net, the boundary of  $\hat{g}(x) \geq \hat{\kappa}$  is piecewise linear and  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  is a union of polytopes, and Theorem 1 applies. Finding all dominating points is done by a sequential ‘‘cutting-plane’’ method that iteratively locates the next dominating point by minimizing  $(x - \mu)^T \Sigma^{-1} (x - \mu)$  over the remaining portion of  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  not covered by the local region of any previously found points  $x_j^*$ . These optimization sequences can be solved via mixed integer program (MIP) formulations for ReLU networks (Tjeng et al. 2017; Huang et al. 2018; see Appendix B). Note that a user can control the size of these MIPs via the neural net architecture. Regardless of the expressiveness of these networks, Algorithm 1 enjoys the following guarantee:

**Theorem 3** (Relaxed efficiency certificate for deep-learning-based mixture IS). *Suppose  $\mathcal{S}_\gamma$  is orthogonally monotone, and  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  satisfies the same conditions for  $\mathcal{S}_\gamma$  in Theorem 1. Then Algorithm 1 attains the upper-bound relaxed efficiency certificate by using a constant number of Stage 1 samples.*

Figure 1(d) shows how our deep-learning-based IS achieves superior efficiency compared to other alternatives. The cross-entropy method can miss a dominating point (1) and result in systematic under-estimation. The lazy-learner IS, on the other hand, generates too many dominating points (64) and degrades efficiency. Algorithm 1 finds the right number (2) and approximate locations of the dominating points.

Moreover, whereas the upper-bound certificate is guaranteed in our design, in practice, the deep-learning-based IS also appears to work well in controlling the conservativeness of the bound, as dictated by the false positive rate  $P(X \in \bar{\mathcal{S}}_\gamma^{\hat{\kappa}}, Y = 0)$  (see our experiments next). We close this section with a finite-sample bound on the false positive rate. Here, in deriving our bound, we assume the use of a sampling distribution  $q$  in generating independent Stage 1 samples, and we use empirical risk minimization (ERM) to train  $\hat{g}$ , i.e.,  $\hat{g} := \operatorname{argmin}_{g \in \mathcal{G}} \{R_{n_1}(g) := \frac{1}{n_1} \sum_{i=1}^{n_1} \ell(g(\tilde{X}_i), Y_i)\}$  where  $\ell$  is a loss function and  $\mathcal{G}$  is the considered hypothesis class. Correspondingly, let  $R(g) := E_{X \sim q} \ell(g(X), I(X \in \mathcal{S}_\gamma))$  be the true risk function and  $g^* := \operatorname{argmin}_{g \in \mathcal{G}} R(g)$  its minimizer. Also let  $\kappa^* := \min_{x \in \mathcal{S}_\gamma} g^*(x)$  be the true threshold associated with  $g^*$  in obtaining the smallest outer rare-event set approximation.

**Theorem 4** (Conservativeness). *Consider  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  obtained in Algorithm 1 where  $\hat{g}$  is trained from an ERM. Suppose the density  $q$  has bounded support  $K \subset [0, M]^d$  and  $0 < q_l \leq q(x) \leq q_u$  for any  $x \in K$ . Also suppose there exists a function  $h$  such that for any  $g \in \mathcal{G}$ ,  $g(x) \geq \kappa$  implies  $\ell(g(x), 0) \geq h(\kappa) > 0$ . (e.g., if  $\ell$  is the*

*squared loss, then  $h(\kappa)$  could be chosen as  $h(\kappa) = \kappa^2$ ). Then, with probability at least  $1 - \delta$ ,*

$$P_{X \sim q} (X \in \bar{\mathcal{S}}_\gamma^{\hat{\kappa}} \setminus \mathcal{S}_\gamma) \leq \frac{R(g^*) + 2 \sup_{g \in \mathcal{G}} |R_{n_1}(g) - R(g)|}{h(\kappa^* - t(\delta, n_1) \sqrt{d} \operatorname{Lip}(g^*) - \|\hat{g} - g^*\|_\infty)}.$$

Here,  $\operatorname{Lip}(g^*)$  is the Lipschitz parameter of  $g^*$ , and  $t(\delta, n_1) = 3 \left( \frac{\log(n_1 q_l) + d \log M + \log \frac{1}{\delta}}{n_1 q_l} \right)^{\frac{1}{d}}$ .

Theorem 4 reveals a tradeoff between overfitting (measured by  $\sup_{g \in \mathcal{G}} |R_{n_1}(g) - R(g)|$  and  $\|\hat{g} - g^*\|_\infty$ ) and underfitting (measured by  $R(g^*) = \inf_{g \in \mathcal{G}} R(g)$ ). Appendix C discusses related results on the sharp estimates of these quantities for deep neural networks, a more sophisticated version of Theorem 4 that applies to the cross-entropy loss, a corresponding bound for the lazy learner, as well as results to interpret Theorem 4 under the original distribution  $p$ .


Finally, we point out some works in the literature that approximate the Pareto frontier of a monotone function (Wu et al. 2018; Legriel et al. 2010). While the boundary of an orthogonally monotone set looks similar to the Pareto frontier, our recipe (outer/inner approximation using piecewise-linear-activation NN) is designed to minimize the number of dominating points while simultaneously achieve the relaxed efficiency certificate for rare-event estimation. Such a guarantee is novelly beyond these previous works.

## 4 Numerical Experiments


We implement and compare the estimated probabilities and the REs of deep-learning-based IS for the upper bound (Deep-PrAE UB) and lazy-learner IS (LL UB). We also show the corresponding lower-bound estimator (Deep-PrAE LB and LL LB) and benchmark with the cross entropy method (CE), adaptive multilevel splitting (AMS), and naive Monte Carlo (NMC). For CE, we run a few variations testing different parametric classes and report two in the following figures: CE that uses a single Gaussian distribution (CE Naive), representing an overly-simplified CE implementation, and CE with Gaussian Mixture Model with  $k$  components (CE GMM- $k$ ), representing a more sophisticated CE implementation. For Deep-PrAE, we use the samples from CE Naive as the Stage 1 samples. We also run a modification of Deep-PrAE (Deep-PrAE Mod) that replaces  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  by  $\mathcal{S}_\gamma$  in the last step of Algorithm 1 as an additional comparison. We use a 2-dimensional numerical example and a safety-testing of an intelligent driving algorithm for a car-following scenario. These two experiments are representative as the former is low-dimensional (visualizable) yet with *extremely* rare

events while the latter is moderately high-dimensional, challenging for most of the existing methods.

**2D Example.** We estimate  $\mu = P(X \in \mathcal{S}_\gamma)$  where  $X \sim N([5, 5]^T, 0.25I_{2 \times 2})$ , and  $\gamma$  ranging from 1.0 to 2.0. We use  $n = 30,000$  (10,000 for Stage 1 and 20,000 for Stage 2), and use 10,000 of the CE samples as our Stage 1 samples. Figure 1 illustrates the shape of  $\mathcal{S}_\gamma$ , which has two dominating points. This probability is microscopically small (e.g.,  $\gamma = 1.8$  gives  $\mu = 4.1 \times 10^{-24}$ ) and serves to investigate our performance in ultra-extreme situations.



(a) Estimated rare-event probability



(b) Estimator’s empirical relative error

Figure 2: 2-dimensional example. Naive Monte Carlo failed in all cases and hence not shown.




Figure 3: Percentage error of CE, AMS, and modified Deep-PrAE (minus % error means under-estimation)

Figure 2 compares all approaches to the true value, which we compute via a proper mixture IS with 50,000

samples assuming the full knowledge of  $\mathcal{S}_\gamma$ . It shows several observations. First, Deep-PrAE and LL (both UB and LB) always provide valid bounds that contain the truth. Second, the UB for LL is more conservative than Deep-PrAE in up to two orders of magnitudes, which is attributed to the overly many (redundant) dominating points. Correspondingly, the RE of LL UB is tremendously high, reaching over 500% when  $\gamma = 2.0$ , compared to around 40% for Deep-PrAE UB, Deep-PrAE LB, and LL LB. Third, CE Naive, which finds only one dominating point, consistently under-estimates the truth by about 50%, yet it gives an over-confident RE, e.g.,  $< 5\%$  when  $\gamma < 2$ . This shows a systematic undetected under-estimation issue when CE is implemented overly-naively. AMS also underestimates the true value by 30%-40%, while CE GMM-2 and Deep-PrAE Mod perform empirically well. Figure 3 summarizes the zoomed-in performances of CE Naive, CE GMM-2, AMS, and Deep-PrAE Mod in terms of percentage error, which is the difference between the estimated and true probability as a percentage of the true value. It shows that while CE performs well when the IS parametric class is well-chosen (CE GMM-2), a poor CE parametric class (CE Naive) as well as AMS could under-estimate. Yet our Deep-PrAE, despite using samples from a poor CE class in Stage 1, can recover valid results: Deep-PrAE provides a valid UB, and Deep-PrAE Mod gives an estimate as good as the good CE class.

**Intelligent Driving Example.** In this example, we evaluate the crash probability of car-following scenarios involving a human-driven lead vehicle (LV) followed by an autonomous vehicle (AV). The AV is controlled by the Intelligent Driver Model (IDM) to maintain safety distance while ensuring smooth ride and maximum efficiency. IDM model is widely used for autonomy evaluation and microscopic transportation simulations (Treiber et al., 2000; Wang et al., 2018; Orzechowski et al., 2019). The state at time  $t$  is given by 6 states consisting of the position, velocity, and acceleration of both LV and AV. The dynamic system has a stochastic input  $u_t$  related to the acceleration of the LV and subject to uncertain human behavior. We consider an evaluation horizon  $T = 60$  seconds and draw a sequence of 15 Gaussian random actions at a 4-second epoch, leading to a 15-dimensional LV action space. A (rare-event) crash occurs at time  $t \leq T$  if the longitudinal distance  $r_t$  between the two vehicles is negative, with  $\gamma$  parameterizing the AV maximum throttle and brake pedals. This rare-event set is analytically challenging (see Zhao et al., 2017 for a similar setting). More details are in Appendix F.

Figure 4 shows the performances of competing approaches, using  $n = 10,000$ . For CE, we use a single






Figure 4: Intelligent driving example. Naive Monte Carlo failed in all cases and hence not shown.

Gaussian (CE Naive) and a large number of mixtures (CE GMM-64). Deep-PrAE and LL (UB and LB) appear consistent in giving upper and lower bounds for the target probability, and Deep-PrAE produces tighter bounds than LL ( $10^{-2}$  vs  $10^{-6}$  in general). LL UB has 5,644 dominating points when  $\gamma = 1$  vs 42 in Deep-PrAE, and needs 4 times more time to search for them than Deep-PrAE. Moreover, the RE of Deep-PrAE is 3 times lower than LL across the range (in both UB and LB). Thus, Deep-PrAE outperforms LL in both tightness, computation speed, and RE. CE Naive and AMS seem to give a stable estimation, but evidence shows they are under-estimated: Deep-PrAE Mod and CE GMM-64 lack efficiency certificates and thus could under-estimate, and the fact that their estimates are higher than CE Naive and AMS suggests both CE Naive and AMS are under-estimated. Lastly, NMC fails to give a single hit in all cases (thus not shown on the graphs). Thus, among all approaches, Deep-PrAE stands out as giving reliable and tight bounds with low RE.

**Summary of Practical Benefits.** Our investigation shows strong evidence of the practical benefits of using Deep-PrAE for rare-event estimation. It generates valid bounds for the target probability with low RE and improved efficiency. The use of classifier prediction helps reduce the computational effort from running more simulations. For example, to assess whether the AV crash rate is below  $10^{-8}$  for  $\gamma = 1.0$ , only 1000

simulation runs would be needed by Deep-PrAE UB or LB to get around 1% RE, which takes about 400 seconds in total. This is in contrast to 3.7 months for naive Monte Carlo.

## 5 Discussion and Future Work

In this paper, we proposed a robust certifiable approach to estimate rare-event probabilities in safety-critical applications. The proposed approach designs efficient IS distribution by combining the dominating point machinery with deep-learning-based rare-event set learning. We study the theoretical guarantees and present numerical examples. The key property that distinguishes our approach with existing black-box rare-event simulation methods is our correctness guarantee. Leveraging on a new notion of relaxed efficiency certificate and the orthogonal monotonicity assumption, our approach avoids the perils of undetected under-estimation as potentially encountered by other methods.

We discuss some key assumptions in our approach and related prospective follow-up works. First, the orthogonal monotonicity assumption appears an important first step to give new theories on black-box rare-event estimation beyond the existing literature. Indeed, we show that even with this assumption, black-box approaches such as CE and splitting can suffer from the dangerous pitfall of undiagnosed under-estimation, and our approach corrects for it. The real-world values of our approach are: (1) We rigorously show why our method has better performances in the orthogonally monotone cases; (2) For tasks close to being orthogonal monotone (e.g., the IDM example), our method is empirically more robust; (3) For non-orthogonally-monotone tasks, though directly using our approach does not provide guarantees, we could potentially train mappings to latent spaces that are orthogonally monotone. We believe such type of geometric assumptions comprises a key ingredient towards a rigorous theory for black-box rare-event estimation that warrants much further developments.

Second, the dominating point search algorithm in our approach assumes Gaussian randomness. To this end we can relax it in two directions: by fitting a GMM with a sufficiently large number of components, or using light-tailed distributions (i.e., with finite exponential moments), since the dominating point machinery applies. These extensions will be left for future work.

Finally, the tightness of the upper bound depends on the sample quality. An ideal method in Stage 1 would generate samples close to the rare-event boundary to produce good approximations. Cutting the Stage 1 effort by, e.g., designing iterative schemes between Stages 1 and 2, will also be a topic for future investigation.

## Acknowledgments

We gratefully acknowledge support from the National Science Foundation under grants CAREER CMMI-1834710, IIS-1849280 and IIS-1849304. Mansur Arief and Wenhao Ding are supported in part by Bosch.

## References

- Amir Dembo, O. Z. (2010). *Large Deviations Techniques and Applications*. Springer-Verlag.
- Anil, C., Lucas, J., and Grosse, R. (2019). Sorting out Lipschitz function approximation. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 291–301, Long Beach, California, USA. PMLR.
- Arief, M., Glynn, P., and Zhao, D. (2018). An accelerated approach to safely and efficiently test pre-production autonomous vehicles on public streets. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2006–2011. IEEE.
- Arouna, B. (2004). Adaptive monte carlo method, a variance reduction technique. *Monte Carlo Methods and Applications*, 10(1):1–24.
- Asmussen, S. and Glynn, P. W. (2007). *Rare-Event Simulation*, pages 158–205. Springer New York, New York, NY.
- Asmussen, S. and Kroese, D. P. (2006). Improved algorithms for rare event simulation with heavy tails. *Advances in Applied Probability*, 38(2):545–558.
- Au, S.-K. and Beck, J. L. (2001). Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic Engineering Mechanics*, 16(4):263–277.
- Blanchet, J. and Lam, H. (2012). State-dependent importance sampling for rare-event simulation: An overview and recent advances. *Surveys in Operations Research and Management Science*, 17(1):38–59.
- Botev, Z. I., L’Ecuyer, P., and Tuffin, B. (2013). Markov chain importance sampling with applications to rare event probability estimation. *Statistics and Computing*, 23(2):271–285.
- Botev, Z. I., Ridder, A., and Rojas-Nandayapa, L. (2016). Semiparametric cross entropy for rare-event simulation. *Journal of Applied Probability*, 53(3):633–649.
- Bucklew, J. (2013). *Introduction to Rare Event Simulation*. Springer Science & Business Media.
- Bucklew, J. A. (2004). *Rare Event Simulation for Level Crossing and Queueing Models*, pages 195–206. Springer New York, New York, NY.
- Cao, Y. and Gu, Q. (2019). Tight sample complexity of learning one-hidden-layer convolutional neural networks. In *Advances in Neural Information Processing Systems 32*, pages 10612–10622. Curran Associates, Inc.
- C erou, F. and Guyader, A. (2007). Adaptive multilevel splitting for rare event analysis. *Stochastic Analysis and Applications*, 25(2):417–443.
- C erou, F., Guyader, A., and Rousset, M. (2019). Adaptive multilevel splitting: Historical perspective and recent results. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(4):043108.
- Chen, B., Blanchet, J., Rhee, C.-H., and Zwart, B. (2019). Efficient rare-event simulation for multiple jump events in regularly varying random walks and compound poisson processes. *Mathematics of Operations Research*, 44(3):919–942.
- Clarke, E. M., Henzinger, T. A., Veith, H., and Bloem, R. (2018). *Handbook of Model Checking*, volume 10. Springer.
- Claybrook, J. and Kildare, S. (2018). Autonomous vehicles: No driver... no regulation? *Science*, 361(6397):36–37.
- Corso, A., Moss, R. J., Koren, M., Lee, R., and Kochenderfer, M. J. (2020). A survey of algorithms for black-box safety validation. *arXiv preprint arXiv:2005.02979*.
- C erou, F. and Guyader, A. (2016). Fluctuation analysis of adaptive multilevel splitting. *The Annals of Applied Probability*, 26(6):3319 – 3380.
- De Boer, P.-T., Kroese, D. P., Mannor, S., and Rubinstein, R. Y. (2005). A tutorial on the cross-entropy method. *Annals of Operations Research*, 134(1):19–67.
- Dean, T. and Dupuis, P. (2009). Splitting for rare event simulation: A large deviation approach to design and analysis. *Stochastic Processes and Their Applications*, 119(2):562–587.
- Dieker, A. B. and Mandjes, M. (2006). Fast simulation of overflow probabilities in a queue with gaussian input. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 16(2):119–151.
- Dupuis, P. and Ellis, R. S. (2011). *A Weak Convergence Approach to the Theory of Large Deviations*, volume 902. John Wiley & Sons.
- Ellis, R. S. (1984). Large deviations for a general class of random vectors. *Ann. Probab.*, 12(1):1–12.

- Evan, A. (2016). Fatal Tesla Self-Driving Car Crash Reminds Us That Robots Aren't Perfect. *IEEE Spectrum*.
- Glasserman, P., Heidelberger, P., Shahabuddin, P., and Zajic, T. (1999). Multilevel splitting for estimating rare event probabilities. *Operations Research*, 47(4):585–600.
- Glynn, P. W. and Iglehart, D. L. (1989). Importance sampling for stochastic simulations. *Management Science*, 35(11):1367–1392.
- Grace, A. W., Kroese, D. P., and Sandmann, W. (2014). Automated state-dependent importance sampling for markov jump processes via sampling from the zero-variance distribution. *Journal of Applied Probability*, 51(3):741–755.
- Guyader, A., Hengartner, N., and Matzner-Løber, E. (2011). Simulation and estimation of extreme quantiles and extreme probabilities. *Applied Mathematics & Optimization*, 64(2):171–196.
- Gärtner, J. (1977). On large deviations from the invariant measure. *Theory of Probability & Its Applications*, 22(1):24–39.
- Harvey, N., Liaw, C., and Mehrabian, A. (2017). Nearly-tight VC-dimension bounds for piecewise linear neural networks. In Kale, S. and Shamir, O., editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 1064–1068, Amsterdam, Netherlands. PMLR.
- Huang, Z., Lam, H., LeBlanc, D. J., and Zhao, D. (2018). Accelerated evaluation of automated vehicles using piecewise mixture models. *IEEE Transactions on Intelligent Transportation Systems*, 19(9):2845–2855.
- Huang, Z., Lam, H., and Zhao, D. (2018). Designing importance samplers to simulate machine learning predictors via optimization. In *2018 Winter Simulation Conference (WSC)*, pages 1730–1741. IEEE.
- Juneja, S. and Shahabuddin, P. (2006). Rare-event simulation techniques: An introduction and recent advances. *Handbooks in Operations Research and Management Science*, 13:291–350.
- Kalra, N. and Paddock, S. M. (2016). Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A: Policy and Practice*, 94:182–193.
- Koopman, P. and Wagner, M. (2017). Autonomous vehicle safety: An interdisciplinary challenge. *IEEE Intelligent Transportation Systems Magazine*, 9(1):90–96.
- Koopman, P. and Wagner, M. (2018). Toward a framework for highly automated vehicle safety validation. Technical report, SAE Technical Paper.
- Koren, M., Alsaif, S., Lee, R., and Kochenderfer, M. J. (2018). Adaptive stress testing for autonomous vehicles. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1–7. IEEE.
- L'ecuyer, P., Blanchet, J. H., Tuffin, B., and Glynn, P. W. (2010). Asymptotic robustness of estimators in rare-event simulation. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 20(1):1–41.
- Legriel, J., Le Guernic, C., Cotton, S., and Maler, O. (2010). Approximating the pareto front of multi-criteria optimization problems. In Esparza, J. and Majumdar, R., editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 69–83, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Lu, Z., Pu, H., Wang, F., Hu, Z., and Wang, L. (2017). The expressive power of neural networks: A view from the width. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 6231–6239. Curran Associates, Inc.
- Müller, T., McWilliams, B., Rousselle, F., Gross, M., and Novák, J. (2019). Neural importance sampling. *ACM Trans. Graph.*, 38(5).
- NHTSA (2007). The new car assessment program suggested approaches for future program enhancements. *DOT HS*, 810:698.
- NTSB (2016). Preliminary Report, Highway HWY16FH018.
- O'Kelly, M., Sinha, A., Namkoong, H., Tedrake, R., and Duchi, J. C. (2018). Scalable end-to-end autonomous vehicle testing via rare-event simulation. In *Advances in Neural Information Processing Systems*, pages 9827–9838.
- Orzechowski, P. F., Li, K., and Lauer, M. (2019). Towards responsibility-sensitive safety of automated vehicles with reachable set analysis. In *2019 IEEE International Conference on Connected Vehicles and Expo (ICCVE)*, pages 1–6.
- Rubinstein, R. Y. (2005). A stochastic minimum cross-entropy method for combinatorial optimization and rare-event estimation. *Methodology and Computing in Applied Probability*, 7(1):5–50.
- Rubinstein, R. Y. and Kroese, D. P. (2013). *The Cross-entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning*. Springer Science & Business Media.

- Sadowsky, J. S. and Bucklew, J. A. (1990). On large deviations theory and asymptotically efficient monte carlo estimation. *IEEE Transactions on Information Theory*, 36(3):579–588.
- Siegmund, D. (1976). Importance Sampling in the Monte Carlo Study of Sequential Tests. *The Annals of Statistics*, 4(4):673 – 684.
- Tjeng, V., Xiao, K., and Tedrake, R. (2017). Evaluating robustness of neural networks with mixed integer programming. *arXiv preprint arXiv:1711.07356*.
- Treiber, M., Hennecke, A., and Helbing, D. (2000). Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2):1805–1824.
- Tuffin, B. and Ridder, A. (2012). Probabilistic bounded relative error for rare event simulation learning techniques. In *Proceedings of the 2012 Winter Simulation Conference (WSC)*, pages 1–12. IEEE.
- Uesato, J., Kumar, A., Szepesvari, C., Erez, T., Ruderman, A., Anderson, K., Heess, N., and Kohli, P. (2018). Rigorous agent evaluation: An adversarial approach to uncover catastrophic failures. *arXiv preprint arXiv:1812.01647*.
- van der Vaart, A. W. and Wellner, J. A. (1996). Weak convergence and empirical processes. *Springer Series in Statistics*.
- Villén-Altamirano, M. and Villén-Altamirano, J. (1994). Restart: a straightforward method for fast simulation of rare events. In *Proceedings of Winter Simulation Conference*, pages 282–289. IEEE.
- Villén-Altamirano, J. (2010). Importance functions for restart simulation of general jackson networks. *European Journal of Operational Research*, 203(1):156–165.
- Wang, X., Jiang, R., Li, L., Lin, Y., Zheng, X., and Wang, F. (2018). Capturing car-following behaviors by deep learning. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):910–920.
- Webb, S., Rainforth, T., Teh, Y. W., and Kumar, M. P. (2018). A statistical approach to assessing neural network robustness. *arXiv preprint arXiv:1811.07209*.
- Wegener, J. and Bühler, O. (2004). Evaluation of different fitness functions for the evolutionary testing of an autonomous parking system. In *Genetic and Evolutionary Computation Conference*, pages 1400–1412. Springer.
- Wu, X., Gomes-Selman, J., Shi, Q., Xue, Y., García-Villacorta, R., Anderson, E., Sethi, S., Steinschneider, S., Flecker, A., and Gomes, C. P. (2018). Efficiently approximating the pareto frontier: Hydropower dam placement in the amazon basin. In *AAAI*.
- Zhao, D., Huang, X., Peng, H., Lam, H., and LeBlanc, D. J. (2017). Accelerated evaluation of automated vehicles in car-following maneuvers. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):733–744.
- Zhao, D., Lam, H., Peng, H., Bao, S., LeBlanc, D. J., Nobukawa, K., and Pan, C. S. (2016). Accelerated evaluation of automated vehicles safety in lane-change scenarios based on importance sampling techniques. *IEEE Transactions on Intelligent Transportation Systems*, 18(3):595–607.

# Appendices

We present supplemental results and discussions. Appendix [A](#) expands Section [2](#) regarding Monte Carlo efficiency and variance reduction. Appendix [B](#) provides further details on Algorithm [1](#), in particular the mixed integer formulation used to solve the underlying optimization problems. Appendix [C](#) expands the efficiency and conservativeness results in Section [3](#). Appendix [D](#) presents the lower-bound relaxed efficiency certificate and estimators in parallel to the upper-bound results in Section [3](#). Appendix [E](#) provides an overview of the cross-entropy method and multi-level splitting (or subset simulation) and discusses their perils for black-box problems. Appendix [F](#) illustrates further experimental results. Finally, Appendix [G](#) shows all technical proofs.

## A Further Details for Section [2](#)

This section expands the discussions in Section [2](#) by explaining in more detail the notion of relative error, challenges in naive Monte Carlo, the concept of dominating points, and the perils of black-box variance reduction algorithms.

### A.1 Explanation of the Role of Relative Error

As described in Section [2](#), to estimate a tiny  $\mu$  using  $\hat{\mu}_n$ , we want to ensure a high accuracy in relative term, namely, [\(1\)](#). Suppose that  $\hat{\mu}_n$  is unbiased and is an average of  $n$  i.i.d. simulation runs, i.e.,  $\hat{\mu}_n = (1/n) \sum_{i=1}^n Z_i$  for some random unbiased output  $Z_i$ . The Markov inequality gives that

$$P(|\hat{\mu}_n - \mu| > \epsilon\mu) \leq \frac{\text{Var}(\hat{\mu}_n)}{\epsilon^2\mu^2} = \frac{\text{Var}(Z_i)}{n\epsilon^2\mu^2}$$

so that

$$\frac{\text{Var}(Z_i)}{n\epsilon^2\mu^2} \leq \delta$$

ensures [\(1\)](#). Equivalently,

$$n \geq \frac{\text{Var}(Z_i)}{\delta\epsilon^2\mu^2} = \frac{RE}{\delta\epsilon^2}$$

is a sufficient condition to achieve [\(1\)](#), where  $RE = \text{Var}(Z_i)/\mu^2$  is the relative error defined as the ratio of variance (per-run) and squared mean.

Note that replacing the second  $\mu$  with  $\hat{\mu}_n$  in the left hand side of [\(1\)](#) does not change the condition fundamentally, as either is equivalent to saying the ratio  $\hat{\mu}_n/\mu$  should be close to 1. Also, note that we focus on the nontrivial case that the target probability  $\mu$  is non-zero but tiny; if  $\mu = 0$ , then no non-zero Monte Carlo estimator can achieve a good relative error.

### A.2 Further Explanation on the Challenges in Naive Monte Carlo

We have seen in Section [2](#) that for the naive Monte Carlo estimator, where  $Z_i = I(X_i \in \mathcal{S}_\gamma)$ , the relative error is  $RE = \mu(1 - \mu)/\mu^2 = (1 - \mu)/\mu$ . Thus, when  $\mu$  is tiny, the sufficient condition for  $n$  to attain [\(1\)](#) scales at least linearly in  $1/\mu$ . In fact, this result can be seen to be tight by analyzing  $n\hat{\mu}_n$  as a binomial variable. To be more specific, we know that  $P(|\hat{\mu}_n - \mu| > \epsilon\mu) = P(|n\hat{\mu}_n - n\mu| > \epsilon n\mu)$  and that  $n\hat{\mu}_n$  takes values in  $\{0, 1, \dots, n\}$ . Therefore, if  $n\mu \rightarrow 0$ , then  $P(|\hat{\mu}_n - \mu| > \epsilon\mu) \rightarrow 1$ , and hence [\(1\)](#) does not hold.

Moreover, the following provides a concrete general statement that an  $n$  that grows only polynomially in  $\gamma$  would fail to estimate  $\mu$  that decays exponentially in  $\gamma$  with enough relative accuracy, of which [\(1\)](#) fails to hold is an implication.

**Proposition 2.** *Suppose that  $\mu = P(X \in \mathcal{S}_\gamma)$  is exponentially decaying in  $\gamma$  and  $n$  is polynomially growing in  $\gamma$ . Define  $\hat{\mu}_n = (1/n) \sum_{i=1}^n I(X_i \in \mathcal{S}_\gamma)$ . Then for any  $0 < \epsilon < 1$ ,*

$$\lim_{\gamma \rightarrow \infty} P(|\hat{\mu}_n - \mu| > \epsilon\mu) = 1.$$

We have used the term efficiency certificate to denote an estimator that achieves [\(1\)](#) with  $n = \tilde{O}(\log(1/\mu))$ . In the rare-event literature, such an estimator is known as “logarithmically efficient” or “weakly efficient” [\(Juneja and Shahabuddin, 2006; Blanchet and Lam, 2012\)](#).

### A.3 Further Explanations of Dominating Points

We have mentioned that a certifiable IS should account for all dominating points, defined in [Definition 2](#). We provide more detailed explanations here. Roughly speaking, for  $X \sim N(\lambda, \Sigma)$  and a rare-event set  $\mathcal{S}_\gamma$ , the Laplace approximation gives  $P(X \in \mathcal{S}_\gamma) \approx e^{-\inf_{a \in \mathcal{S}_\gamma} \frac{1}{2}(a-\lambda)^T \Sigma^{-1}(a-\lambda)}$  (see the proof of [Theorem 5](#)). Thus, to obtain an efficiency certificate, IS estimator given by  $Z = L(X)I(X \in \mathcal{S}_\gamma)$ , where  $X \sim \tilde{p}$  and  $L = dp/d\tilde{p}$ , needs to have  $\widetilde{Var}(Z) \leq \tilde{E}[Z^2] \approx e^{-\inf_{a \in \mathcal{S}_\gamma} (a-\lambda)^T \Sigma^{-1}(a-\lambda)}$  (where  $\widetilde{Var}(\cdot)$  and  $\tilde{E}[\cdot]$  denote the variance and expectation under  $\tilde{p}$ , and  $\approx$  is up to some factor polynomial in  $\inf_{a \in \mathcal{S}_\gamma} (a-\lambda)^T \Sigma^{-1}(a-\lambda)$ ; note that the last equality relation cannot be improved, as otherwise it would imply that  $\widetilde{Var}(Z) = \tilde{E}[Z^2] - (\tilde{E}[Z])^2 < 0$ ).

Now consider an IS that translates the mean of the distribution from  $\mu$  to  $a^* = \operatorname{argmin}_{a \in \mathcal{S}_\gamma} (a-\lambda)^T \Sigma^{-1}(a-\lambda)$ , an intuitive choice since  $a^*$  contributes the highest density among all points in  $\mathcal{S}_\gamma$  (this mean translation also bears the natural interpretation as an exponential change of measure; [\(Bucklew, 2013\)](#)). The likelihood ratio is  $L(x) = e^{(\mu-a^*)^T \Sigma^{-1}(x-\lambda) + \frac{1}{2}(\lambda-a^*)^T \Sigma^{-1}(\lambda-a^*)}$ , giving

$$\tilde{E}[Z^2] = \tilde{E}[L(X)^2 I(X \in \mathcal{S}_\gamma)] = e^{-(a^*-\lambda)^T \Sigma^{-1}(a^*-\lambda)} \tilde{E}[e^{-2(a^*-\lambda)^T \Sigma^{-1}(x-a^*)} I(X \in \mathcal{S}_\gamma)] \quad (3)$$

If the “overshoot”  $(a^*-\lambda)^T \Sigma^{-1}(x-a^*)$ , i.e., the remaining term in the exponent of  $L(x)$  after moving out  $-(a^*-\lambda)^T \Sigma^{-1}(a^*-\lambda)$ , satisfies  $(a^*-\lambda)^T \Sigma^{-1}(x-a^*) \geq 0$  for all  $x \in \mathcal{S}_\gamma$ , then the expectation in the right hand side of [\(3\)](#) is bounded by 1, and an efficiency certificate is achieved. This, however, is not true for all set  $\mathcal{S}_\gamma$ , which motivates the following definition of the dominant set and points in [Definition 2](#).

For instance, if  $\mathcal{S}_\gamma$  is convex, then, noting that  $(x-\lambda)^T \Sigma^{-1}$  is precisely the gradient of the function  $(1/2)(x-\lambda)^T \Sigma^{-1}(x-\lambda)$ , we get that  $a^*$  gives a singleton dominant set since  $(a^*-\lambda)^T \Sigma^{-1}(x-a^*) \geq 0$  for all  $x \in \mathcal{S}_\gamma$  is precisely the first order optimality condition of the involved quadratic optimization. In general, if we can decompose  $\mathcal{S}_\gamma = \bigcup_j \mathcal{S}_\gamma^j$  where  $\mathcal{S}_\gamma^j = \{x : (a_j - \lambda)^T \Sigma^{-1}(x - a_j) \geq 0\}$  for a dominating point  $a_j \in A_\gamma$ , then each  $\mathcal{S}_\gamma^j$  can be viewed as a “local” region where the dominating point  $a_j$  is the highest-density, or the most likely point such that the rare event occurs.

The following is the detailed version of [Theorem 1](#).

**Theorem 5** (Certifiable IS). *Suppose that  $A_\gamma$  is the dominant set for  $\mathcal{S}_\gamma$  associated with the distribution  $N(\lambda, \Sigma)$ . Then we can decompose  $\mathcal{S}_\gamma = \bigcup_j \mathcal{S}_\gamma^j$  where  $\mathcal{S}_\gamma^j$ ’s are disjoint,  $a_j \in \mathcal{S}_\gamma^j$  and  $\mathcal{S}_\gamma^j \subset \{x : (a_j - \lambda)^T \Sigma^{-1}(x - a_j) \geq 0\}$  for  $a_j \in A_\gamma$ . Denote  $a^* = \operatorname{argmin}\{(a_j - \lambda)^T \Sigma^{-1}(a_j - \lambda) : a_j \in A_\gamma\}$ . Assume that each component of  $a^*$  is of polynomial growth in  $\gamma$ . Moreover, assume that there exist invertible matrix  $B$  and positive constant  $\varepsilon$  such that  $\{x : B(x - a^*) \geq 0, (x - a^*)^T \Sigma^{-1}(x - a^*) \leq \varepsilon^2\} \subset \mathcal{S}_\gamma$ . Then the IS distribution  $\sum_j \alpha_j N(a_j, \Sigma)$  achieves an efficiency certificate in estimating  $\mu = P(X \in \mathcal{S}_\gamma)$ , i.e., if we let  $Z = I(X \in \mathcal{S}_\gamma)L(X)$  where  $L$  is the corresponding likelihood ratio, then  $\tilde{E}[Z^2]/\tilde{E}[Z]^2$  is at most polynomially growing in  $\gamma$ . This applies in particular to  $\mathcal{S}_\gamma = \{x : f(x) \geq \gamma\}$  where  $f(x)$  is a piecewise linear function.*

We contrast [Theorem 5](#) with existing works on dominating points. The latter machinery has been studied in [\(Sadowsky and Bucklew, 1990; Dieker and Mandjes, 2006\)](#). These papers, however, consider regimes where the Gärtner-Ellis Theorem [\(Gärtner, 1977; Ellis, 1984\)](#) can be applied, which requires the considered rare-event set to scale proportionately with the rarity parameter. This is in contrast to the general conditions on the dominating points used in [Theorem 5](#).

### A.4 Further Explanation of the Example in [Theorem 2](#)

In the theorem, there are two dominating points  $\gamma$  and  $-\kappa\gamma$  but the IS design only considers the first one. As a result, there could exist “unlucky” scenario where the sample falls into the rare-event set, so that  $I(X \in \mathcal{S}_\gamma) = 1$ , while the likelihood ratio  $L(X)$  explodes, which leads to a tremendous estimation variance. Part 2 of the theorem further shows how this issue is undetected empirically, as the empirical RE appears small (polynomially in  $n$  and hence  $\gamma$  by our choice of  $n$ ) while the estimation concentrates at a value that can be severely under the correct

one (especially when  $k < 1$ ). This is because the samples all land on the neighborhood of the solely considered dominating point. If the missed dominating point is a significant contributor to the rare-event probability, then the empirical performance would look as if the rare-event set is *smaller*, leading to a systematic under-estimation. Note that this phenomenon occurs even if the estimator is unbiased, which is guaranteed by IS by default.

## B Further Details on Implementing Algorithm 1

We provide further details on implementing Algorithm 1. In particular, we present how to solve the optimization problem

$$x^* = \arg \min_x (x - \lambda)^T \Sigma^{-1} (x - \lambda) \quad \text{s.t.} \quad \hat{g}(x) \geq \hat{\kappa}, \quad (x_j^* - \lambda)^T \Sigma^{-1} (x - x_j^*) < 0 \quad \forall x_j^* \in \hat{A}_\gamma \quad (4)$$

to obtain the next dominating point in the sequential cutting-plane approach in Stage 2. Moreover, we also present how to tune

$$\hat{\kappa} = \max\{\kappa \in \mathbb{R} : (\bar{\mathcal{S}}_\gamma^\kappa)^c \subset \mathcal{H}(T_0)\} \quad (5)$$

in Stage 1.

**MIP formulations for ReLU-activated neural net classifier.** The problem (4) can be reformulated into a mixed integer program (MIP), in the case where  $\hat{g}(x)$  is trained via a ReLU-activated neural net classifier, which is used in our deep-learning-based IS. Since the objective is convex quadratic and second set of constraints is linear in (4), we focus on the first constraint  $\hat{g}(x) \geq \gamma$ . The neural net structure  $\hat{g}(x)$  in our approach (say with  $n_g$  layers) can be represented as  $\hat{g}(x) = (\hat{g}_{n_g} \circ \dots \circ \hat{g}_1)(x)$ , where each  $\hat{g}_i(\cdot)$  denotes a ReLU-activated layer with linear transformation, i.e.  $\hat{g}_i(\cdot) = \max\{LT(\cdot), 0\}$ , where  $LT(\cdot)$  denotes a certain linear transformation in the input. In order to convert  $\hat{g}(\cdot)$  into an MIP constraint, we introduce  $M$  as a practical upper bound for  $x_1, \dots, x_n$  such that  $|x_i| < M$ . The key step is to reformulate the ReLU function  $y = \max\{x, 0\}$  into

$$\begin{aligned} y &\leq x + M(1 - z) \\ y &\geq x \\ y &\leq Mz \\ y &\geq 0 \\ z &\in \{0, 1\}. \end{aligned}$$

For simple ReLU networks, the size of the resulting MIP formulation depends linearly on the number of neurons in the neural network. In particular, the number of binary decision variables is linearly dependent on the number of ReLU neurons, and the number of constraints is linearly dependent the total number of all neurons (here we consider the linear transformations as independent neurons).

The MIP reformulation we discussed can be generalized to many other popular piecewise linear structures in deep learning. For instance, linear operation layers, such as normalization and convolutional layers, can be directly used as constraints; some non-linear layers, such as ReLU and max-pooling layers, introduce non-linearity by the “max” functions. A general reformulation for the max functions can be used to convert these non-linear layers to mixed integer constraints.

Consider the following equality defined by a max operation  $y = \max\{x_1, x_2, \dots, x_n\}$ . Then the equality is equivalent to

$$\begin{aligned} y &\leq x_i + 2M(1 - z_i), i = 1, \dots, n \\ y &\geq x_i, i = 1, \dots, n \\ \sum_{i=1, \dots, n} z_i &= 1 \\ z_i &\in \{0, 1\}. \end{aligned}$$

**Tuning  $\hat{\kappa}$ .** We illustrate how to tune  $\hat{\kappa}$  to achieve (5). This requires checking, for a given  $\kappa$ , whether  $(\bar{\mathcal{S}}_\gamma^\kappa)^c \subset \mathcal{H}(T_0)$ . Then, by discretizing the range of  $\kappa$  or using a bisection algorithm, we can leverage this check to obtain (5).

We use an MIP to check  $(\bar{\mathcal{S}}_\gamma^\kappa)^c \subset \mathcal{H}(T_0)$ . Recall that  $\mathcal{H}(T_0) = \bigcup_{i:Y_i=0} \{x \in \mathbb{R}_+^d : x \leq \tilde{X}_i\}$ . We want to check if  $\{x \in \mathbb{R}_+^d : \hat{g}(x) \leq \kappa\}$  for a given  $\kappa$  lies completely inside the hull, where  $\hat{g}(x)$  is trained with a ReLU-activated neural net. This can be done by solving an optimization problem as follows. First, we rewrite  $\mathcal{H}(T_0)$  as  $\{x \in \mathbb{R}_+^d : \min_{i=1,\dots,n} \max_{j=1,\dots,d} \{x^j - \tilde{X}_i^j\} \leq 0\}$ , where  $x^j$  and  $\tilde{X}_i^j$  refer to the  $j$ -th components of  $x$  and  $\tilde{X}_i$  respectively. Then we solve

$$\begin{aligned} & \max_{x \in \mathbb{R}^d} && \min_{i=1,\dots,n} \max_{j=1,\dots,d} \{x^j - \tilde{X}_i^j\} \\ & \text{subject to} && \hat{g}(x) \leq \kappa \\ & && x \geq 0 \end{aligned} \quad (6)$$

If the optimal value is greater than 0, this means  $\{x \in \mathbb{R}_+^d : \hat{g}(x) \leq \kappa\}$  is not completely inside  $\mathcal{H}(T_0)$ , and vice versa. Now, we rewrite (6) as

$$\begin{aligned} & \max_{x \in \mathbb{R}^d, \beta \in \mathbb{R}} && \beta \\ & \text{subject to} && \max_{j=1,\dots,d} \{x^j - \tilde{X}_i^j\} \geq \beta \quad \forall i = 1, \dots, n \\ & && \hat{g}(x) \leq \kappa \\ & && x \geq 0 \end{aligned} \quad (7)$$

We then rewrite (7) as an MIP by introducing a large real number  $M$  as a practical upper bound for all coordinates of  $x$ :

$$\begin{aligned} & \max_{x \in \mathbb{R}^d, \beta \in \mathbb{R}} && \beta \\ & \text{subject to} && x^j - \tilde{X}_i^j + 4M(1 - z_{ij}) \geq \beta \quad \forall i = 1, \dots, n, j = 1, \dots, d \\ & && \sum_{j=1,\dots,d} z_{ij} \geq 1 \quad \forall i = 1, \dots, n \\ & && z_{ij} \in \{0, 1\} \quad \forall i = 1, \dots, n, j = 1, \dots, d \\ & && \hat{g}(x) \leq \kappa \\ & && x \geq 0 \end{aligned} \quad (8)$$

Note that the set of points  $T_0$  to be considered in constructing  $\mathcal{H}(T_0)$  can be reduced to its ‘‘extreme points’’. More concretely, we call a point  $x \in T_0$  an extreme point if there does not exist any other point  $x' \in T_0$  such that  $x \leq x'$ . We can eliminate all points  $x \in T_0$  such that  $x \leq x'$  for another  $x' \in T_0$ , and the resulting orthogonal monotone hull would remain the same. If we carry out this elimination, then in (7) we need only consider  $\tilde{X}_i$  that are extreme points in  $\mathcal{H}(T_0)$ , which can reduce the number of integer variables needed to add. In practice, we can also randomly remove points in  $T_0$  to further reduce the number of integer variables. This would not affect the correctness of our approach, but would increase the conservativeness of the final estimate.

## C Further Results for Section 3

Here we present and discuss several additional results for Section 3 regarding estimation efficiency and conservativeness. The latter includes further theorems on the lazy-learner classifier and classifiers constructed using the difference of two functions, translation of the false positive rate under the Stage 1 sampling distribution to under the original distribution, and interpretations and refinements of the conservativeness results.

### C.1 Extending Upper-Bound Relaxed Efficiency Certificate to Two-Stage Procedures

We present an extension of Proposition 1 to two-stage procedures, which is needed to set up Corollary 1

**Proposition 3** (Extended relaxed efficiency certificate). *Suppose constructing  $\hat{\mu}_n = \hat{\mu}_{n_2}(D_{n_1})$  consists of two stages, with  $n = n_1 + n_2$ : First we sample  $D_{n_1} = \{\tilde{X}_1, \dots, \tilde{X}_{n_1}\}$ , where  $\tilde{X}_i$  are i.i.d. (following some sampling distribution), and given  $D_{n_1}$ , we construct  $\hat{\mu}_{n_2}(D_{n_1}) = (1/n_2) \sum_{i=1}^{n_2} Z_i$  where  $Z_i$  are i.i.d. conditional on  $D_{n_1}$  (following some distribution). Suppose  $\hat{\mu}_n$  is conditionally upward biased almost surely, i.e.,  $\bar{\mu}(D_{n_1}) := E[\hat{\mu}_n | D_{n_1}] \geq \mu$ , and the conditional relative error given  $D_{n_1}$  in the second stage satisfies  $RE(D_{n_1}) := \text{Var}(Z_i | D_{n_1}) / \bar{\mu}(D_{n_1})^2 = \tilde{O}(\log(1/\bar{\mu}(D_{n_1})))$ . If  $n_1 = \tilde{O}(\log(1/\mu))$  (such as a constant number), then  $\hat{\mu}_n$  possesses the upper-bound relaxed efficiency certificate.*



### C.2 Conservativeness of Lazy Learner

We provide a result to quantify the conservativeness of the lazy-learner IS in terms of the false positive rate. Recall that the lazy learner constructs the outer approximation of the rare-event set using  $\mathcal{H}(T_0)^c$ , which is the complement of the orthogonal monotone hull of the set of all non-rare-event samples. The conservativeness is measured concretely by the set difference between  $\mathcal{H}(T_0)^c$  and  $\mathcal{S}_\gamma$ , for which we have the following result:

**Theorem 6** (Conservativeness of lazy learner). *Suppose that the density  $q$  has bounded support  $K \subset [0, M]^d$ , and  $0 < q_l \leq q(x) \leq q_u$  for any  $x \in K$ . Then, with probability at least  $1 - \delta$ ,*

$$\begin{aligned} P_{X \sim q}(X \in \mathcal{H}(T_0)^c \setminus \mathcal{S}_\gamma) &\leq M^{d-1} q_u \left( \frac{\sqrt{d}}{2} \right)^{d-1} w_{d-1} t(\delta, n_1) \\ &= \sqrt{\frac{e}{\pi(d-1)}} \left( \frac{1}{2} \pi e \right)^{\frac{d-1}{2}} q_u t(\delta, n_1) (1 + O(d^{-1})). \end{aligned}$$

Here  $t(\delta, n_1) = 3 \left( \frac{\log(n_1 q_l) + d \log M + \log \frac{1}{\delta}}{n_1 q_l} \right)^{\frac{1}{d}}$ ,  $w_d$  is the volume of a  $d$ -dimensional Euclidean ball of radius 1, and the last  $O(\cdot)$  is as  $d$  increases.

### C.3 Translating the False Positive Rate to under the Original distribution

Theorems 4 and 6 are stated with respect to  $q$ , the sampling distribution used in the first stage. We explain how to translate the false positive rate results to under the original distribution  $p$ . In the discussion below, we will consider Theorem 4 (and Theorem 6 can be handled similarly). In this case, our target is to give an upper bound to  $P_{X \sim p}(X \in \mathcal{S}_\gamma^c \setminus \mathcal{S}_\gamma)$  based on the result of Theorem 4.

If the true input distribution  $p$  does not have a bounded support, we can first choose  $M$  to be large to make sure that  $P_{X \sim p}(X \notin [0, M]^d)$  is small compared to the probability of  $\mathcal{S}_\gamma$ . We argue that we do not need  $M$  to be too large here. Indeed, if  $p$  is light tail (e.g., a distribution with tail probability exponential in  $M$ ), then the required  $M$  grows at most polynomially in  $\gamma$ .

Having selected  $M$ , and with the freedom in selecting  $q$  in Stage 1, we could make sure that in  $[0, M]^d$ ,  $q(x)$  is bounded away from 0 (e.g., we can choose  $q$  to be the uniform distribution over  $[0, M]^d$ ). Then, by Theorem 4 and a change of measure argument, we can give a bound for  $P_{X \sim p}(X \in [0, M]^d, X \in \mathcal{S}_\gamma^c \setminus \mathcal{S}_\gamma)$ . Finally, we bound the false positive rate with respect to  $p$  by  $P_{X \sim p}(X \in \mathcal{H}(T_0)^c \setminus \mathcal{S}_\gamma) \leq P_{X \sim p}(X \notin [0, M]^d) + P_{X \sim p}(X \in [0, M]^d, X \in \mathcal{H}(T_0)^c \setminus \mathcal{S}_\gamma)$ .

### C.4 Conservativeness Results for Classifiers Constructed Using Differences of Two Trained Functions

Theorem 4 presents a conservativeness result when  $\hat{g}$  is trained with an empirical risk minimization (ERM). In this subsection, we will show a more sophisticated version of Theorem 4 which corresponds more closely to the  $\hat{g}$  that we implemented in our experiments. Suppose that the Stage 1 samples are generated in the same way as in Algorithm 1. We let  $\mathcal{F} := \{f_\theta\}$  denote the function class induced by the model. Here a main difference with previously is that we allow functions in  $\mathcal{F}$  to be 2-dimensional, and both the loss function and the classification boundary will be constructed from these 2-dimensional functions.

Suppose that  $f_\theta$  is the output a neural network with 2 neurons in the output layer, and denote them as  $f_{\theta,0}, f_{\theta,1}$ . Let the loss function evaluated at the  $i$ -th sample be  $\ell(f_\theta(\tilde{X}_i), Y_i)$ . For example, the cross-entropy loss is given by  $- \left[ I(Y_i = 0) \log \frac{e^{f_{\theta,0}(\tilde{X}_i)}}{e^{f_{\theta,0}(\tilde{X}_i)} + e^{f_{\theta,1}(\tilde{X}_i)}} + I(Y_i = 1) \log \frac{e^{f_{\theta,1}(\tilde{X}_i)}}{e^{f_{\theta,0}(\tilde{X}_i)} + e^{f_{\theta,1}(\tilde{X}_i)}} \right]$ . Like in the ERM approach in Theorem 4, we compute  $\hat{f} = f_{\hat{\theta}} \in \mathcal{F}$  which is the minimizer of the empirical risk, i.e.,  $\hat{f} = \arg \min_{f_\theta \in \mathcal{F}} R_{n_1}(f_\theta)$ . For each function  $f_\theta \in \mathcal{F}$ , define function  $g_\theta$  as  $g_\theta := f_{\theta,1} - f_{\theta,0}$ . In this modified approach, the learned rare-event set would be given by  $\tilde{\mathcal{S}}_\gamma^\kappa := \{x : g_{\hat{\theta}}(x) \geq \kappa\}$ , and to make sure that  $\mathcal{S}_\gamma \subset \tilde{\mathcal{S}}_\gamma^\kappa$ , we would replace  $\kappa$  by  $\hat{\kappa} = \min\{g_{\hat{\theta}}(x) : x \notin \mathcal{H}(T_0)\}$  as in Step 1 of Algorithm 1.

We give a theorem similar to Theorem 4 for this more sophisticated procedure. To this end, we begin by giving some definitions similar to the set up of Theorem 4. Let  $R(f_\theta) := E_{X \sim q} \ell(f_\theta(X), I(X \in \mathcal{S}_\gamma))$  denote the true risk

function. Let  $f^* = \arg \min_{f \in \mathcal{F}} R(f)$  denote the true risk minimizer within function class  $\mathcal{F}$ . Define  $g^* = f_1^* - f_0^*$  accordingly and let  $\kappa^* := \min_{x \notin \mathcal{S}_\gamma} g^*(x)$  denote the true threshold associated with  $f^*$  in obtaining the smallest outer rare-event approximation.

**Theorem 7.** *Suppose that the density  $q$  has bounded support  $K \subset [0, M]^d$  and  $0 < q_l \leq q(x) \leq q_u$  for any  $x \in K$ . Also suppose that there exists a function  $h$  such that for any  $f_\theta \in \mathcal{F}$ , if  $g_\theta(x) \geq \kappa$ , we have  $\ell(f_\theta(x), 0) \geq h(\kappa) > 0$  (for the cross entropy loss, this happens if we know that  $f_\theta$  has a bounded range). Then, for the set  $\tilde{\mathcal{S}}_\gamma^{\hat{\kappa}}$ , with probability at least  $1 - \delta$ ,*

$$\begin{aligned} & P_{X \sim q} \left( X \in \tilde{\mathcal{S}}_\gamma^{\hat{\kappa}}, X \in \mathcal{S}_\gamma^c \right) \\ & \leq \left( h(\kappa^*) - t(\delta, n_1) \sqrt{d} \text{Lip}(g^*) - \|\hat{g} - g^*\|_\infty \right)^{-1} \left( R(f^*) + 2 \sup_{f_\theta \in \mathcal{F}} |R_{n_1}(f_\theta) - R(f_\theta)| \right) \end{aligned}$$

Here  $\text{Lip}(g^*)$  is the Lipschitz parameter of  $g^*$ , and  $t(\delta, n_1)$  is defined as in Theorem 4.

### C.5 Implications of Theorem 4 and Related Results in the Literature

First, we explain the trade-off between overfitting and underfitting. If the function class  $\mathcal{G}$  is not rich, then  $R(g^*) = \inf_{g \in \mathcal{G}} R(g)$  may be big because of the lack of expressive power. On the other hand, if the function class is too rich, then the generalization error will be huge. Here, the generalization error is represented by  $\sup_{g_\theta \in \mathcal{G}} |R_{n_1}(g_\theta) - R(g_\theta)|$  as well as  $t(\delta, n_1) \sqrt{d} \text{lip}(g^*) + \|\hat{g} - g^*\|_\infty$ , which characterize the difference between the right hand side of the bound in the theorem and its limit as  $n_1 \rightarrow \infty$ .

Another question is how to give a more refined bound for the false positive rate based on Theorem 4 that depends on explicit constants of the classification model or training process. This would involve theoretical results for deep neural networks that are under active research. Let us examine the terms appearing in Theorem 4 and give some related results. In machine learning theory, the term  $\sup_{g_\theta \in \mathcal{G}} |R_{n_1}(g_\theta) - R(g_\theta)|$  is often bounded by the Rademacher complexity of the function class (some results about the Rademacher complexity for neural networks are in Harvey et al. 2017; Cao and Gu 2019). The convergence of  $\|\hat{g} - g^*\|_\infty$  to 0 as  $n_1 \rightarrow \infty$  is implied by the convergence of the parameters, which is in turn justified by the empirical process theory (van der Vaart and Wellner, 1996). A bound for  $\text{Lip}(g^*)$  could be potentially derived by adding norm constraints to the parameters in the neural network (Anil et al., 2019). On the other hand, if we let the network size grow to infinity, the class of neural networks can approximate any continuous function (Lu et al., 2017), and hence  $R(g^*)$  can be arbitrarily small when the neural network is complex enough. However, if we restrict the choices of networks, for instance by the Lipschitz constant, then no results regarding the sufficiency of its expressive power for arbitrary functions are available in the literature to our knowledge, and thus it appears open how to simultaneously give bounds for  $\text{Lip}(g^*)$  and  $R(g^*)$ . Future investigations on the expressive power of restricted classes of neural networks would help refining our conservativeness results further.

## D Lower-Bound Efficiency Certificate and Estimators

In Section 3, we described an approach that gives an estimator for the rare-event probability with an upper-bound relaxed efficiency certificate. Here we present analogous definitions and results on the lower-bound relaxed efficiency certificate. This lower-bound estimator gives an estimation gap for the upper-bound estimator. Moreover, by combining both of them, we can obtain an interval for the target rare-event probability.

The lower-bound relaxed efficiency certificate is defined as follows (compare with Definition 3):

**Definition 6.** *We say an estimator  $\hat{\mu}_n$  satisfies a lower-bound relaxed efficiency certificate to estimate  $\mu$  if  $P(\hat{\mu}_n - \mu > \epsilon\mu) \leq \delta$  with  $n \geq \tilde{O}(\log(1/\mu))$ , for given  $0 < \epsilon, \delta < 1$ .*

This definition requires that, with high probability,  $\hat{\mu}_n$  is a lower bound of  $\mu$  up to an error of  $\epsilon\mu$ . We have the following analog to Proposition 1:

**Corollary 2.** *Suppose  $\hat{\mu}_n$  is downward biased, i.e.,  $\bar{\mu} := E[\hat{\mu}_n] \leq \mu$ . Moreover, suppose  $\hat{\mu}_n$  takes the form of an average of  $n$  i.i.d. simulation runs  $Z_i$ , with  $RE = \text{Var}(Z_i)/\bar{\mu}^2 = \tilde{O}(\log(1/\bar{\mu}))$ . Then  $\hat{\mu}_n$  possesses the lower-bound relaxed efficiency certificate.*

This motivates us to learn an inner approximation of the rare-event set in Stage 1 and then in Stage 2, we use IS as in Theorem 1 to estimate the probability of this inner approximation set. For the inner set approximation, like the outer approximation case, we use our Stage 1 samples  $\{(\tilde{X}_i, Y_i)\}_{i=1, \dots, n_1}$  to construct an approximation set  $\bar{\mathcal{S}}_\gamma$  that has zero false positive rate, i.e.,

$$P(X \in \bar{\mathcal{S}}_\gamma, Y = 0) = 0. \quad (9)$$

To make sure of (9), we again exploit the knowledge that the rare event set  $\mathcal{S}_\gamma$  is orthogonally monotone. Indeed, denote  $T_1 := \{\tilde{X}_i : Y_i = 1\}$  as the rare-event sampled points and for each point  $x \in \mathbb{R}_+^d$ , let  $\mathcal{Q}(x) := \{x' : x' \geq x\}$ . We construct  $\mathcal{J}(T_1) := \cup_{x \in T_1} \mathcal{Q}(x)$  which serves as the ‘‘upper orthogonal monotone hull’’ of  $T_1$ . The orthogonal monotonicity property of  $\mathcal{S}_\gamma$  implies that  $\mathcal{J}(T_1) \subset \mathcal{S}_\gamma$ . Moreover,  $\mathcal{J}(T_1)$  is the largest choice of  $\bar{\mathcal{S}}_\gamma$  such that (9) is guaranteed. Based on this observation, in parallel to Section 3, depending on how we construct the inner approximation to the rare-event set, we propose the following two approaches.

**Lazy-Learner IS (Lower Bound).** We now consider an estimator for  $\mu$  where in Stage 1, we sample a constant  $n_1$  i.i.d. random points from some density, say  $q$ . Then, we use the mixture IS depicted in Theorem 1 to estimate  $P(X \in \mathcal{J}(T_1))$  in Stage 2. Since  $\mathcal{J}(T_1)$  takes the form  $\cup_{x \in T_1} \mathcal{Q}(x)$ , it has a finite number of dominating points, which can be found by a sequential algorithm. But as explained in Section 3, this leads to a large number of mixture components that degrades the IS efficiency.

**Deep-Learning-Based IS (Lower Bound).** We train a neural network classifier, say  $\hat{g}$ , using all the Stage 1 samples  $\{(\tilde{X}_i, Y_i)\}$ , and obtain an approximate rare-event region  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}} = \{x : \hat{g}(x) \geq \hat{\kappa}\}$ , where  $\hat{\kappa}$  is say 1/2. Then we adjust  $\hat{\kappa}$  minimally away from 1/2, say to  $\hat{\kappa}$ , so that  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}} \subset \mathcal{J}(T_1)$ , i.e.,  $\hat{\kappa} = \min\{\kappa \in \mathbb{R} : \bar{\mathcal{S}}_\gamma^{\hat{\kappa}} \subset \mathcal{J}(T_1)\}$ . Then  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  is an inner approximation for  $\mathcal{S}_\gamma$  (see Figure 1(c), where  $\hat{\kappa} = 0.83$ ). Stage 1 in Algorithm 2 shows this procedure. With this, we can run mixture IS to estimate  $P(X \in \bar{\mathcal{S}}_\gamma^{\hat{\kappa}})$  in Stage 2.

---

**Algorithm 2: Deep-PrAE to estimate  $\mu = P(X \in \mathcal{S}_\gamma)$  (lower bound).**

---

**Input:** Black-box evaluator  $I(\cdot \in \mathcal{S}_\gamma)$ , initial Stage 1 samples  $\{(\tilde{X}_i, Y_i)\}_{i=1, \dots, n_1}$ , Stage 2 sampling budget  $n_2$ , input distribution  $N(\lambda, \Sigma)$ .

**Output:** IS estimate  $\hat{\mu}_n$ .

1 **Stage 1 (Set Learning):**

2 Train classifier with positive decision region  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}} = \{x : \hat{g}(x) \geq \hat{\kappa}\}$  using  $\{(\tilde{X}_i, Y_i)\}_{i=1, \dots, n_1}$ ;

3 Replace  $\hat{\kappa}$  by  $\hat{\kappa} = \min\{\kappa \in \mathbb{R} : \bar{\mathcal{S}}_\gamma^{\hat{\kappa}} \subset \mathcal{J}(T_1)\}$ ;

4 **Stage 2 (Mixture IS based on Searched dominating points):**

5 The same as Stage 2 of Algorithm 1.

---

As we can see, compared with Algorithm 1, the only difference is how we adjust  $\kappa$  in Stage 1. And similar to Theorem 3, we also have that Algorithm 2 attains the lower-bound relaxed efficiency certificate:

**Theorem 8** (Lower-bound relaxed efficiency certificate for deep-learning-based mixture IS). *Suppose  $\mathcal{S}_\gamma$  is orthogonally monotone, and  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  satisfies the same conditions for  $\mathcal{S}_\gamma$  in Theorem 1. Then Algorithm 2 attains the lower-bound relaxed efficiency certificate by using a constant number of Stage 1 samples.*

Finally, we investigate the conservativeness of this bound, which is measured by the false negative rate  $P(X \notin \bar{\mathcal{S}}_\gamma^{\hat{\kappa}}, Y = 1)$ . Like in Section 3, we use ERM to train  $\hat{g}$ , i.e.,  $\hat{g} := \operatorname{argmin}_{g \in \mathcal{G}} \{R_{n_1}(g) := \frac{1}{n_1} \sum_{i=1}^{n_1} \ell(g(\tilde{X}_i), Y_i)\}$  where  $\ell$  is a loss function and  $\mathcal{G}$  is the considered hypothesis class. Let  $g^*$  be the true risk minimizer as described in Section 3. For inner approximation, we let  $\kappa^* := \max_{x \in \mathcal{S}_\gamma} g^*(x)$  be the true threshold associated with  $g^*$  in obtaining the largest inner rare-event set approximation. Then we have the following result analogous to Theorem 4.

**Theorem 9** (Lower-bound estimation conservativeness). *Consider  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  obtained in Algorithm 2 where  $\hat{g}$  is trained from an ERM. Suppose the density  $q$  has bounded support  $K \subset [0, M]^d$  and  $0 < q_l \leq q(x) \leq q_u$  for any  $x \in K$ . Also suppose there exists a function  $h$  such that for any  $g \in \mathcal{G}$ ,  $g(x) \leq \kappa$  implies  $\ell(g(x), 1) \geq h(\kappa) > 0$ . (e.g., if  $\ell$  is the squared loss, then  $h(\kappa)$  could be chosen as  $h(\kappa) = (1 - \kappa)^2$ ). Then, with probability at least  $1 - \delta$ ,*

$$P_{X \sim q} \left( X \in \bar{\mathcal{S}}_\gamma^{\hat{\kappa}} \setminus \mathcal{S}_\gamma \right) \leq \frac{R(g^*) + 2 \sup_{g \in \mathcal{G}} |R_{n_1}(g) - R(g)|}{h(\kappa^* + t(\delta, n_1)) \sqrt{d} \operatorname{Lip}(g^*) + \|\hat{g} - g^*\|_\infty}.$$

Here,  $Lip(g^*)$  is the Lipschitz parameter of  $g^*$ , and  $t(\delta, n_1) = 3 \left( \frac{\log(n_1 q_1) + d \log M + \log \frac{1}{\delta}}{n_1 q_1} \right)^{\frac{1}{d}}$ .

## E Cross Entropy and Adaptive Multilevel Splitting

We provide some details on the cross-entropy method and adaptive multilevel splitting (or subset simulation), and also discuss their challenges in black-box problems.

**Cross Entropy.** The cross-entropy (CE) method (De Boer et al., 2005; Rubinstein and Kroese, 2013) uses a sequential optimization approach to iteratively solve for the optimal parameter in a parametric class of IS distributions. The objective in this optimization sequence is to minimize the Kullback–Leibler divergence between the IS distribution and the zero-variance IS distribution (the latter is theoretically known to be the conditional distribution given the occurrence of the rare event, but is unimplementable as it requires knowing the rare-event probability itself). Specifically, assume we are interested in estimating  $P(g(X) > \gamma)$  and a parametric class  $p_\theta$  is considered. The cross-entropy method adaptively chooses  $\gamma_1 < \gamma_2 < \dots < \gamma$ . At each intermediate level  $k$ , we use the updated IS distribution  $p_{\theta_k^*}$ , designed for simulating  $P(g(X) > \gamma_k)$ , as the sampling distribution to draw samples of  $X$  that sets up an empirical optimization, from which the next  $\theta_{k+1}^*$  is obtained.

While flexible and easy to use, the efficiency of CE depends crucially on the expressiveness of the parametric class  $p_\theta$  and the parameter convergence induced by the empirical optimization sequence. There are good approaches to determine the parametric classes (e.g., Botev et al., 2016), and also studies on the efficiency of IS distributions parametrized by empirical optimization (Tuffin and Ridder, 2012). However, it is challenging to obtain an efficiency certificate for CE that requires iterative empirical optimization in the common form depicted above. Insufficiency on either the choice of the parametric class or the parameter convergence may lead to the undetectable under-estimation issue (e.g., as in Theorem 2).

**Adaptive Multilevel Splitting.** Adaptive multilevel splitting (AMS) (or subset simulation) (C erou and Guyader, 2007; Au and Beck, 2001) decomposes the rare-event estimation problem into estimating a sequence of conditional probabilities. We adaptively choose a threshold sequence  $\gamma_1 < \gamma_2 < \dots < \gamma_K = \gamma$ . Then  $P(g(x) > \gamma)$  can be rewritten as  $P(g(x) > \gamma) = P(g(x) > \gamma_1) \prod_{k=2}^K P(g(x) > \gamma_k | g(x) > \gamma_{k-1})$ . AMS then aims to estimate  $P(g(x) > \gamma_1)$  and  $P(g(x) > \gamma_k | g(x) > \gamma_{k-1})$  for each intermediate level  $k = 2, \dots, K$ . In standard implementation, these conditional probabilities are estimated using samples from  $p(g(x) > \gamma_k | g(x) > \gamma_{k-1})$  through variants of the Metropolis-Hasting (MH) algorithms.

Theoretical studies have shown some nice properties of AMS, including unbiasedness and asymptotic normality (e.g., see C erou et al., 2019). However, the variance of the estimator depends on the mixing property of the proposal distribution in the MH steps (C erou and Guyader, 2016). Under ideal settings when direct sampling from  $P(g(x) > \gamma_k | g(x) > \gamma_{k-1})$  is possible, it is shown that AMS is ‘‘almost’’ asymptotically optimal (Guyader et al., 2011). However, to our best knowledge, there is yet any study on provable efficiency of rare-event estimators with consideration of both AMS and MH sampling. In practice, to achieve a good performance, AMS requires a proposal distribution in the MH algorithm that can efficiently generate samples with low correlations.

## F Further Details for Numerical Experiments

This section provides more details on the two experimental examples in Section 4.

### F.1 2D Example

In the 2D example, the rarity parameter  $\gamma$  governs the shape of the rare-event set  $\mathcal{S}_\gamma = \{x : g(x) \geq \gamma\}$ . We consider a linear combination of sigmoid functions  $g(x) = \|\theta_1 \psi(x - c_1 - \gamma) + \theta_2 \psi(x - c_2 - \gamma) + \theta_3 \psi(x - c_3 - \gamma) + \theta_4 \psi(x - c_4 - \gamma)\|$  where  $\theta, c$  are some constant vectors and  $\psi(x) = \frac{\exp(x)}{1 + \exp(x)}$ . A point  $x$  is a rare-event if  $g(x) > \gamma$ , where we take  $\gamma = 1.8$  in Section 4. We use  $p = N([5, 5]^T, 0.25I_{2 \times 2})$ . Figure 5 shows the rare-event set and its approximations for various  $\gamma$ 's. The Deep-PrAE boundaries seem tight in most cases, attributed to both the sufficiently trained NN classifier and the bisection algorithm implemented for tuning  $\hat{\kappa}$  after the NN training.




Figure 5: The contour of  $p$ , rare-event set  $\mathcal{S}_\gamma$  (reddish region), outer- and inner- approximation boundaries (black lines) and Deep-PrAE UB and LB decision boundaries (white lines) for some  $\gamma$  values in the 2D example.

## F.2 Intelligent Driving Safety Testing Example

We provide more details about the self-driving example, which simulates the interaction of an autonomous vehicle (AV) model that follows a human-driven lead vehicle (LV). The AV is controlled by the Intelligent Driver Model (IDM), widely used for autonomy evaluation and microscopic transportation simulation, that maintains a safety distance while ensuring smooth ride and maximum efficiency. The states of the AV are  $s_t = [x_{\text{follow}}, x_{\text{lead}}, v_{\text{follow}}, v_{\text{lead}}, a_{\text{follow}}, a_{\text{lead}}]_t$  which are the position, velocity and acceleration of the AV and LV respectively. The throttle input to the AV is defined as  $u_t$  which has an affine relationship with the acceleration of the vehicle. Similarly, the randomized throttle of the LV is represented by  $w_t$ . With a car length of  $L$ , the distance between the LV and AV at time  $t$  is given by  $r_t = x_{\text{lead},t} - x_{\text{follow},t} - L$ , which has to remain below the crash threshold for safety. We describe the dynamics in more detail below. Figure 6 gives a pictorial overview of the interaction.




Figure 6: The states  $s_t$  and input  $u_t$  of the self-driving safety-testing simulation.  $w_t$  denotes the throttle input of the AV from the IDM.

Table 1: Parameters of the Intelligent Drivers Model (IDM)

Parameters	Value
Safety distance, $s_0$	2 m
Speed of AV in free traffic, $v_0$	30 m/s
Maximum acceleration of AV, $a$	$2\gamma$ m/s <sup>2</sup>
Comfortable deceleration of AV, $b$	1.67 m/s <sup>2</sup>
Maximum deceleration of AV, $d$	$2\gamma$ m/s <sup>2</sup>
Safe time headway, $\bar{T}$	1.5 s
Acceleration exponent parameter, $\delta$	4
Car length, $L$	4 m

**LV actions.** The LV action contains human-driving uncertainty in decision-making modeled as Gaussian increments. For every  $\Delta t$  time-steps, a Gaussian random variable is generated with the mean centered at the previous action  $u_{t-\Delta t}$ . We initialize  $u_0 = 10$  (unitless) and  $\Delta t = 4$  sec, which corresponds to zero initial acceleration and an acceleration change in the LV once every 4 seconds.

**Intelligent Driver Model (IDM) for AV.** The IDM is governed by the following equations (the subscripts “follow” and “lead” defined in Figure 6 is abbreviated to “f” and “l” for conciseness):

$$\begin{aligned}
 \dot{x}_f &= v_f \\
 \dot{x}_l &= v_l \\
 \dot{v}_f &= \max \left( a \left( 1 - \left( \frac{v_f}{v_0} \right)^\delta - \left( \frac{s^*(v_f, \Delta v_f)}{s_f} \right)^2 \right), -d \right) \\
 s^*(v_f, \Delta v_f) &= s_0 + v_f \bar{T} + \frac{v_f \Delta v_f}{2\sqrt{ab}} \\
 s_f &= x_l - x_f - L \\
 \Delta v_f &= v_f - v_l,
 \end{aligned}$$

The parameters are presented in Table 1 and  $v_l \propto u_t$  and  $v_f \propto w_t$ . The randomness of LV actions  $u_t$ ’s propagates into the system and affects all the simulation states  $s_t$ . The IDM is governed by simple first-order kinematic equations for the position and velocity of the vehicles. The acceleration of the AV is the decision variable where it is defined by a sum of non-linear terms which dictate the “free-road” and “interaction” behaviors of the AV and LV. The acceleration of the AV is constructed in such a way that certain terms of the equations dominate when the LV is far away from the AV to influence its actions and other terms dominate when the LV is in close proximity to the AV.

**Rarity parameter  $\gamma$ .** Parameter  $\gamma$  signifies the range invoked by the AV acceleration and deceleration pedals. Increasing  $\gamma$  implies that the AV can have sudden high deceleration and hence avoid crash scenarios better and making crashes rarer. In contrast, decreasing  $\gamma$  reduces the braking capability of the AV and more easily leads to crashes. For instance,  $\gamma = 1.0$  corresponds to AV actions in the range  $[5, 15]$  or correspondingly  $a_{\text{follow}, t} \in [-2, 2]$ , and  $\gamma = 2.0$  corresponds to  $a_{\text{follow}, t} \in [-4, 4]$ . Figure 7 shows the approximate rare-event set by randomly sampling points and evaluating the inclusion in the set, for the two cases of  $\gamma = 1.0$  and  $\gamma = 2.0$ . In particular, we slice the 15-dimensional space onto pairs from five of the dimensions. In all plots, we see that the crash set (red) are monotone, thus supporting the use of our Deep-PrAE framework. Although the crash set is not located in the “upper-right corner”, we can implement Deep-PrAE framework for such problems by simple re-orientation.

**Sample trajectories.** Figure 8 shows two examples of sample trajectories, one successfully maintaining a safe distance, and the other leading to a crash. In Figure 8(e)-(h) where we show the crash case, the AV maintains a safe distance behind the LV until the latter starts rapidly decelerating (Figure 8(h)). Here the action corresponds to the throttle input that has an affine relationship with the acceleration of the vehicle. The LV ultimately decelerates at a rate that the AV cannot attain and its deceleration saturates after a point which leads to the crash.

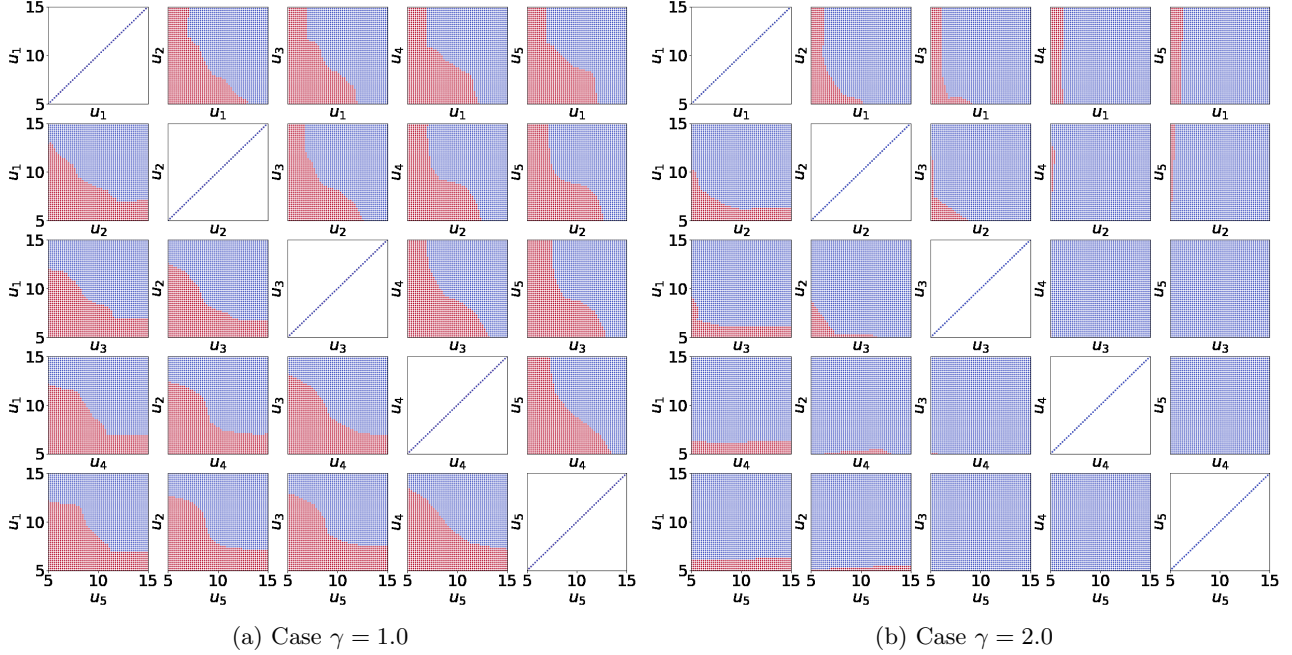


Figure 7: Slice of pairs of the first 5 dimensions of LV action space. For any  $(u_i, u_{i'})$  shown,  $u_j, j \notin \{i, i'\}$  is fixed at a constant value. Blue dots = non-crash cases, red dots = crash cases.

### F.3 Code

The code and environment settings for the experiments are available at <https://github.com/safeai-lab/DeepPrAE/>.

## G Proofs

### G.1 Proofs for the Dominating Point Methodologies

*Proof of Proposition 2.* Since  $\mu$  is exponentially decaying in  $\gamma$  while  $n$  is polynomially growing in  $\gamma$ , we know that  $\lim_{\gamma \rightarrow \infty} n\mu = 0$ . Since  $n\hat{\mu}_n$  takes values in  $\{0, 1, \dots, n\}$ , we get that  $P(|\hat{\mu}_n - \mu| > \varepsilon\mu) = P(|n\hat{\mu}_n - n\mu| > \varepsilon n\mu) \rightarrow 1$  as  $\gamma \rightarrow \infty$ .  $\square$

*Proof of Theorem 5.* Throughout this proof, we write  $f(\gamma) \sim g(\gamma)$  if  $f(\gamma)/g(\gamma)$  changes at most polynomially in  $\gamma$ . We know that

$$\tilde{E}[Z^2] = \sum_j \tilde{E}[I(X \in \mathcal{S}_\gamma^j) L^2(X)] \leq \sum_j e^{-(a_j - \lambda)^T \Sigma^{-1} (a_j - \lambda)} / \alpha_j \sim e^{-(a^* - \lambda)^T \Sigma^{-1} (a^* - \lambda)}.$$

Denote  $Y = B(X - \lambda) \sim N(0, B\Sigma B^T)$  and  $s = B(a^* - \lambda)$ . Define  $\tilde{\varepsilon} = \varepsilon \min_{u: u^T (B\Sigma B^T)^{-1} u = 1} \|u\|_\infty$ . Then we also

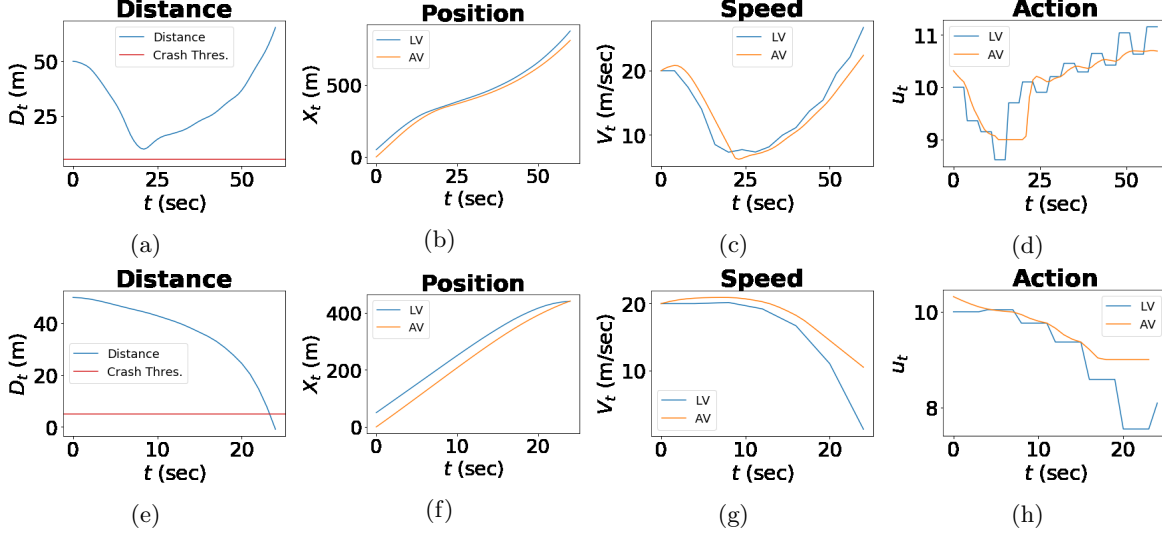


Figure 8: Autonomous Car Following Experiment Trajectories. Figures (a) - (d) represent a simulation episode without a crash occurring where the AV follows the LV successfully at a safe distance. Figures (e) - (h) represents a simulation episode where crash occurs at  $t = 23$  seconds due to the repeated deceleration of the LV.

know that

$$\begin{aligned}
 & \tilde{E}[I(X \in \mathcal{S}_\gamma)L(X)] \\
 & \geq P(B(X - a^*) \geq 0, (X - a^*)^T \Sigma^{-1}(X - a^*) \leq \varepsilon^2) \\
 & = P(Y \geq s, (Y - s)^T (B \Sigma B^T)^{-1}(Y - s) \leq \varepsilon^2) \\
 & = \int_{y \geq s, (y-s)^T (B \Sigma B^T)^{-1}(y-s) \leq \varepsilon^2} (2\pi)^{-d/2} |B \Sigma B^T|^{-1/2} e^{-y^T (B \Sigma B^T)^{-1} y / 2} dy \\
 & \geq (2\pi)^{-d/2} |B \Sigma B^T|^{-1/2} e^{-\varepsilon^2 / 2} e^{-(a^* - \lambda)^T \Sigma^{-1}(a^* - \lambda) / 2} \\
 & \quad \int_{y \geq s, (y-s)^T (B \Sigma B^T)^{-1}(y-s) \leq \varepsilon^2} e^{-s^T (B \Sigma B^T)^{-1}(y-s)} dy \\
 & \geq (2\pi)^{-d/2} |B \Sigma B^T|^{-1/2} e^{-\varepsilon^2 / 2} e^{-(a^* - \lambda)^T \Sigma^{-1}(a^* - \lambda) / 2} \prod_{i=1}^d \int_0^{\tilde{\varepsilon}} e^{-s^T (B \Sigma B^T)^{-1} e_i u_i} du_i \\
 & = (2\pi)^{-d/2} |B \Sigma B^T|^{-1/2} e^{-\varepsilon^2 / 2} e^{-(a^* - \lambda)^T \Sigma^{-1}(a^* - \lambda) / 2} \prod_{i=1}^d \frac{1 - e^{-s^T (B \Sigma B^T)^{-1} e_i \tilde{\varepsilon}}}{s^T (B \Sigma B^T)^{-1} e_i}.
 \end{aligned}$$

Note that it is easy to verify that  $s^T (B \Sigma B^T)^{-1} e_i \geq 0$ . If  $s^T (B \Sigma B^T)^{-1} e_i = 0$ , then we naturally use  $\tilde{\varepsilon}$  to substitute  $\frac{1 - e^{-s^T (B \Sigma B^T)^{-1} e_i \tilde{\varepsilon}}}{s^T (B \Sigma B^T)^{-1} e_i}$ . Since we have assumed that the components of  $a^*$  are at most polynomially growing in  $\gamma$ , finally we get that

$$\tilde{E}[I(X \in \mathcal{S}_\gamma)L(X)] \sim e^{-(a^* - \lambda)^T \Sigma^{-1}(a^* - \lambda) / 2}$$

and hence  $\tilde{E}[Z^2] / \tilde{E}[Z]^2$  is at most polynomially growing in  $\gamma$ . □

*Proof of Theorem 2.* We know that  $\tilde{E}[Z] = \bar{\Phi}(\gamma) + \bar{\Phi}(k\gamma)$ . Moreover,

$$\tilde{E}[Z^2] = e^{\gamma^2} (\bar{\Phi}(2\gamma) + \bar{\Phi}((k-1)\gamma)).$$

If  $0 < k \leq 1$ , then  $\tilde{E}[Z] = O(e^{-k^2 \gamma^2 / 2} / \gamma)$  and  $\tilde{E}[Z^2] = O(e^{\gamma^2})$  as  $\gamma \rightarrow \infty$ . If  $1 < k < 3$ , then  $\tilde{E}[Z] = O(e^{-\gamma^2 / 2} / \gamma)$  and  $\tilde{E}[Z^2] = O(e^{(1-(k-1)^2/2)\gamma^2} / \gamma)$  as  $\gamma \rightarrow \infty$ . In both cases, we get that  $\tilde{E}[Z^2] / \tilde{E}[Z]^2$  grows



exponentially in  $\gamma$ . On the other hand, we know that

$$\begin{aligned} & \tilde{P} \left( \left| \frac{1}{n} \sum_i Z_i - \bar{\Phi}(\gamma) \right| > \varepsilon \bar{\Phi}(\gamma) \right) \\ & \leq \tilde{P}(\exists i : X_i \leq -k\gamma) + \tilde{P} \left( \left| \frac{1}{n} \sum_i I(X_i \geq \gamma) e^{\gamma^2/2 - \gamma X_i} - \bar{\Phi}(\gamma) \right| > \varepsilon \bar{\Phi}(\gamma) \right). \end{aligned}$$

Clearly  $\tilde{P}(\exists i : X_i \leq -k\gamma) = 1 - (1 - \bar{\Phi}((k+1)\gamma))^n = O(n\bar{\Phi}((k+1)\gamma))$ , which is exponentially decreasing in  $\gamma$  as  $n$  is polynomial in  $\gamma$ . Moreover, by Chebyshev's inequality,

$$\begin{aligned} & \tilde{P} \left( \left| \frac{1}{n} \sum_i I(X_i \geq \gamma) e^{\gamma^2/2 - \gamma X_i} - \bar{\Phi}(\gamma) \right| > \varepsilon \bar{\Phi}(\gamma) \right) \\ & \leq \frac{\tilde{E}[I(X_i \geq \gamma) e^{\gamma^2 - 2\gamma X_i}]}{n\varepsilon^2 \bar{\Phi}^2(\gamma)} = \frac{e^{\gamma^2} \bar{\Phi}(2\gamma)}{n\varepsilon^2 \bar{\Phi}^2(\gamma)} = O\left(\frac{\gamma}{n\varepsilon^2}\right). \end{aligned}$$

Thus  $P(|\hat{\mu}_n - \bar{\Phi}(\gamma)| > \varepsilon \bar{\Phi}(\gamma)) = O\left(\frac{\gamma}{n\varepsilon^2}\right)$ . Moreover, we know that  $P(\exists i : Z_i > 0) \geq 1 - 1/2^n$  and if  $Z_i > 0$  for some  $i$ , then we have that

$$\frac{\sum_i Z_i^2/n}{(\sum_i Z_i/n)^2} \leq n^2. \quad \square$$

## G.2 Proofs for the Relaxed Efficiency Certificate

*Proof of Proposition 1.* We have

$$P(\hat{\mu}_n - \mu < -\varepsilon\mu) \leq P(\hat{\mu}_n - \bar{\mu} < -\varepsilon\bar{\mu})$$

since  $\bar{\mu} \geq \mu$  and  $1 - \varepsilon > 0$ . Note that the Markov inequality gives

$$P(\hat{\mu}_n - \bar{\mu} < -\varepsilon\bar{\mu}) \leq \frac{\widetilde{Var}(Z_i)}{n\varepsilon^2\bar{\mu}^2}$$

so that

$$n \geq \frac{\widetilde{Var}(Z_i)}{\delta\varepsilon^2\bar{\mu}^2} = \frac{RE}{\delta\varepsilon^2} = \tilde{O}\left(\log \frac{1}{\mu}\right) = \tilde{O}\left(\log \frac{1}{\mu}\right)$$

achieves the relaxed efficiency certificate. □

*Proof of Proposition 3.* The proof follows from that of Proposition 1 with a conditioning on  $D_{n_1}$ . We have

$$P(\hat{\mu}_n - \mu < -\varepsilon\mu | D_{n_1}) \leq P(\hat{\mu}_n - \bar{\mu}(D_{n_1}) < -\varepsilon\bar{\mu}(D_{n_1}) | D_{n_1})$$

since  $\bar{\mu}(D_{n_1}) \geq \mu$  almost surely and  $1 - \varepsilon > 0$ . Note that the Markov inequality gives

$$P(\hat{\mu}_n - \bar{\mu}(D_{n_1}) < -\varepsilon\bar{\mu}(D_{n_1}) | D_{n_1}) \leq \frac{Var(Z_i | D_{n_1})}{n_2\varepsilon^2\bar{\mu}(D_{n_1})^2}$$

so that

$$n_2 \geq \frac{Var(Z_i | D_{n_1})}{\delta\varepsilon^2\bar{\mu}(D_{n_1})^2} = \frac{RE(D_{n_1})}{\delta\varepsilon^2} = \tilde{O}\left(\log \left(\frac{1}{\bar{\mu}(D_{n_1})}\right)\right) = \tilde{O}\left(\log \frac{1}{\mu}\right)$$

almost surely. Thus,

$$n = n_1 + n_2 \geq \tilde{O}\left(\log \frac{1}{\mu}\right)$$

achieves the relaxed efficiency certificate. □

*Proof of Corollary 1.* Follows directly from Proposition 3, since  $\bar{\mathcal{S}}_\gamma \supset \mathcal{S}_\gamma$  implies  $\bar{\mu}(D_{n_1}) \geq \mu$  almost surely.  $\square$

*Proof of Theorem 3.* We have assumed that  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$  satisfies the assumptions for  $\mathcal{S}_\gamma$  in Theorem 5. Then following the proof of Theorem 5, we obtain the efficiency certificate for the IS estimator in estimating its mean. Theorem 3 is then proved by directly applying Corollary 1.  $\square$

### G.3 Proofs for Conservativeness

Recall that  $T_0 = \{\tilde{X}_i : Y_i = 0\}$  where the samples are generated as in Algorithm 1. By some combinatorial argument, we can prove the following lemma which says that with high probability, each point in  $\mathcal{S}_\gamma^c$  that has sufficient distance to its boundary could be covered by  $\mathcal{H}(T_0)$ .

**Lemma 1.** *Suppose that the density  $q$  has bounded support  $K \subset [0, M]^d$ , and for any  $x \in K$ , suppose that  $0 < q_l \leq q(x) \leq q_u$ . Define  $B_t := \{x \in \mathcal{S}_\gamma^c : x + t\mathbf{1}_{d \times 1} \in \mathcal{S}_\gamma^c\}$ . Then with probability at least  $1 - \delta$ , we have that  $B_{t(\delta, n_1)} \subset \mathcal{H}(T_0)$ . Here  $t(\delta, n_1) = 3 \left( \frac{\log(n_1 q_l) + d \log M + \log \frac{1}{\delta}}{n_1 q_l} \right)^{\frac{1}{d}}$ .*

*Proof.* The basic idea is to construct a finite number of regions, such that when there is at least one sample point in each of these regions, we would have that  $B_t \subset \mathcal{H}(T_0)$ . Then we could give a lower bound to the probability of  $B_t \subset \mathcal{H}(T_0)$  in terms of the number of regions and the volume of each of these regions.

By dividing the first  $d - 1$  coordinates into  $\frac{M}{\delta}$  equal parts, we partition the region  $[0, M]^d$  into rectangles, each with side length  $\delta$ , except for the  $d$ -th dimension (the  $\delta$  here is not exactly the  $\delta$  in the statement of the lemma, since we will do a change of variable in the last step). To be more precise, the rectangles are given by

$$Z_j = \left( \prod_{i=1}^{d-1} [(j_i - 1)\delta, j_i\delta] \right) \times [0, M].$$

Here  $j \in J$  and  $J$  is defined by

$$J := \left\{ j = (j_1, \dots, j_{d-1}), j_i = 1, 2, \dots, \frac{M}{\delta} \right\}.$$

Denote by  $J_0$  the set which consists of  $j \in J$  such that there exist a point in  $B_{2\delta}$  whose first  $d - 1$  coordinates are  $j_1\delta, j_2\delta, \dots, j_{d-1}\delta$  respectively, i.e.,  $J_0 = \left\{ j \in J : B_{2\delta} \cap \left( \left( \prod_{i=1}^{d-1} [j_i\delta, (j_i+1)\delta] \right) \times [0, M] \right) \neq \emptyset \right\}$ . For all  $j \in J_0$ , let  $p_j$  be the point such that

- i)  $p_j \in B_\delta$
- ii) The first  $d - 1$  coordinates of  $p_j$  are  $j_1\delta, j_2\delta, \dots, j_{d-1}\delta$  respectively
- iii)  $p_j$  has  $d$ -th coordinate larger than  $-\delta + \sup_{p \text{ satisfies i), ii)}} (d\text{-th coordinate of } p)$ .

From the definition of  $J_0$  and the fact that  $B_\delta \supset B_{2\delta}$ ,  $p_j$  is guaranteed to exist. We claim that  $B_{2\delta} \cap Z_j \subset \mathcal{R}(p_j)$ , where  $\mathcal{R}(p_j)$  is the rectangle that contains 0 and  $p_j$  as two of its corners. Clearly, from the definition of  $Z_j$ , for any point  $x \in B_{2\delta} \cap Z_j$ , its first  $d - 1$  coordinates are smaller than  $j_1\delta, j_2\delta, \dots, j_{d-1}\delta$  respectively. For the  $d$ -th coordinate, suppose on the contrary that there exists  $x \in B_{2\delta} \cap Z_j$  with  $d$ -th coordinate greater than the  $d$ -th coordinate of  $p_j$ . Since  $x \in Z_j$ , the first  $d - 1$  coordinates of  $x$  are at least  $(j_1 - 1)\delta, (j_2 - 1)\delta, \dots, (j_{d-1} - 1)\delta$ , so we have that  $x + \delta\mathbf{1}_{d \times 1} \geq p_j + \delta e_d$ . Since  $x \in B_{2\delta}$ , we know that  $x + 2\delta\mathbf{1}_{d \times 1} \in \mathcal{S}_\gamma^c$ . Hence by the previous inequality and the orthogonal monotonicity of  $\mathcal{S}_\gamma$ ,  $p_j + \delta e_d + \delta\mathbf{1}_{d \times 1} \in \mathcal{S}_\gamma^c$ . By definition of  $B_\delta$ , this implies  $p_j + \delta e_d \in B_\delta$ . This contradicts iii) in the definition of  $p_j$ . By contradiction, we have shown that each point in  $B_{2\delta} \cap Z_j$  has  $d$ -th coordinate smaller than the  $d$ -th coordinate of  $p_j$ . So the claim that  $B_{2\delta} \cap Z_j \subset \mathcal{R}(p_j)$  for any  $j \in J_0$  is proved.

Then we consider those  $j$  such that  $j \in J - J_0$ . For any point  $x \in Z_j$ , the first  $d - 1$  coordinates of  $x + \delta\mathbf{1}_{d \times 1}$  are at least  $j_1\delta, j_2\delta, \dots, j_{d-1}\delta$  respectively. Since  $j \notin J_0$ , we have that  $x + \delta\mathbf{1}_{d \times 1} \notin B_{2\delta}$ . This implies  $x + 3\delta\mathbf{1}_{d \times 1} \notin \mathcal{S}_\gamma^c$ , or  $x \notin B_{3\delta}$ . So we have shown that for any  $j \notin J_0$ ,  $B_{3\delta} \cap Z_j = \emptyset$ . This implies  $B_{3\delta}$  has a partition given by  $B_{3\delta} = \cup_{j \in J} (B_{3\delta} \cap Z_j) = \cup_{j \in J_0} (B_{3\delta} \cap Z_j)$ . Notice that  $B_{3\delta} \subset B_{2\delta}$ , from the result in the preceding paragraph, we conclude that  $B_{3\delta} \subset \cup_{j \in J_0} \mathcal{R}(p_j)$ .

For each  $j \in J_0$  and the constructed  $p_j$ , consider the region

$$G_j := \{x \in S_\gamma^c : x \geq p_j\}.$$

Observe that, if there exists a sample point in  $T_0$  that lies in  $G_j$ , then we have  $p_j \subset \mathcal{H}(T_0)$  which implies  $\mathcal{R}(p_j) \subset \mathcal{H}(T_0)$ . Since  $p_j \in B_\delta$  and  $\mathcal{S}_\gamma$  is orthogonally monotone, we have that  $G_j$  contains the rectangle which contains  $p_j$  and  $p_j + \delta \mathbf{1}_{d \times 1}$  as two of its corners, so  $\text{Vol}(G_j) \geq \delta^d$ . Hence the probability that  $\mathcal{R}(p_j) \subset \mathcal{H}(T_0)$  has a lower bound given by

$$P(\mathcal{R}(p_j) \subset \mathcal{H}(T_0)) \geq P(T_0 \cap G_j \neq \emptyset) \geq 1 - (1 - \delta^d q_l)^{n_1} \geq 1 - e^{-n_1 q_l \delta^d}.$$

Notice that  $|J_0| \leq \left(\frac{M}{\delta}\right)^{d-1}$ , by union bound we have that

$$P(\cup_{j \in J_0} \mathcal{R}(p_j) \subset \mathcal{H}(T_0)) \geq 1 - \frac{M^{d-1}}{\delta^{d-1}} e^{-n_1 q_l \delta^d}.$$

Since we have shown that  $B_{3\delta} \subset \cup_{j \in J_0} \mathcal{R}(p_j)$ , this implies

$$P(B_{3\delta} \subset \mathcal{H}(T_0)) \geq 1 - \frac{M^{d-1}}{\delta^{d-1}} e^{-n_1 q_l \delta^d}.$$

Based on this inequality, it is not hard to check that for  $t(\delta, n_1) = 3 \left( \frac{\log(n_1 q_l) + d \log M + \log \frac{1}{\delta}}{n_1 q_l} \right)^{\frac{1}{d}}$ , we have that  $P(B_{t(\delta)} \subset \mathcal{H}(T_0)) \geq 1 - \delta$ .  $\square$

*Proof of Theorem [6](#).* First, we show the inequality in the theorem, i.e.,  $P_{X \sim q}(X \in \mathcal{H}(T_0)^c \setminus \mathcal{S}_\gamma) \leq M^{d-1} q_u \left(\frac{\sqrt{d}}{2}\right)^{d-1} w_{d-1} t(\delta, n_1)$ . It suffices to show that with probability at least  $1 - \delta$ ,  $\text{Vol}(\mathcal{H}(T_0)^c \setminus \mathcal{S}_\gamma) \leq M^{d-1} \left(\frac{\sqrt{d}}{2}\right)^{d-1} w_{d-1} t(\delta, n_1)$ , or equivalently  $\text{Vol}(\mathcal{S}_\gamma^c \setminus \mathcal{H}(T_0)) \leq M^{d-1} \left(\frac{\sqrt{d}}{2}\right)^{d-1} w_{d-1} t(\delta, n_1)$ . Since by lemma [1](#) we have that  $B_{t(\delta, n_1)} \subset \mathcal{H}(T_0)$  with probability at least  $1 - \delta$ , it suffices to show that  $\text{Vol}(\mathcal{S}_\gamma^c \setminus B_{t(\delta, n_1)}) \leq M^{d-1} \left(\frac{\sqrt{d}}{2}\right)^{d-1} w_{d-1} t(\delta, n_1)$ . This latter inequality actually follows from the definition of  $B_{t(\delta, n_1)}$  and some geometric argument. Indeed, by definition of  $B_{t(\delta, n_1)}$ , for each  $x \in \mathcal{S}_\gamma^c \setminus B_{t(\delta, n_1)}$ ,  $x$  belongs to the area which is obtained by moving the boundary of  $\mathcal{S}_\gamma$  in direction  $-\frac{\mathbf{1}_{d \times 1}}{\sqrt{d}}$  for a distance of  $t(\delta, n_1) \sqrt{d}$ . So the volume of  $\mathcal{S}_\gamma^c \setminus B_{t(\delta, n_1)}$  is bounded by

$$\begin{aligned} & t(\delta, n_1) \sqrt{d} \times \text{Vol}_{d-1}(\text{projection of the boundary of } S_0 \text{ in direction } \mathbf{1}_{d \times 1}) \\ & \leq t(\delta, n_1) \sqrt{d} \times \text{Vol}_{d-1}(\text{projection of } [0, M]^d \text{ in direction } \mathbf{1}_{d \times 1}) \end{aligned}$$

Here  $\text{Vol}_{d-1}$  means computing volume in the  $d - 1$  dimensional space. Notice that  $[0, M]^d$  is contained in a ball with radius  $\frac{M\sqrt{d}}{2}$ , we have that

$$\text{Vol}_{d-1}(\text{projection of } [0, M]^d \text{ in direction } \mathbf{1}_{d \times 1}) \leq M^{d-1} \left(\frac{\sqrt{d}}{2}\right)^{d-1} w_{d-1}.$$

Combining the preceding two inequalities, we have proved the inequality in the theorem. Next we show the equality in the theorem. Indeed, when  $d$  is large, we have asymptotic formula  $w_d = \frac{1}{\sqrt{d\pi}} \left(\frac{2\pi e}{d}\right)^{\frac{d}{2}} (1 + O(d^{-1}))$ . Plugging this into the RHS above, we will obtain the asymptotic bound as stated in the theorem.  $\square$

*Proof of Theorem [4](#).* By Markov inequality and the definition of  $h, \bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$ , we know that

$$P_{X \sim q}(X \in \bar{\mathcal{S}}_\gamma^{\hat{\kappa}}, X \in \mathcal{S}_\gamma^c) = P_{X \sim q}(\hat{g}(X) \geq \hat{\kappa}, X \in \mathcal{S}_\gamma^c) \leq \frac{R(\hat{g})}{h(\hat{\kappa})}. \quad (10)$$

We will compare the numerator and denominator of the RHS of (10) with their counterparts for the true minimizer  $g^*$ . For the numerator, since  $\hat{g}$  is the empirical risk minimizer, we have that

$$\begin{aligned} R(\hat{g}) &\leq R_{n_1}(g) + \sup_{g_\theta \in \mathcal{G}} |R_{n_1}(g_\theta) - R(g_\theta)| \leq R_{n_1}(g^*) + \sup_{g_\theta \in \mathcal{G}} |R_{n_1}(g_\theta) - R(g_\theta)| \\ &\leq R(g^*) + 2 \sup_{g_\theta \in \mathcal{G}} |R_{n_1}(g_\theta) - R(g_\theta)|. \end{aligned}$$

For the denominator, from the definition of  $\bar{\mathcal{S}}_\gamma^{\hat{\kappa}}$ , it is not hard to verify that, in Algorithm 1, our choice of  $\hat{\kappa}$  is given by  $\hat{\kappa} = \min\{\hat{g}(x) : x \in \mathcal{H}(T_0)^c\}$ . By lemma 1, we have that with probability at least  $1 - \delta$ ,  $B_{t(\delta, n_1)} \subset \mathcal{H}(T_0)$ , which implies that with probability at least  $1 - \delta$ ,

$$\begin{aligned} \hat{\kappa} &\geq \min\{\hat{g}(x) : x \in B_{t(\delta, n_1)}^c\} \geq \min\{g^*(x) : x \in B_{t(\delta, n_1)}^c\} - \|\hat{g} - g^*\|_\infty \\ &\geq \min\{g^*(x) : x \in \mathcal{S}_\gamma\} - t(\delta, n_1)\sqrt{d}\text{Lip}(g^*) - \|\hat{g} - g^*\|_\infty \\ &= \kappa^* - t(\delta, n_1)\sqrt{d}\text{Lip}(g^*) - \|\hat{g} - g^*\|_\infty. \end{aligned}$$

Putting the preceding two inequalities into the Markov inequality (10), and notice that  $h$  is non decreasing by its definition, the theorem is proved.  $\square$