

# A Survey of Learning Criteria Going Beyond the Usual Risk

Matthew J. Holland

MATTHEW-H@AR.SANKEN.OSAKA-U.AC.JP

Kazuki Tanabe

TANABE@AR.SANKEN.OSAKA-U.AC.JP

Osaka University

Ibaraki, Osaka 567-0047 Japan

## Abstract

Virtually all machine learning tasks are characterized using some form of loss function, and “good performance” is typically stated in terms of a sufficiently small *average* loss, taken over the random draw of test data. While optimizing for performance on average is intuitive, convenient to analyze in theory, and easy to implement in practice, such a choice brings about trade-offs. In this work, we survey and introduce a wide variety of non-traditional criteria used to design and evaluate machine learning algorithms, place the classical paradigm within the proper historical context, and propose a view of learning problems which emphasizes the question of “what makes for a desirable loss distribution?” in place of tacit use of the expected loss.

## 1. Introduction

Over the past two decades, the rapid and widespread adoption of machine learning tools has provided a major impetus to design new methods and develop a more refined methodology. It is only natural that as machine learning techniques are brought to bear on new and increasingly diverse problems, it becomes necessary to re-define what we mean by “good” performance. Put roughly, within the classical paradigm of statistical machine learning, “good” invariably means “small average test loss.” While this is a perfectly natural definition, and in some cases precisely what we want to achieve, prioritizing the average test loss means that a lot of other important aspects of performance are abstracted away, and left without any guarantee. As a result, even when a solution is optimal on average, the results may be highly unsatisfactory. Here are a few examples to emphasize this point:

- **High-confidence guarantees.** Assuming just finite variance, in terms of the *average* squared estimation error of the mean, the sample mean is minimax optimal (Bickel and Lehmann, 1981; Lugosi and Mendelson, 2019a). However, in terms of *high-probability* error bounds, the sample mean has exponentially worse dependence on the confidence level than more robust alternatives (Catoni, 2012; Devroye et al., 2016).
- **Balanced classification.** Under extreme class imbalance, even the best possible misclassification probability can result in very poor performance on minority classes, leading to a large *balanced error* (Menon et al., 2021).
- **Fair decisions.** Even if a dataset does not include sensitive demographic information such as ethnicity, gender, or economic status, optimal performance on average can lead to highly sub-optimal performance for under-represented groups (Hashimoto et al., 2018).

- **Distribution shift.** Superior performance on average during training could be disastrous at test time if the underlying data distribution has changed (Curi et al., 2020; Quiñonero Candela et al., 2009; Sutter et al., 2021).
- **Resilience to rare events.** Decisions which are optimal on average need not be resilient to rare but catastrophic “black swan” events (Lee et al., 2020; Taleb, 2008).

These issues highlight the fact that a successful machine learning application need not always be characterized by achieving the minimum average loss. This broadened perspective has driven the machine learning community to develop a wide variety of new techniques to ensure learning algorithms reliably provide desirable solutions under more diverse notions of success. In particular, over the past five years, clear trends in algorithm design have emerged in which the underlying loss function is left abstract, and *properties of the loss distribution* are the primary focus. Our chief goal in this paper is to elucidate and bring attention to this trend. We survey the recent literature, provide some historical context, and use a general formulation that allows us to concisely express a wide variety of learning problems as special cases.

### 1.1 The Classical Paradigm

To make it clear precisely what this paper is about, we start with a concise formulation of the traditional learning problem. In the classical paradigm, one seeks the best possible performance *on average*. Put simply, this refers to the fact that in the vast majority of statistical learning problems, the ultimate goal is formulated in terms of minimizing the expected “test” loss. More formally, letting  $\mathbf{Z} \sim \mu$  denote an as-yet unobserved random (test) data point,  $\mathcal{H}$  a set of possible decisions, and  $\ell(\cdot)$  a loss function which assigns a numerical penalty to (decision, data) pairs, the ultimate goal is to find a solution  $h^*$ , which may or may not belong to  $\mathcal{H}$ , that satisfies

$$\mathbf{E}_\mu \ell(h^*; \mathbf{Z}) \leq \inf_{h \in \mathcal{H}} \mathbf{E}_\mu \ell(h; \mathbf{Z}) + \varepsilon \quad (1)$$

for some  $\varepsilon \geq 0$ , which represents our tolerance level. The expected value of a random loss is traditionally called the *risk*, a term with a long history (see §4 for details). Depending on the actual realization of  $\mathbf{Z}$ , the quality of any  $h$  as measured by  $\ell(h; \mathbf{Z})$  is subject to variability. While performance could certainly be catastrophic in the worst case, the milder condition (1) simply requires that on average, the decision made is optimal.<sup>1</sup> Even so, the problem is still far from trivial, since  $\mu$  is unknown. This is where “training” data comes in. Let us denote a sequence of random training data by  $\mathbf{Z}_n := (\mathbf{Z}_1, \dots, \mathbf{Z}_n)$ . Obviously, if there is no relation between the distributions of  $\mathbf{Z}_n$  and  $\mathbf{Z}$ , we have no hope of finding a solution satisfying (1). A strong relation can be made by the following assumption:

$$\mathbf{Z}_n \text{ is an independent random sample of } n \text{ points drawn from } \mu. \quad (2)$$

Seeking a solution to (1) under the assumption (2) is precisely what Vapnik (1999, Ch. 1) calls the “general setting of the learning problem.” With this link between training and

---

1. It is sometimes meaningless to ask for worst-case optimality over the random draw of test data. For example, consider estimating the mean of  $\mathbf{X} \sim \text{Normal}(0, 1)$ . For any choice of  $\theta \in \mathbb{R}$ , we always have  $\sup_{\mathbf{X}} (\theta - \mathbf{X})^2 = \infty$ .

test data in place, the most natural strategy is to replace the unknown  $\mu$  in (1) with the empirical distribution induced by  $\mathbf{Z}_n$ , and optimize that objective instead; this is *empirical risk minimization (ERM)*. More formally, an ERM solution is any decision  $H_n$  satisfying the following property:

$$H_n \in \arg \min_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \ell(h; Z_i). \quad (3)$$

Intuitively, the hope is that a solution which is optimal at training time will also be optimal at test time. Of course, this is an idealized goal, since with limited samples, true optimality is not statistically plausible. Furthermore, the off-sample quality depends entirely on the nature of the random training data that happens to materialize. It is always possible to get a very unlucky sample  $\mathbf{Z}_n$  which is not at all representative of  $\mu$  and leads to a poor candidate  $H_n$ , but under the assumption (2), this cannot occur with high probability. With this in mind, it is common to relax the requirement (1), and ask that for some  $\varepsilon > 0$  we have

$$\mathbf{P} \left\{ \mathbf{E}_\mu \ell(H_n; \mathbf{Z}) \leq \inf_{h \in \mathcal{H}} \mathbf{E}_\mu \ell(h; \mathbf{Z}) + \varepsilon \right\} \geq 1 - \delta \quad (4)$$

with  $0 < \delta < 1$  controlling the desired level of confidence, noting that probability is with respect to  $\mathbf{Z}_n$  and any algorithmic randomness that choosing  $H_n$  entails.<sup>2</sup> A great deal of statistical learning theory amounts to finding a bound on the sample size  $n$  that is sufficient or necessary for  $H_n$ , the output of a  $\mathbf{Z}_n$ -dependent learning algorithm, to satisfy the performance requirement (4) over various classes of data distributions  $\mu$ , decision spaces  $\mathcal{H}$ , and loss functions  $\ell$  (Anthony and Bartlett, 1999; Kearns and Vazirani, 1994; Mohri et al., 2012; Vapnik, 1998, 1999).

## 1.2 Our Viewpoint

On one hand, the breadth of learning problems captured within the classical paradigm of §1.1 is tremendous. Even with  $\mathcal{H}$  fixed, different tasks can be characterized by distinct loss functions, and a massive body of knowledge regarding effective loss function design has been developed in the statistical machine learning community, with countless refinements and extensions made to suit the needs of applied fields such as computer vision and natural language processing.

On the other hand, the class of problems captured within §1.1 can also be seen as being severely limited by the (often tacit) assumption that  $\mathbf{E}_\mu \ell(\cdot; \mathbf{Z})$  is the appropriate criterion to use for a given loss function. Taking a broader viewpoint, the expected value  $\mathbf{E}_\mu \ell(h; \mathbf{Z})$  is just one “property” of the distribution of the random loss  $\ell(h; \mathbf{Z})$  incurred by decision candidate  $h \in \mathcal{H}$ . As the machine learning methodology develops and matures, we believe that the question of *which property to emphasize as the criterion for learning* will play a critical role, in some cases orthogonal to the choice of  $\mathcal{H}$  and  $\ell$ .

---

2. When  $H_n \notin \mathcal{H}$  is allowed under (4), the learning problem is sometimes called *improper* (Shalev-Shwartz and Ben-David, 2014, Ch. 3).

Making this all a bit more explicit, we take  $\ell$  and  $\mathcal{H}$  as given, and focus on the random losses under the test data distribution  $\mu$ , denoted by

$$\mathcal{L}_\mu(\mathcal{H}) := \{\ell(h; Z) : Z \sim \mu, h \in \mathcal{H}\}. \tag{5}$$

Note also that the learning problem described in (1) can be equivalently re-stated in terms of *choosing a desirable loss distribution*, that is, seeking  $L_*$  such that

$$\mathbf{E}_\mu L_* \leq \mathbf{E}_\mu L + \varepsilon, \quad \forall L \in \mathcal{L}_\mu(\mathcal{H}). \tag{6}$$

Here “desirability” is measured by reducing a random variable to a one-number summary, i.e., the mapping  $L \mapsto \mathbf{E}_\mu L \in \bar{\mathbb{R}}$ . If learning is broadly understood as an attempt at decision-making which leads to a desirable off-sample loss distribution, then aside from technical convenience and historical precedent, there is in principle no reason to restrict ourselves to the mean. For example, any valid median, the variance, skewness, kurtosis, inter-quartile range, weighted sums of these quantities; these are all elementary examples of properties which might be considered when making a broadened notion of “desirable loss distribution” precise.

In this paper, we generalize the class of learning problems described in §1.1 as follows. Let  $C_\mu : \mathcal{L}_\mu(\mathcal{H}) \rightarrow \bar{\mathbb{R}}$  be any function which maps random variables  $L \in \mathcal{L}_\mu(\mathcal{H})$  to the extended real line, and for readability let us overload the notation setting  $C_\mu(h) := C_\mu(\ell(h; Z))$  for each  $h \in \mathcal{H}$ . We call  $C_\mu(\cdot)$  a *learning criterion* function, and refer to the value  $C_\mu(L)$  as a *property* of the distribution of  $L$ . Finally, fix a subset  $\mathcal{C} \subset \bar{\mathbb{R}}$  that we call the *desirability set*. With these parts in place, we expand upon the traditional learning problem of (4) using the following general condition on algorithm output  $H_n$ :

$$\mathbf{P} \{C_\mu(H_n) \in \mathcal{C}\} \geq 1 - \delta. \tag{7}$$

In words, this generalized setup can be succinctly summarized as follows.

**Generalized learning problem:**

“Reliably achieve a test loss distribution with sufficiently desirable properties.”  
 $\underbrace{\geq 1-\delta}_{\geq 1-\delta} \quad \underbrace{L \in \mathcal{L}_\mu(\mathcal{H})}_{L \in \mathcal{L}_\mu(\mathcal{H})} \quad \underbrace{\mathcal{C}}_{\mathcal{C}} \quad \underbrace{C_\mu}_{C_\mu}$

Moving forward, we shall use this general formulation as a vantage point from which we can view recent conceptual and technical trends in machine learning as special cases. Our focus in this paper is purely on the criterion  $C_\mu(\cdot)$ . Issues of loss function design, regularization, and model properties/constraints are all abstracted away, implicitly captured by  $\ell$ ,  $\mathcal{H}$ , and  $\mathcal{C}$ . For reference, however, here are some simple and concrete examples illustrating the flexibility of the formulation given here.

*Example 1* (Supervised learning). Let  $Z = (X, Y)$ , with  $X$  a random feature vector, and  $Y$  a random binary label taking values in  $\{-1, 1\}$  for example. Given a classifier  $h$ , the (random) hinge loss incurred by  $h$  is  $\max\{0, 1 - Yh(X)\}$ , and the zero-one loss is  $\mathbf{I}\{h(X) \neq Y\}$ . ■

*Example 2* (Unsupervised learning). Let  $X$  and  $h(X)$  be random vectors taking values on a linear space equipped with a norm  $\|\cdot\|$ , where  $h(X)$  satisfies some special constraints, e.g., a neural network whose intermediate layers provide a low-memory encoding of  $X$ , and use the distortion  $\|h(X) - X\|$  as a loss. ■

*Example 3* (Standard setup). By setting the learning criterion and desirability set as

$$C_\mu(L) = \mathbf{E}_\mu L, \quad \mathcal{C} = (-\infty, C_\mu^* + \varepsilon]$$

we get a familiar formulation of the learning problem, with some flexibility left in specifying  $\mathcal{C}$ , depending on how we set the threshold  $C_\mu^*$ . If we set

$$C_\mu^* = \inf \{C_\mu(h) : h \in \mathcal{H}\}$$

then we recover (4) as a special case of (7). If instead constraining ourselves to  $\mathcal{H}$ , we set

$$C_\mu^* = \inf \{C_\mu(h) : h \text{ satisfies } |C_\mu(h)| < \infty\}$$

then the desirability set  $\mathcal{C}$  encodes a requirement to achieve a small *approximation error* on top of a small *estimation error*, following the terminology of Cucker and Smale (2002). ■

*Example 4* (A less typical setup). One can in principle consider setting a threshold  $C_\mu^*$  about which we hope performance does not deviate too far, namely

$$\mathcal{C} = [C_\mu^* - \underline{\varepsilon}, C_\mu^* + \bar{\varepsilon}], \quad \bar{\varepsilon}, \underline{\varepsilon} \geq 0$$

where the width  $\bar{\varepsilon} + \underline{\varepsilon}$  controls the amount of deviation that we tolerate. Note that while modifying the design of  $C_\mu$  lets us encode preferences on the deviations of  $L$ , introducing  $\mathcal{C}$  in this form enables us to request small deviations over the random draw of  $\mathbf{Z}_n$  as well. ■

*Example 5* (Relative performance). Another interesting way to define desirability through  $C_\mu(\cdot)$  and a threshold  $C_\mu^*$  is to look at relative values. For example, letting  $a, b, c$  be non-negative real values, consider the desirability set

$$\mathcal{C} = \left\{ x \in \mathbb{R} : \frac{|x - C_\mu^*|}{a + b|x| + c|C_\mu^*|} \leq \varepsilon \right\}$$

where  $\varepsilon > 0$  once again specifies a tolerance level. See Haussler (1992) for an example of learning theoretical results stated in terms of this notion of desirability, with  $a > 0$  and  $b = c = 1$ . ■

### 1.3 Related Work

Now that we have a picture of what this paper is about, let us first take a bird’s-eye view of the context in which this survey was born, and highlight some topics that we do *not* cover.

Starting with some clear omissions, we do not delve into reinforcement learning at all, although the basic viewpoint put forward in §1.2 can clearly be extended to such problem settings; see Chow et al. (2016, 2017) for some examples. We also do not explicitly treat any Bayesian learning theory, but a great deal of such theory can indeed be captured by our formulation when  $\mathcal{H}$  is a set of sufficiently regular probability distributions over a parameter space, and  $\ell$  or  $\mathcal{H}$  are designed to enforce some degree of “proximity” to a prior distribution. To see this, consider for example the rich examples derived from the general “Bayesian Learning Rule” of Khan and Rue (2021). We do not delve into any details regarding the duality theory underlying convex learning criteria; see the very recent survey

of Royset (2022) for more detailed background. We also do not provide a comprehensive list of applied research featuring machine learning driven by novel learning criteria; the recent survey of Hu et al. (2022) provides a useful reference in this respect. Finally, since the underlying base loss function is left abstract, we do not consider adversarial methods which depend on particular forms of losses; see recent work by Robey et al. (2022) which interpolates between average and worst-case criteria by controlling perturbations of the inputs used in supervised learning tasks.

We used the broad term “property” to refer to the output of criterion  $C_\mu$  in §1.2, and while we believe this general notion is not yet widespread in the machine learning community, the idea of defining various characteristics of interest for an unknown (non-parametric) distribution is certainly very well-established in statistics.<sup>3</sup> It should also be noted that the term *property* has been used in a formal way, and that there is a broad literature on the *elicitation* of properties, i.e., characterizing properties which can be expressed as the solution to an expected loss minimization problem. An important subset of such properties are the “M-parameters” as described lucidly by Koltchinskii (1997), an extension of classical M-estimators (Huber, 1964). The literature on property elicitation begins in statistics and economics (Lambert et al., 2008; Savage, 1971), but important connections to surrogate complexity in machine learning are being established (Finocchiaro et al., 2021). Elicitation and surrogate function design are deep and interesting topics, especially in the generalized learning setup of §1.2, but in this paper, we will not delve further in this direction.

As emphasized in the previous sub-section, the central idea here is that of putting the loss distribution at the forefront, making “desirable properties” a key design decision, rather than falling back on the tacit assumption that a small mean is always good enough. While up until recently it has remained rather obscure, this notion has existed in the machine learning community for decades. For a lucid example in the context of binary classification, Breiman (1999) gave an example showing how a uniformly higher margin distribution need not imply better off-sample classification performance. Reyzin and Schapire (2006) built upon this example, showing how naive margin maximization tends to be inferior to algorithms which “do a better job with the overall distribution.” Two decades later, the idea of taking a broader view of loss distributions is starting to receive widespread attention, in terms of both designing and evaluating algorithms. One stark case is the work of Agarwal et al. (2021), who used tail-trimmed confidence intervals to effectively infer about deep reinforcement learning algorithm performance after just “a handful of runs,” receiving an Outstanding Paper Award at NeurIPS 2021.

The learning criterion  $C_\mu$  is something to be explicitly *designed*, and coupled with desirability conditions encoded by  $\mathcal{C}$ , it ultimately encodes an *optimization* task to be solved by the machine. In this direction, the massive literature on *risk functions* (also *risk measures*) in financial settings is of great relevance. Letting  $L \in \mathcal{L}_\mu(\mathcal{H})$  represent an uncertain financial outcome (e.g., negative returns of a certain portfolio), a succinct numerical summary of this uncertainty via  $L \mapsto C_\mu(L)$  is important for guiding principled decision-making under uncertainty. The notion of *coherent* risk functions as axiomatized by Artzner et al. (1999) has been extremely influential, and important special cases such as *conditional value-at-risk (CVaR)* have been central to both theory and practice of financial risk design. Sufficiently

---

3. Work by Lehmann, Bickel, and Huber in the 1960s and 1970s was particularly influential in this respect. See the commentary of Oja (2012) for some illuminating historical context.

regular risk functions (e.g., coherent risks) can be used to characterize *deviation* properties (Rockafellar et al., 2006), and the “risk quadrangle” of Rockafellar and Uryasev (2013) gives a detailed systematic account of functions for quantifying risk and deviations, linking them up with notions of “regret” and “error.” In a sense, the basic viewpoint and degree of generality in our formulation from §1.2 is also present in (Rockafellar and Uryasev, 2013), as is evident in the following quote:

A decision [results in] *random variables*... which can only be *shaped in their distributions* through the [decision made], not pinned down to specific values. Now there is no longer a single, evident answer to how optimization should be viewed, but risk measures can come to the rescue.

— Rockafellar and Uryasev (2013, §5)

Traditionally, decisions are made so as to minimize risk, however it may be defined. Our viewpoint differs slightly in that all four “corners” of the risk quadrangle are homogenized and placed on equal footing as potential candidates for the learning criterion  $C_\mu$ .

Our general formulation of  $\mathbf{L}$  being “sufficiently desirable” in terms of  $C_\mu(\mathbf{L}) \in \mathcal{C}$  has been influenced by the well-known learning model of Haussler (1992). In this model,  $C_\mu(\mathbf{L}) = \mathbf{E}_\mu \mathbf{L}$  is assumed throughout (called the “risk”), but another notion called “regret” is introduced, a non-negative value which evaluates the desirability of  $\mathbf{L}$  as measured by  $\mathbf{E}_\mu \mathbf{L}$ . Using Haussler’s definition, the (random) regret in the formulation (4) is the Bernoulli random variable

$$\mathbf{I} \left\{ \mathbf{E}_\mu \ell(\mathbf{H}_n; \mathbf{Z}) \leq \inf_{h \in \mathcal{H}} \mathbf{E}_\mu \ell(h; \mathbf{Z}) + \varepsilon \right\},$$

whereas in our generalized condition (7), it is

$$\mathbf{I} \{ C_\mu(\mathbf{H}_n) \in \mathcal{C} \}.$$

Since Haussler’s full learning problem requires us to choose  $\mathbf{H}_n$  to minimize the expected regret over the draw of  $\mathbf{Z}_n$ , our setup in §1.2 is a special case of the framework in (Haussler, 1992) whenever we restrict ourselves to  $C_\mu(\mathbf{L}) = \mathbf{E}_\mu \mathbf{L}$ . On the other hand, with the freedom to design  $C_\mu(\cdot)$  as we like, we are able to go well beyond the original Haussler model.

### 1.4 Clerical Matters

We use *plain italics* to indicate emphasis, and *colored italics* to indicate terms that have formal definitions, either within this paper or in the works we have cited. We use “■” to indicate the end of examples and remarks. We distinguish random quantities by using sans-serif typeface, e.g.,  $\mathbf{X}$ ,  $\mathbf{Z}$ ,  $\mathbf{Z}_i$ ,  $\mathbf{H}_n$ ,  $\widehat{\mathbf{C}}_n$ , and so forth. Throughout this paper,  $\mu$  denotes the probability distribution of the underlying random data of interest; sometimes this will be  $\mathbf{X}$ , sometimes this will be  $(\mathbf{X}, \mathbf{Y})$ , sometimes this will be an abstract  $\mathbf{Z}$ . In any case, expectation will be denoted by  $\mathbf{E}_\mu$ . Probabilities on the other hand are expressed using a generic symbol  $\mathbf{P}$ , regardless of the underlying distribution. The event “random variable  $\mathbf{X}$  equals  $x$ ” is written  $\{\mathbf{X} = x\}$ , and the corresponding Bernoulli random variable (via indicator function) is  $\mathbf{I}\{\mathbf{X} = x\}$ . For any positive integer  $k$ , we write  $[k] := \{1, \dots, k\}$ , useful for indexing.

## 2. Core Concepts

From a historical perspective, there have always been clear motivations, both theoretical and practical, for introducing new learning criteria into the picture. In this section, we briefly highlight some of the key scenarios within which these functions have arisen. A more systematic and general exposition of important criterion classes will follow in §3.

### 2.1 Robustness Under Heavy-Tailed Data

As mentioned in §1, one of the most natural reasons to design novel objective functions is that the empirical mean is a sub-optimal estimator of the true mean when the data distribution is heavy-tailed. For random  $\mathbf{X} \sim \mu$  on  $\mathbb{R}$ , assuming finite variance, setting  $\theta^* = \mathbf{E}_\mu \mathbf{X}$ , we have

$$\mathbf{E}_\mu (\theta^* - \mathbf{X})^2 \leq \mathbf{E}_\mu (\theta - \mathbf{X})^2, \quad \forall \theta \in \mathbb{R}.$$

That is, the mean minimizes the expected squared deviations from  $\mathbf{X}$ . Given an independent random sample  $\mathbf{X}_1, \dots, \mathbf{X}_n$ , let  $\widehat{\mathbf{X}}_n$  denote an arbitrary estimator of  $\mathbf{E}_\mu \mathbf{X}$ . Under just the finite variance assumption, writing  $\sigma_\mu^2 := \text{var}_\mu \mathbf{X}$ , ideally we would like to have “sub-Gaussian” guarantees of the form

$$\mathbf{P} \left\{ |\widehat{\mathbf{X}}_n - \mathbf{E}_\mu \mathbf{X}| \leq a\sigma_\mu \sqrt{\frac{\log(1/\delta)}{n}} \right\} \geq 1 - \delta \tag{8}$$

for a sufficiently wide range of  $\delta \in (0, 1)$ , where  $a$  is a constant free of  $\mu$ ,  $n$ , and  $\delta$ ; for more on this topic, see Devroye et al. (2016). Unfortunately, it is well-known that under such a general setting, for small values of  $a$  (e.g.,  $a = 1$ ), the usual sample mean fails to satisfy such a property (Lugosi and Mendelson, 2019a, Thm. 1). The sample mean is of course an ERM solution under the squared error, namely

$$\bar{\mathbf{X}}_n \in \arg \min_{\theta \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n (\theta - \mathbf{X}_i)^2, \quad \bar{\mathbf{X}}_n := \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i.$$

If we consider these squared errors as losses  $\ell(\theta; \mathbf{X}_i) = (\theta - \mathbf{X}_i)^2$ , then a well-established approach to obtain a more robust estimator is to modify the loss function to have sub-quadratic growth in the limit. A typical approach is to choose an estimator  $\widehat{\mathbf{X}}_n$  satisfying

$$\widehat{\mathbf{X}}_n \in \arg \min_{\theta \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \rho \left( \frac{\theta - \mathbf{X}_i}{s} \right) \tag{9}$$

where  $\rho : \mathbb{R} \rightarrow [0, \infty)$  is convex, approximately quadratic around zero, but with linear growth in the limit ( $\rho(x) = \mathcal{O}(x)$  as  $|x| \rightarrow \infty$ ), and  $s > 0$  is a scaling parameter used to control the bias incurred by using a non-quadratic loss function (larger  $s$  means smaller bias). This basic “M-estimation” idea of course underlies many core techniques in robust statistics (Huber and Ronchetti, 2009), and an important study from Catoni (2012) showed how with appropriate settings of  $s$ , estimators given by (9) can be shown to satisfy (8) with nearly optimal constant  $a$ . Many other robust mean estimators satisfying the sub-Gaussian



property (8) have since been studied, including the re-discovered mean-of-means (Lerasle and Oliveira, 2011), the well-known trimmed mean (Lugosi and Mendelson, 2019b), and others; see Lugosi and Mendelson (2019a) for an up-to-date survey.

Since we know that there are practical algorithms for achieving the sub-Gaussian error bounds of the form (8), it is natural to consider such algorithms as *sub-routines* that let us modify the naive ERM given by (3). That is, going back to our general setup described in §1, with  $Z \sim \mu$  and  $\mathbf{Z}_n$  satisfying (2), writing  $L(h) := \ell(h; Z)$  and  $L_i(h) := \ell(h; Z_i)$  now for an arbitrary “base” loss function  $\ell(\cdot)$ , at the very least we know that pointwise in  $h \in \mathcal{H}$ , we can construct an estimator using

$$\widehat{C}_n(h) \in \arg \min_{\theta \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \rho \left( \frac{\theta - L_i(h)}{s} \right) \tag{10}$$

which letting  $\sigma_\mu^2(h) := \text{var}_\mu L(h)$ , we know satisfies

$$\mathbf{P} \left\{ |\widehat{C}_n(h) - \mathbf{E}_\mu L(h)| \leq a \sigma_\mu(h) \sqrt{\frac{\log(1/\delta)}{n}} \right\} \geq 1 - \delta. \tag{11}$$

This suggests the idea of replacing naive ERM algorithms (3) with alternatives of the form

$$H_n \in \arg \min_{h \in \mathcal{H}} \widehat{C}_n(h) \tag{12}$$

where  $\widehat{C}_n(h)$  is computed using (10), or any of the other robust mean estimation procedures mentioned earlier. This novel approach to designing a learning algorithm was studied in seminal work by Brownlees et al. (2015). This work is conceptually important because it rigorously demonstrates the statistical merits of using biased, robust estimators as objective functions. Statistically speaking, for learning tasks under heavy-tailed data, it is possible to show that any solution satisfying (12) enjoys strong guarantees *in the classical paradigm*, i.e., in terms of the expected loss condition (4), that are not possible for naive ERM (3).

Of course, we can also view the objective in (12) from the population level as well, namely we can define a learning criterion class as follows:

$$C_\mu(h; s) := \arg \min_{\theta \in \mathbb{R}} \mathbf{E}_\mu \rho \left( \frac{\theta - L(h)}{s} \right). \tag{13}$$

With proper choice of  $\rho$  and  $s > 0$ , one can control the bias such that good performance in terms of  $C_\mu(h; s)$  implies good performance *on average*. More generally, however, it is perfectly plausible to fix  $s$  and use  $C_\mu(h; s)$  as the primary criterion for performance. In any case, the chief limitation of this approach is computational; actual computation of  $H_n$  in (12) is left abstract, since the minimization of  $\widehat{C}_n(\cdot)$  given by (10) is a challenging optimization problem.<sup>4</sup> Since M-parameters (including all quantiles) are easy to compute but difficult to minimize directly, they are perhaps best suited as *evaluation* metrics to be coupled with a more convenient surrogate objective (Holland and Ikeda, 2017). For direct

---

4. Since the M-estimator is itself defined as the solution of a minimization problem, minimizing that quantity as a function of  $h \in \mathcal{H}$  is what is termed a *bi-level program*.

minimization, a natural alternative is that of criteria which offer a smooth approximation to quantiles, which are easy to optimize under typical base loss functions. Exponential smoothing or “tilting” of the loss distribution (Li et al., 2020, 2021b) can be used effectively to this end, and “inverted” variants of CVaR (Lee et al., 2020) can also be used to obtain a similar effect. We will revisit both of these examples in more detail in §3.

## 2.2 Variance Control

All else equal, smaller variance tends to make statistical inference easier. Once again we will use one-dimensional mean estimation as a motivating example, retaining the notation from §2.1, as well as the assumption (2). Assuming  $\mathbf{X} \sim \mu$  has finite variance  $\sigma_\mu^2$ , Chebyshev’s inequality gives us a simple error bound on the sample mean  $\bar{\mathbf{X}}_n$ , namely that

$$|\bar{\mathbf{X}}_n - \mathbf{E}_\mu \mathbf{X}| \leq \sigma_\mu \sqrt{\frac{1}{n\delta}} \tag{14}$$

holds with probability no less than  $1 - \delta$ , for any choice of  $0 < \delta \leq 1$ . If we further assume that  $\mathbf{X}$  takes values on the unit interval  $[0, 1] \subset \mathbb{R}$  with probability one, then once again with probability no less than  $1 - \delta$ , we have

$$|\bar{\mathbf{X}}_n - \mathbf{E}_\mu \mathbf{X}| \leq \sigma_\mu \sqrt{\frac{2 \log(2/\delta)}{n}} + \frac{\log(2/\delta)}{3n}. \tag{15}$$

This follows from a result known as Bennett’s inequality.<sup>5</sup> Both (14) and (15) show us that all else equal, a smaller variance can enable sharper guarantees. With this in mind, recalling the ERM condition (3) and the fact that in general we may have many ERM solutions, it is natural to prefer those with smaller variance (in terms of the loss distribution). Indeed, as the bound (15) suggests, it is well-established that achieving a small variance can be used to guarantee “fast rates” (i.e., better than  $\mathcal{O}(\sqrt{1/n})$ ) for empirical risk minimizers.<sup>6</sup> In terms of designing learning criteria, the simplest choice is just the sum

$$C_\mu(h) = \mathbf{E}_\mu L(h) + \sqrt{\sigma_\mu^2(h)}. \tag{16}$$

This property of the loss distribution dates back to classical work on portfolio optimization by Markowitz (1952) (originally without the square root on  $\sigma_\mu^2(h)$ ), and replacing  $\mu$  with the empirical distribution induced by  $\mathbf{Z}_n$ , solutions to the resulting learning criterion minimization were analyzed in influential work by Maurer and Pontil (2009). While statistically easy to interpret and advantageous in theory, much like the robust objective (10) from §2.1, naively trying to optimize for (16) can be computationally challenging. In particular, even when the loss  $h \mapsto \ell(h; \mathbf{Z})$  is convex, we cannot in general guarantee that  $h \mapsto C_\mu(h)$  in (16) is convex. Recently, important insights have been established linking mean-variance to certain objective functions from the field of *distributionally robust optimization (DRO)* (Duchi and Namkoong, 2019; Gotoh et al., 2018), in which the goal is to minimize the worst-case expected loss to be incurred over a *set* of distributions, rather than just one pre-specified  $\mu$ . These methods use asymmetric measures of deviation, which place more weight

5. See Boucheron et al. (2013, Ch. 2) for an authoritative modern reference.

6. See for example Duchi and Namkoong (2019) and the references therein.

on worst-case tails and preserve convexity. This DRO approach to stochastic optimization also has important applications to deal with data distribution shift, and will be introduced formally in §3.

### 2.3 Risk-Averse Learning

In contrast to robustness concerns discussed in §2.1, where the expected loss was too sensitive to the loss tails to allow for high-confidence guarantees, sometimes the tails are precisely what we want to prioritize. This is an example of *risk aversion*, a tendency to overweight bad events and underweight good events, where the terms “bad” and “good” may be subjective or have different meanings depending on the task at hand.<sup>7</sup> One particularly intuitive way to define a “bad” event is to use the quantiles of the loss distribution as a threshold. Formally, let us denote the distribution function of  $\mathbf{L}$  by  $F_\mu(u) := \mathbf{P}\{\mathbf{L} \leq u\}$ , and using this, we define the  $\beta$ -level *quantile* of  $\mathbf{L}$  as follows:<sup>8</sup>

$$Q_\beta(\mathbf{L}) := \min\{u \in \mathbb{R} : F_\mu(u) \geq \beta\}, \quad 0 < \beta < 1. \tag{17}$$

Since losses come with the interpretation of larger being worse, a natural “bad” event is something like  $\{\mathbf{L} \geq Q_{0.95}(\mathbf{L})\}$ . As such, the average loss incurred beyond this threshold is a perfectly natural class of learning criteria that generalizes beyond the mean:

$$C_\mu(h; \beta) = \mathbf{E}_\mu \left[ \mathbf{L}(h) \mathbf{I}\{\mathbf{L}(h) \geq Q_\beta(h)\} \right], \quad 0 \leq \beta < 1. \tag{18}$$

Once again, in (18) we overload our formal notation  $Q_\beta$  by defining  $Q_\beta(h) := Q_\beta(\mathbf{L}(h))$ , with the special case of  $\beta = 0$  defined as

$$C_\mu(h; 0) := \lim_{\beta \rightarrow 0^+} C_\mu(h; \beta) = \mathbf{E}_\mu \mathbf{L}(h).$$

Given a training sample  $\mathbf{Z}_n = (Z_1, \dots, Z_n)$  and sorting the training losses as before such that  $L_{[1]} \geq \dots \geq L_{[n]}$ , for any quantile level  $0 \leq \beta \leq 1$ , there exists a  $k \in [n]$  such that the empirical objective analogous to (18) can be written

$$\widehat{C}_n(h; k) = \frac{1}{n} \sum_{i=1}^k L_{[i]}(h).$$

A re-scaled variant known as the “average top- $k$ ” objective is given as

$$\frac{n}{k} \widehat{C}_n(h; k) = \frac{1}{k} \sum_{i=1}^k L_{[i]}(h) \tag{19}$$

and algorithms based on (19) have been studied by Shalev-Shwartz and Wexler (2016) as well as Fan et al. (2017a,b), though the issues of robustness mentioned in §2.1 are of course still relevant here, which has naturally led to robustified variants (Holland and Haress, 2021). Setting  $k = n$  returns the usual empirical mean, and  $k = 1$  yields a maximum value.

---

7. A classical reference from economics is Pratt (1964). See also Tversky and Kahneman (1992).  
 8. When  $F_\mu(\cdot)$  is not continuous,  $Q_\beta$  in (17) is sometimes distinguished as the “left” or “lower” quantile.

Assuming the distribution function  $F_\mu$  is continuous, we have  $\mathbf{P}\{L(h) > Q_\beta(h)\} = \mathbf{P}\{L(h) \geq Q_\beta(h)\} = 1 - \beta$ , and we can easily construct a (population) criterion corresponding to (19) as

$$C_\mu^{\text{CVaR}}(h; \beta) := \frac{C_\mu(h; \beta)}{\mathbf{P}\{L(h) \geq Q_\beta(h)\}} = \frac{\mathbf{E}_\mu L(h) \mathbf{I}\{L(h) \geq Q_\beta(h)\}}{1 - \beta} = \mathbf{E}_\mu [L(h) \mid L(h) \geq Q_\beta(h)]. \quad (20)$$

In both (19) and (20), we are clearly considering the mean of a pre-specified fraction of the largest losses. This quantity has been called the *expected shortfall* (Acerbi and Tasche, 2002a), and it is also well-known as *conditional value-at-risk (CVaR)* (Rockafellar and Uryasev, 2000), a natural name given the right-most form in (20).<sup>9</sup> The quantity given in (20) can be seen as a special case of a much larger class of criteria under the name *optimized certainty equivalent (OCE)*, and is in fact also a limiting case of DRO risk functions (Duchi and Namkoong, 2018). Both will be introduced with more generality in §3.

## 2.4 Fairness-Aware Learning

Since machine learning methods are increasingly being utilized to automate or support real-world decisions, it is natural to run into issues of ethics and in particular *fairness* to individuals or groups with certain traits. To give this broad concept an intuitive and clear formalization, assuming the underlying data  $Z \sim \mu$  is a tuple of one or more random variables, let  $S$  denote a “sensitive” variable, representing one or more of the random variables composing  $Z$ . In the context of supervised learning, where  $Z = (\mathbf{X}, Y)$  and we want to choose some  $h \in \mathcal{H}$  such that we reliably get a sufficiently good prediction  $h(\mathbf{X}) \approx Y$ , the sensitive variable  $S$  might be a subset of features in  $\mathbf{X}$ , or indeed the class label  $Y$ . For example, say  $Y$  is binary-valued and indicates success at graduating from university, and  $\mathbf{X}$  represents various attributes of an individual such as test grades in primary and secondary school, extra-curricular honors, gender, ethnicity, and family income. In this case,  $S$  might correspond to gender, ethnicity, and family income. If we take the traditional approach of optimizing for prediction accuracy (i.e., minimizing  $\mathbf{P}\{h(\mathbf{X}) \neq Y\}$ ), then it is of course plausible that  $S$  plays a key role in accurately predicting  $Y$ , but without intervention, this could result in individuals with certain traits being assigned predictions that are intuitively unfair, even if they are “optimal” in terms of accuracy.

To mitigate this tendency, we must either change the ultimate objective directly, or enforce some fairness constraints. There are numerous specific examples of fairness constraints in the literature, but one of the strongest is (statistical) *parity* (Dwork et al., 2012), which asks that  $h(\mathbf{X})$  and  $S$  be statistically independent. This is the same as requiring that  $h(\mathbf{X})$  and  $h(\mathbf{X}) \mid S$  have the same distribution. A closely related notion looks at *performance* through the lens of a loss function, i.e., asking that  $\ell(h; Z)$  and  $\ell(h; Z) \mid S$  have the same distribution. Note that this lets us go well beyond the supervised learning paradigm.<sup>10</sup>

9. Nomenclature in the financial risk literature is somewhat convoluted; the paper of Acerbi and Tasche (2002b) is a useful reference. For extensions beyond the continuous case, see Rockafellar and Uryasev (2002).

10. Consider the example of Samadi et al. (2018), where dimension reduction using standard PCA can be unfair in that low-education sub-groups incur a disproportionately high reconstruction error.

Asking for the marginal and conditional distributions to be identical is an extremely strong requirement; a much weaker constraint just asks that their expected values align, i.e., that

$$\mathbf{E}_\mu \ell(h; \mathbf{Z}) = \mathbf{E}_\mu [\ell(h; \mathbf{Z}) | \mathbf{S}] \quad (21)$$

hold with probability 1 over the random draw of  $\mathbf{S}$ . This notion of “fairness via subgroup risks” is formulated in a lucid way by Williamson and Menon (2019), though as they mention, the underlying notion is present in several works preceding theirs (e.g., Hashimoto et al. (2018)). While (21) is weaker than asking the marginal and conditional distributions to align, the goal of choosing  $h \in \mathcal{H}$  such that this equality holds is still not very realistic, and further relaxations are required to link fairness constraints up to practical learning algorithms.

In this direction, note that the right-hand side of (21) is a random variable depending on the draw of  $\mathbf{S}$ , but the left-hand side is a constant. As such, for (21) to hold, the random variable on the right-hand side must have zero variance:  $\text{var}_\mathbf{S}(\mathbf{E}_\mu [\ell(h; \mathbf{Z}) | \mathbf{S}]) = 0$ . Relaxing this, we could ask for  $\varepsilon$ -small variance, i.e., that  $\text{var}_\mathbf{S}(\mathbf{E}_\mu [\ell(h; \mathbf{Z}) | \mathbf{S}]) \leq \varepsilon$  holds for some  $\varepsilon > 0$  that we specify. Recalling our previous discussion of the benefits of small variance in §2.2, as a direct approach to this problem, we could naively introduce a learning criterion of the form

$$C_\mu(h) = \mathbf{E}_\mu \mathbf{L}(h) + \sqrt{\text{var}_\mathbf{S}(\mathbf{E}_\mu [\mathbf{L}(h) | \mathbf{S}])} \quad (22)$$

where once again we write  $\mathbf{L}(h) := \ell(h; \mathbf{Z})$  for readability. Just as in §2.2, using the symmetric variance as-is leads to computational difficulties, and more congenial alternatives such as DRO and CVaR have been put to effective use in this direction (Hashimoto et al., 2018; Williamson and Menon, 2019). A full description of these classes will be given in §3 to follow.

### 3. Key Classes of Learning Criteria

With the general learning problem formulated in §1.2 coupled with the motivating examples discussed in §2, we will now dive into an exposition of some of the most important classes of learning criteria that have arisen in and around the machine learning literature. For readability, we shall divide this section into several sub-sections, each of which dedicated to a certain class of interest.

#### 3.1 Order-Based Criteria

Arguably the most intuitive strategy to design criteria going beyond the usual expected value is to sort the (empirical) losses and assign weights depending on rank before summation. Sorting can be done in either ascending or descending order, as long as we are consistent with the weighting. We already saw examples of losses sorted in descending order in §1.3 and §2.3, and here we will consider sorting in ascending order, since that aligns us with the literature we take as context. With  $n$  observations, let  $\mathbf{L}_{(i)}(h)$  denote the  $i$ th-smallest loss under  $h$ . That is, the losses are sorted such that  $\mathbf{L}_{(1)}(h) \leq \dots \leq \mathbf{L}_{(n)}(h)$  holds. Letting  $\mathbf{w} := (w_1, \dots, w_n)$  denote non-negative weights, we can define a new objective function as

follows:

$$\widehat{C}_n(h; \mathbf{w}) = \sum_{i=1}^n w_i L_{(i)}(h). \tag{23}$$

One classical motivation for this is to discard errant values to aid in estimating the population location, typically the mean or median.<sup>11</sup> This can be achieved in a strict sense by setting  $w_i = 0$  for rank  $i$  above or below certain thresholds, such that a pre-specified fraction of the data is discarded. More generally, the idea of setting weights as a function of rank is an important idea with a long history in statistics.<sup>12</sup> Making this a bit more explicit, letting  $W : [0, 1] \rightarrow [0, \infty)$  denote a weighting function, the weights in (23) can be set as  $w_i = W(i/n)$ , for each  $i \in [n]$ . Denoting the empirical distribution function of the losses by  $F_n(u; h) := \sum_{i=1}^n I\{L_i(h) \leq u\}/n$ , note that we have  $F_n(L_{(i)}(h); h) = i/n$ , and thus (23) can be re-written as follows:

$$\widehat{C}_n(h; \mathbf{w}) = \sum_{i=1}^n W(F_n(L_{(i)})) L_{(i)} = \sum_{i=1}^n W(F_n(L_i)) L_i \tag{24}$$

where we have suppressed the dependence on  $h$  in the notation just for readability. Reading left to right, the second equality in (24) clearly holds since passing unsorted losses through the empirical CDF ensures they get the correct rank-based weighting. In moving from the preceding empirical objective to (population) learning criteria, the two forms given in (24) are indicative of two paths we can take. Since both paths are meaningful and have relevant literature, we will look at both of them, one at a time.

The first path is centered around sorted losses. To link up sorted losses to quantiles, recall our (population) definition of  $Q_\beta(L)$  in (17), and replace the true distribution function  $F_\mu$  with its empirical analogue  $F_n$ . This yields the empirical quantile (viewed as a function of  $\beta$ ) as follows:

$$Q_{\beta,n} := \min \{u \in \mathbb{R} : F_n(u) \geq \beta\}, \quad 0 < \beta \leq 1. \tag{25}$$

Ranked losses have a clear relation to the empirical quantiles. Note that for any  $0 < \beta \leq 1$ , we have  $Q_{\beta,n} = L_{(\lceil n\beta \rceil)}$ . Since  $\beta \mapsto Q_{\beta,n}(L)$  is just a simple step function, if we assume that  $W(\beta) \geq 0$  for  $\beta \in \{1/n, 2/n, \dots, 1\}$  and  $W(\beta) = 0$  elsewhere, we can once again re-write (23) in terms of quantiles using an integral form, namely

$$\widehat{C}_n(h; \mathbf{w}) = \int_0^1 W(\beta) Q_{\beta,n}(L) \, d\beta \tag{26}$$

where we set  $w_i = W(i/n)$  for each  $i \in [n]$ . The equality in (26) uses the fact that  $F_n(Q_{\beta,n}(L)) = \beta$  for all  $\beta \in \{1/n, 2/n, \dots, 1\}$ , plus the first equality in (24). With this (empirical) form in mind, a population variant follows:

$$C_\mu(L; W) := \int_0^1 W(\beta) Q_\beta(L) \, d\beta. \tag{27}$$

---

11. This is the “discard-average” in classic work by Daniell (1920), who also cites a 1912 work of Poincaré.  
 12. See for example Bickel and Lehmann (1975) and the references therein. A more modern take on the optimality of trimmed-means is given by Lugosi and Mendelson (2019b).

Clearly, the criterion in (27) is completely characterized by the weighting function  $W(\cdot)$ . It is also natural to extend the criterion definition to allow for any function  $W(\cdot)$  such that the integral on the right-hand side of (27) is finite (or at least well-defined). The general form given by (27) can be used to capture a wide variety of well-known properties. Linear combinations of order statistics, also known as “L-estimates,” have played an important role in robust statistics (Huber and Ronchetti, 2009). Applications to robust unsupervised learning, replacing average distortion with an “L-statistic” have also been studied recently by Maurer et al. (2020).

The second path, motivated by the right-most form in (24), places its emphasis on rank-based weighting (via a  $W$ -transformed distribution function) applied to *unsorted* losses. Learning criteria of the form

$$C_\mu(\mathbf{L}; W) = \mathbf{E}_\mu W(F_\mu(\mathbf{L})) \mathbf{L} \tag{28}$$

play the central role here. This form has intuitive appeal, since we are working with losses rather than quantiles, and concentration properties of the empirical CDF are well-known. That said, we cannot naively cast this as a traditional ERM-style learning problem. This is because for any  $i \in [n]$ , the empirical quantity  $W(F_n(L_i)) L_i$  depends on the entire sample  $\mathbf{Z}_n$ , and thus independence is lost. Getting around this requires  $W(\cdot)$  to satisfy certain continuity properties, and sharp concentration of  $F_n(\cdot; h)$  around  $F_\mu(\cdot; h)$ , uniform over  $h \in \mathcal{H}$ . Conditions for the uniform convergence of (24) to  $C_\mu(h; W)$  in (28) were studied by Khim et al. (2020) under bounded losses and Lipschitz continuous  $W(\cdot)$ , with  $C_\mu(\cdot)$  called the “L-risk.” An extension centered around Hölder-continuous functions defined on families of distribution functions was developed recently by Liu et al. (2022). Finally, it should be noted that when the underlying distribution function  $F_\mu$  is invertible, the two paths coincide, i.e.,  $C_\mu(\mathbf{L}; W)$  given in (27) and (28) are equal (Holland and Haress, 2022, Lem. 1).

To complement the preceding general exposition, we look at several prominent special cases that arise in the machine learning literature.

*Example 6* (CVaR). It is well-established that CVaR as seen in (20) is a special case of the form (27). In particular, we have

$$C_\mu^{\text{CVaR}}(h; \beta) = \int_0^1 W_\beta(u) Q_u(h) du, \quad W_\beta(u) := \frac{\mathbf{I}\{u \geq \beta\}}{1 - \beta} \tag{29}$$

for any choice of  $0 \leq \beta < 1$ .<sup>13</sup> Since this weighting function determines a probability measure over the unit interval, we are “averaging the quantile,” and since the left quantile is often referred to as the *value-at-risk*, CVaR is also sometimes referred to as the *average value-at-risk (AVaR)* (Ruszczyński and Shapiro, 2006a). In the context of machine learning, CVaR is perhaps the best-known alternative to the mean. An early application can be found in Kashima (2007), but analysis of stochastic learning algorithms under CVaR is far more recent; see Cardoso and Xu (2019); Soma and Yoshida (2020); Curi et al. (2020); Holland and Haress (2021) for representative works. ■

13. See for example Corollary 4.3 and Eqn. 3.3 of Acerbi and Tasche (2002b), or Shapiro (2013, Eqn. 6).

*Example 7* (Spectral risk). A natural extension of the form (29) is to consider a weighting function that places progressively more weight on the upper tails, satisfying

$$W(\cdot) \text{ is non-decreasing on } [0, 1] \text{ and } \int_0^1 W(u) \, du = 1. \tag{30}$$

Under the specifications of (30), learning criteria of the form (27) are called *spectral risks* (Acerbi, 2002). Spectral risk estimators are studied by Pandey et al. (2019), and learning theory for algorithms based on spectral risks is developed by Holland and Haress (2022) and Liu et al. (2022). ■

*Example 8* (Average top- $k$ ). Given a finite sample  $\mathbf{Z}_n = (Z_1, \dots, Z_n)$  and utilizing a weight function defined as

$$W_k(u) := \frac{\mathbb{I}\{u \geq (n - k)/n\}}{(k/n)}$$

ensures we cut off all losses below a certain rank  $k \in [n]$ , and then re-scale to get the arithmetic mean of the  $k$  largest losses. Spelling this all out a bit more explicitly, setting the weight function to  $W = W_k$  as just described in (26), we obtain

$$\int_0^1 W_k(u) \mathbf{Q}_{u,n}(h) \, du = \left(\frac{n}{k}\right) \int_{(n-k)/n}^1 \mathbf{Q}_{u,n}(h) \, du = \left(\frac{n}{k}\right) \sum_{i=n-k+1}^n \frac{\mathbf{L}_{(i)}(h)}{n} = \frac{1}{k} \sum_{i=1}^k \mathbf{L}_{[i]}(h).$$

That is, we recover the average top- $k$  empirical risk (19) as seen earlier in §2.3. As mentioned before, algorithms optimizing this objective has been studied in the recent machine learning literature (Fan et al., 2017a,b; Shalev-Shwartz and Wexler, 2016). A handy computational fact is that the map  $(x_1, \dots, x_n) \mapsto \sum_{i=1}^k x_{[i]}$  is convex (Ogryczak and Tamir, 2003), although in general it is not differentiable. For smooth approximations, see Example 9 in the next sub-section. ■

### 3.2 OCE Risks

We proceed by looking at a class of learning criteria with motivations quite different from the order-based criteria seen in the preceding §3.1. To provide this class with some appropriate historical context, let  $\mathbf{X} \sim \mu$  be a random variable representing some uncertain future “payout” (larger is better), and let  $f : \mathbb{R} \rightarrow \mathbb{R}$  be in charge of assigning “utility” values to raw payouts. Decision making under uncertainty is traditionally centered around the expected utility, namely  $\mathbf{E}_\mu f(\mathbf{X})$ . This fact is of course closely related to the classical paradigm of statistical learning theory discussed in §1.1.<sup>14</sup> With these elements in place, the *certainty equivalent* of  $\mathbf{X}$  under  $f$  is formally defined as

$$\text{CE}_\mu(\mathbf{X}; f) := \{\theta \in \mathbb{R} : f(\theta) = \mathbf{E}_\mu f(\mathbf{X})\}. \tag{31}$$

In words, from the perspective of  $f(\cdot)$ , receiving  $\theta \in \text{CE}_\mu(\mathbf{X}; f)$  with certainty is just as appealing as receiving  $\mathbf{X}$  with uncertainty, where the uncertainty is summarized by averaging.<sup>15</sup> Building upon this concept, if  $\theta \in \mathbb{R}$  represents a certain payment, there are many

14. For more on expected utility and decision-making under certainty, see §4.1.

15. See Mas-Colell et al. (1995, Defn. 6.C.2) for more background on this notion in the context of preferences.



situations in which committing to  $\theta$  leaves us with an uncertain payment of  $X - \theta$ . When  $\theta$  is something we can control, it is natural to “optimize” this value such that the overall outcome is as preferable as possible, at least on average. Arguably the simplest way to evaluate such a situation is purely in terms of expected utility, with the best possible outcome quantified as

$$\sup_{\theta \in \mathbb{R}} [f(\theta) + \mathbf{E}_\mu f(X - \theta)].$$

A slightly modified version has us evaluate certain outcomes in their raw units, and uncertain outcomes using utility, taking the expectation of their sum. More explicitly, we define

$$\overline{\text{OCE}}_\mu(X; f) := \sup_{\theta \in \mathbb{R}} [\theta + \mathbf{E}_\mu f(X - \theta)] \tag{32}$$

known as the *optimized certainty equivalent (OCE)*, originally proposed and studied by Ben-Tal and Teboulle (1986) assuming that  $f(\cdot)$  is concave, monotonic (non-decreasing), and “normalized” such that  $f(0) = 0$  and  $1 \in \partial f(0)$ .<sup>16</sup> Regarding the naming, (32) was originally called NCE (new certainty equivalent) in (Ben-Tal and Teboulle, 1986), but two decades later, the authors published an updated work on this quantity (Ben-Tal and Teboulle, 2007), in which it was re-named OCE.<sup>17</sup>

Shifting our mindset from positively-oriented “payouts” to negatively oriented losses gives us an analogous property for loss distributions. Letting  $L$  represent a random loss as usual, the *OCE risk* associated with  $L$  under  $\phi$  is defined by

$$\text{OCE}_\mu(L; \phi) := \inf_{\theta \in \mathbb{R}} [\theta + \mathbf{E}_\mu \phi(L - \theta)] \tag{33}$$

where  $\phi : \mathbb{R} \rightarrow \mathbb{R}$  is assumed to have the same properties as  $f$  in (32), except that  $\phi$  is convex rather than concave, but still non-decreasing. We can easily link the OCE risk of (33) with the OCE for payouts given by (32); given a utility function  $f$ , a “disutility” function is naturally created by taking  $\phi$  such that  $\phi(x) = -f(-x)$ . Note that there are two minus signs here. First,  $x$  passed to  $\phi$  will be a loss, and thus since we need to pass  $f$  a payout, we pass it  $-x$ . Second, since  $f(\cdot)$  returns utility (larger is better),  $-f(\cdot)$  is something to be minimized. This setting yields

$$\begin{aligned} \text{OCE}_\mu(L; \phi) &= \inf_{\theta \in \mathbb{R}} [\theta - \mathbf{E}_\mu f(-L + \theta)] \\ &= -\sup_{\theta \in \mathbb{R}} (-1) [\theta - \mathbf{E}_\mu f(-L + \theta)] \\ &= -\sup_{\theta \in \mathbb{R}} [-\theta + \mathbf{E}_\mu f(-L + \theta)] \\ &= -\sup_{\theta \in \mathbb{R}} [\theta + \mathbf{E}_\mu f(-L - \theta)] \\ &= -\overline{\text{OCE}}_\mu(-L; f) \end{aligned}$$

recalling that  $\text{OCE}_\mu(\cdot)$  on the left-most side follows the definition (33), and  $\overline{\text{OCE}}_\mu(\cdot)$  on the right-most side follows (32). Regardless of which form we use, the notion of a *trade-off* between certain and uncertain outcomes is quite fundamental in the theory of risk functions.

16. See Ben-Tal and Teboulle (2007, Defn. 2.1) for example.

17. Still, the naming is peculiar; one would expect the  $\arg \max_{\theta \in \mathbb{R}}$  rather than the  $\sup_{\theta \in \mathbb{R}}$  to be called OCE.

If we overload the notation in (33) setting  $\underline{\text{OCE}}_\mu(h; \phi) := \underline{\text{OCE}}_\mu(\mathbf{L}(h); \phi)$  for each  $h \in \mathcal{H}$ , this becomes another natural criterion for machine learning problems. One very important quality of this class is the ease of preserving *convexity*. If  $h \mapsto \mathbf{L}(h) = \ell(h; \mathbf{Z})$  is convex on  $\mathcal{H}$  with probability one, then  $h \mapsto \underline{\text{OCE}}_\mu(h; \phi)$  is also convex on  $\mathcal{H}$ . Preserving convexity depends critically on the monotonicity of  $\phi$ ; the convexity of  $\phi$  alone is not sufficient.<sup>18</sup> The OCE risk class as a whole was tackled using traditional statistical learning theory techniques and assumptions in an elegant way by Lee et al. (2020). Following their notation, by setting  $\varphi(x) := \phi(x) - x$ , we can re-write any OCE risk as

$$\underline{\text{OCE}}_\mu(\mathbf{L}; \phi) = \mathbf{E}_\mu \mathbf{L} + \inf_{\theta \in \mathbb{R}} [\mathbf{E}_\mu \varphi(\mathbf{L} - \theta)] \tag{34}$$

which gives this learning criterion a natural interpretation as a “mean + deviation” sum.<sup>19</sup> With the assumptions placed on  $\phi$  (and thus  $\varphi$ ), it follows that  $\varphi(x) \geq 0$  for all  $x \in \mathbb{R}$ , and thus (34) shows us that  $\mathbf{E}_\mu \mathbf{L} \leq \underline{\text{OCE}}_\mu(\mathbf{L}; \phi)$  always holds. As we will see in the examples to follow, OCE risks include important sub-classes that are well-known and well-studied in their own right.

*Example 9* (Tilted risk). For  $\gamma \neq 0$ , let us define

$$C_\mu(h; \gamma) := \frac{1}{\gamma} \log \left( \mathbf{E}_\mu e^{\gamma \mathbf{L}(h)} \right) \tag{35}$$

and call this the  $\gamma$ -tilted risk, a term used recently in work by Li et al. (2020, 2021a). First and foremost, when  $\gamma > 0$ , this is a special case of OCE risk (33), with  $\phi_\gamma(x) := (e^{\gamma x} - 1)/\gamma$ . It also has a rather special property in that it can be seen as a certainty equivalent under  $\phi_\gamma$ , which is to say that for any choice of  $\gamma \neq 0$ , one can easily verify that

$$\phi(C_\mu(h; \gamma)) = \mathbf{E}_\mu \phi_\gamma(\mathbf{L}(h)). \tag{36}$$

It must be noted however that when  $\gamma < 0$ , this  $C_\mu(\cdot; \gamma)$  is not a valid OCE risk, since  $\phi_\gamma(\cdot)$  is concave, and plugging  $\phi = \phi_\gamma$  into (33) clearly leads to  $\underline{\text{OCE}}_\mu(\mathbf{L}; \phi_\gamma) = -\infty$ . That said, while the OCE form is meaningless when  $\gamma < 0$ , the tilted risk form (35) is perfectly valid, and for general choices of  $\gamma \neq 0$ , this objective has a long history of use in risk-sensitive reinforcement learning; see García and Fernández (2015, §3.2.1). As described by Föllmer and Knispel (2011), the special case of  $\gamma > 0$  is often referred to as the *entropic risk* in financial contexts. Viewed as a function of  $\gamma$ , the quantity  $\gamma C_\mu(h; \gamma)$  is the logarithm of the *moment generating function* of  $\mathbf{L}(h)$ , a quantity that plays an important role in statistical theory (Stuart and Ord, 1994, Ch. 3). Replacing  $\mu$  in (35) with the empirical distribution on  $\mathbf{Z}_n = (Z_1, \dots, Z_n)$ , we get

$$\widehat{C}_n(h; \gamma) := \frac{1}{\gamma} \log \left( \frac{1}{n} \sum_{i=1}^n e^{\gamma \mathbf{L}_i(h)} \right) \tag{37}$$

and taking  $\gamma > 0$  sufficiently large gives us a smooth approximation to the “largest loss” objective  $\mathbf{L}_{(n)} = \mathbf{L}_{[1]}$ , a well-known technique from the optimization literature (Pee and Royset, 2011). ■

18. See the “Convexity Theorem” of Rockafellar and Uryasev (2013).  
 19. Using terminology from the fundamental risk quadrangle of Rockafellar and Uryasev (2013),  $\mathbf{E}_\mu \phi$  in (33) is a “regret measure,” and  $\mathbf{E}_\mu \varphi$  in (34) is an “error measure” quantifying the “nonzeroness” of  $\mathbf{L}$ .

*Example 10* (CVaR). A very important special case of OCE risk is CVaR, given earlier in the form (20) in §2.3. One of the central results in the seminal work of Rockafellar and Uryasev (2000, 2002) is that CVaR as given by (20) can be represented using a convex program, namely that

$$C_{\mu}^{\text{CVaR}}(h; \beta) = \inf_{\theta \in \mathbb{R}} \left[ \theta + \frac{1}{1-\beta} \mathbf{E}_{\mu}(\mathbf{L}(h) - \theta)_+ \right] = Q_{\beta}(h) + \frac{1}{1-\beta} \mathbf{E}_{\mu}(\mathbf{L}(h) - Q_{\beta}(h))_+ \quad (38)$$

holds for any  $0 \leq \beta < 1$ , where  $(x)_+ := \max\{0, x\}$  denotes the positive part. In particular, this tells us that CVaR is an OCE risk with  $\phi$  set to  $\phi(x) = (x)_+/(1-\beta)$ . In the special case of binary classification using margin errors  $\ell(h; x, y) = -y\langle h, x \rangle / \|h\|$  with  $y \in \{-1, +1\}$ , Takeda and Sugiyama (2008) showed that minimizing the CVaR objective (38) is equivalent to solving the extended  $\nu$ -SVM objective. On the other hand, when the underlying training losses  $L_1(h), \dots, L_n(h)$  are Bernoulli random variables (e.g., 0-1 error), Zhai et al. (2021b) show how the CVaR objective collapses into the expected 0-1 error objective, though randomizing the choice of  $h$  can be used to obtain distinct performance guarantees. ■

*Remark 11* (Inverted OCE). While traditional OCE risks are always at least as large as the expected loss, giving us flexibility over sensitivity to tails in the upward direction, one naturally may be interested in prioritizing tails in the *downward* direction, i.e., placing more emphasis on “easy” examples and ignoring hard examples. Obviously this can be realized directly using order-based risks as discussed in §3.1, but a slight modification to the OCE risk class can also be used to achieve such an effect. Letting  $\phi$  be as before, define the *inverted OCE risk* by

$$\underline{\text{IOCE}}_{\mu}(\mathbf{L}; \phi) := \sup_{\theta \in \mathbb{R}} [\theta - \mathbf{E}_{\mu} \phi(\theta - \mathbf{L})] \quad (39)$$

$$= - \inf_{\theta \in \mathbb{R}} [\mathbf{E}_{\mu} \phi(\theta - \mathbf{L}) - \theta] \quad (40)$$

$$= \mathbf{E}_{\mu} \mathbf{L} - \inf_{\theta \in \mathbb{R}} \mathbf{E}_{\mu} \varphi(\theta - \mathbf{L}) \quad (41)$$

where  $\varphi(x) = \phi(x) - x$  as before, meaning that since  $\varphi(x) \geq 0$  for all  $x \in \mathbb{R}$ , the inverted OCE can be no greater than the mean, i.e.,  $\underline{\text{IOCE}}_{\mu}(\mathbf{L}; \phi) \leq \mathbf{E}_{\mu} \mathbf{L}$  holds for all valid  $\phi$ . The naming is due to Lee et al. (2020), who also show that the special case of “inverted CVaR” under the empirical distribution yields an “average bottom- $k$ ” risk (Lee et al., 2020, Prop. 1), a fact analogous to the link seen earlier between traditional CVaR and average top- $k$  risk (cf. §2.3 and Ex. 8). In particular, this shows us that inverted OCE can be used to completely ignore outliers in the upward direction; of course, when  $\mathbf{L}$  is unbounded below, sensitivity to outliers in the downward direction is amplified. Similarly, recalling the special case of the tilted/entropic risk introduced earlier with  $\phi_{\gamma}(x) = (e^{\gamma x} - 1)/\gamma$  for  $\gamma > 0$ , plugging this in to the inverted OCE (39), a bit of calculus quickly shows us that

$$\underline{\text{IOCE}}_{\mu}(\mathbf{L}; \phi_{\gamma}) = \frac{1}{-\gamma} \log \left( \mathbf{E}_{\mu} e^{-\gamma \mathbf{L}} \right)$$

meaning that the tilted risk (35) using a *negative* tilting parameter is an inverted OCE risk, as we might intuitively expect. Note that in general, the inverted OCE does not inherit the convexity properties enjoyed by the original OCE. ■

*Remark 12* (Dispersion and OCE). In light of the alternate form (34) for OCE risk, one natural entry point when designing new learning criteria is to start with  $\varphi$  used to measure “dispersion” in the underlying distribution. The canonical measure of dispersion is the standard deviation, and noting that

$$\inf_{\theta \in \mathbb{R}} \mathbf{E}_\mu(\mathbf{L} - \theta)^2 = \mathbf{E}_\mu(\mathbf{L} - \mathbf{E}_\mu \mathbf{L})^2 = \text{var}_\mu \mathbf{L}$$

it seems intuitive to try setting  $\varphi(x) = x^2/2$ . In this case,  $\phi(x) = \varphi(x) + x$  indeed satisfies  $1 \in \partial\phi(0)$  and  $\phi(0) = 0$ , but a lack of monotonicity means this choice, namely the mean-variance  $\mathbf{E}_\mu \mathbf{L} + \text{var}_\mu \mathbf{L}$ , is not a valid OCE risk.<sup>20</sup> Naively replacing this with  $\phi(x) = (x)_+^2/2 + x$  does yield a valid OCE risk, but unfortunately it is not very interesting, since

$$\inf_{\theta \in \mathbb{R}} \mathbf{E}_\mu(\mathbf{L} - \theta)_+^2 = 0$$

and thus the OCE risk collapses into the mean, i.e.,  $\text{OCE}_\mu(\mathbf{L}; \phi) = \mathbf{E}_\mu \mathbf{L}$ . This also means that criteria based on asymmetric measures of dispersion relative to a fixed threshold  $\theta$ , such as

$$\mathbf{E}_\mu \mathbf{L} + \left(\frac{1}{p}\right) \mathbf{E}_\mu (\mathbf{L} - \theta)_+^p \tag{42}$$

and

$$\mathbf{E}_\mu \mathbf{L} + \left(\mathbf{E}_\mu (\mathbf{L} - \theta)_+^p\right)^{1/p} \tag{43}$$

cannot be OCE risks for any value of  $p \geq 1$  and  $\theta \in \mathbb{R}$  such that  $\mathbf{E}_\mu (\mathbf{L} - \theta)_+^p > 0$ . In the special case where we set  $\theta = \mathbf{E}_\mu \mathbf{L}$ , the criterion (43) is known as *mean-upper-semi-deviation*, and is well-studied in the literature, but not an OCE risk; see the examples in Ruszczyński and Shapiro (2006b); Rockafellar and Uryasev (2013) for more background. ■

### 3.3 DRO Risks

To motivate the next class of learning criteria, let us re-consider “worst case scenarios” in a learning problem. Given a fixed data distribution  $\mathbf{Z} \sim \mu$  and random test loss  $\mathbf{L} \in \mathcal{L}_\mu(\mathcal{H})$  with  $\mathcal{L}_\mu(\mathcal{H})$  as defined in (5) as usual, in the preceding §3.2 we saw OCE risks such as entropic risk and CVaR which let us control the degree of sensitivity to the loss tails on the upside, i.e., a bad event stated in terms of the random draw of  $\mathbf{L}$ . This scenario is perfectly natural, but recalling the key assumption (2) it usually assumes that we have training data  $\mathbf{Z}_n = (\mathbf{Z}_1, \dots, \mathbf{Z}_n)$  which are independently sampled from  $\mu$ , and does not consider the possibility that the data distribution at training time might differ from that at test time; a case of *distribution shift*. In order to make a decision with some degree of resilience to a shift in the data distribution, letting  $\mathcal{P}$  denote a set of probability distributions defined on the same domain as  $\mu$ , one can formalize a new “worst-case” scenario using

$$\bar{C}_{\mathcal{P}}(h) := \sup \{C_\nu(h) : \nu \in \mathcal{P}\} \tag{44}$$

---

20. Table 1 of Lee et al. (2020) incorrectly includes the mean-variance as an OCE risk.

where  $L \mapsto C_\nu(L)$  is an arbitrary learning criterion defined on  $\mathcal{L}_\nu(\mathcal{H})$  for each  $\nu \in \mathcal{P}$ , and we overload our notation with  $C_\nu(h) := C_\nu(L(h))$  as usual. In the special case where  $C_\nu L = \mathbf{E}_\nu L$ , the objective (44) and algorithms designed to minimize it are categorized under the term *distributionally robust optimization (DRO)*, with (44) referred to as a *DRO risk*. See Shapiro (2017) and the references therein for more background from the optimization community. In the context of machine learning, the work of Duchi and Namkoong (2021) has been particularly influential.

It goes without saying that the critical element of (44) is the set  $\mathcal{P}$ , often referred to as the *uncertainty set* (Duchi and Namkoong, 2021; Shapiro, 2017). Still assuming the training data are sampled from  $\mu$ , there must be some connection between  $\mu$  and the elements of  $\mathcal{P}$  in order for the learning problem to be tractable. Even when  $\mu \in \mathcal{P}$ , if all the distributions in  $\mathcal{P} \setminus \{\mu\}$  are completely unrelated to  $\mu$ , then we have no hope of minimizing  $\bar{C}_\mathcal{P}(h)$  in (44), even approximately. We only have access to a sample  $\mathbf{Z}_n$  from  $\mu$ , and so the other elements in  $\mathcal{P}$  must be “similar” or “close” enough to  $\mu$  that a straightforward perturbation of the sample  $\mathbf{Z}_n$  can effectively simulate a sample taken from the most “unlucky” distribution included in  $\mathcal{P}$ .

A natural way of formalizing this notion is to introduce a measure of discrepancy between probability distributions. Letting  $\text{Div}(\cdot)$  be a function which takes a pair of probability measures and returns a non-negative value from the extended real line, given the special role of  $\mu$  here, it is intuitive to constrain  $\mathcal{P}$  in (44) to contain only distributions which are  $\varepsilon$ -similar to  $\mu$ , where discrepancy is measured by  $\text{Div}(\cdot)$ . More explicitly, we set  $\mathcal{P} = \mathcal{P}_\varepsilon(\mu; \text{Div})$  defined by

$$\mathcal{P}_\varepsilon(\mu; \text{Div}) := \{\nu : \text{Div}(\nu; \mu) \leq \varepsilon\}. \tag{45}$$

If in addition to  $\text{Div}(\nu; \mu) \geq 0$ , we have  $\text{Div}(\nu; \mu) = 0$  if and only if  $\nu = \mu$ , it is traditional to call  $\text{Div}(\cdot)$  a *divergence*.<sup>21</sup> The best-known special case of a divergence is the *Kullback-Leibler (KL) divergence* of  $\nu$  from  $\mu$ , defined by

$$\text{Div}_{\text{KL}}(\nu; \mu) := \mathbf{E}_\nu \log \left( \frac{d\nu}{d\mu} \right) \tag{46}$$

whenever  $\nu \ll \mu$  (absolute continuity), with  $\text{Div}_{\text{KL}}(\nu; \mu) := \infty$  otherwise. Here  $d\nu/d\mu$  denotes the Radon-Nikodym density of  $\nu$  with respect to  $\mu$ , a random variable.<sup>22</sup> Note that  $\text{Div}_{\text{KL}}$  in (46) can be re-written as

$$\text{Div}_{\text{KL}}(\nu; \mu) = \mathbf{E}_\mu \left( \frac{d\nu}{d\mu} \right) \log \left( \frac{d\nu}{d\mu} \right) = \mathbf{E}_\mu \left( \frac{d\nu}{d\mu} \right) \log \left( \frac{d\nu}{d\mu} \right) - \mathbf{E}_\mu \left( \frac{d\nu}{d\mu} \right) + 1. \tag{47}$$

In the right-most expression of (47), the density  $d\nu/d\mu$  is being passed through the function  $x \mapsto x \log(x) - x + 1$  before taking expectation. This function is convex and non-negative on  $[0, \infty)$ , taking its minimum (zero) at  $x = 1$ . An important generalization of the KL divergence simply replaces this map with an arbitrary function  $x \mapsto f(x)$  that retains the

21. There is a tremendous literature on divergences and their role in the problems of statistics, machine learning, and information theory. See Grünwald and Dawid (2004) for a stimulating entry point.

22. For a textbook introduction to Radon-Nikodym densities, see Ash and Doléans-Dade (2000, §2.2).

key properties just mentioned. More formally, given a convex function  $f : [0, \infty) \rightarrow [0, \infty]$  satisfying  $f(1) = 0$ , we define the  $f$ -divergence of  $\nu$  from  $\mu$  by

$$\text{Div}_f(\nu; \mu) := \mathbf{E}_\mu f\left(\frac{d\nu}{d\mu}\right) \quad (48)$$

whenever  $\nu \ll \mu$ , otherwise setting  $\text{Div}_f(\nu; \mu) := \infty$  as before. While in principle all sorts of different uncertainty sets can be conceived of, typical DRO risks in the machine learning literature are all special cases of  $\mathcal{P}_\varepsilon(\mu; \text{Div}_f)$ . See recent work by Fröhlich and Williamson (2023), who demonstrate sharp connections between the design of  $\mathcal{P}_\varepsilon(\mu; \text{Div}_f)$  and distributions with “desirable” tail properties. Representative special cases of DRO risk will be introduced in the examples to follow.

*Example 13* (KL, DRO, and tilted risk). A natural first example is  $f(x) = x \log(x) - x + 1$ , so that  $\text{Div}_f = \text{Div}_{\text{KL}}$  as given above. In this case, assuming that  $C_\nu(h) = \mathbf{E}_\nu L(h)$  and the absolute loss  $|L(h)|$  is bounded, when we set  $\mathcal{P} = \mathcal{P}_\varepsilon(\mu; \text{Div}_{\text{KL}})$  for  $\varepsilon > 0$ , the DRO risk (44) can be expressed as

$$\bar{C}_\mathcal{P}(h) = \inf_{\gamma > 0} \left[ \frac{\varepsilon}{\gamma} + \frac{1}{\gamma} \log\left(\mathbf{E}_\mu e^{\gamma L(h)}\right) \right]. \quad (49)$$

See Shapiro (2017, Ex. 3.11) and the references therein for more details on this derivation. The equality (49) illustrates a close relation between the tilted risk under  $\mu$  (cf. Example 9) and the worst-case *mean* over all  $\nu$  which are at least  $\varepsilon$ -similar to  $\mu$  in terms of KL divergence. Taking  $\varepsilon > 0$  large means we need to take  $\gamma > 0$  correspondingly large in order to minimize the objective in (49). Taking  $\varepsilon \rightarrow 0_+$  implies that the optimal tilting parameter correspondingly decreases as  $\gamma \rightarrow 0_+$ . In this limit we have  $\bar{C}_\mathcal{P}(h) \rightarrow \mathbf{E}_\mu L(h)$ . Note however that minimizing (49) in  $h$  is not in general equivalent to minimizing the  $\gamma$ -tilted risk (35) for any fixed choice of  $\gamma$ , since the optimal  $\gamma$  in (49) changes with the choice of  $h$ . ■

*Example 14* (Cressie-Read DRO). A flexible and useful sub-class of  $f$ -divergences is given by setting  $f = f_c$  for functions of the form

$$f_c(x) := \frac{x^c - cx + c - 1}{c(c - 1)} \quad (50)$$

parameterized by any  $c \in \mathbb{R} \setminus \{0, 1\}$ . In the limit of  $c \rightarrow 1$ , we have  $f_c(x) \rightarrow x \log(x) - x + 1$  and thus we recover KL divergence as a special case. This class is called the *Cressie-Read family* of  $f$ -divergences by Duchi and Namkoong (2021), and restricting ourselves to  $c > 1$ , setting  $\mathcal{P} = \mathcal{P}_\varepsilon(\mu; \text{Div}_{f_c})$  we have

$$\bar{C}_\mathcal{P}(h) = \inf_{\theta \in \mathbb{R}} \left[ \theta + (1 + c(c - 1)\varepsilon)^{1/c} (\mathbf{E}_\mu (L(h) - \theta)_+^{c_*})^{1/c_*} \right] \quad (51)$$

for any choice of  $\varepsilon > 0$  (Duchi and Namkoong, 2021, Lem. 1), where  $c_* := c/(c - 1)$ . The expression in (51) is quite reminiscent of the OCE risk classes seen earlier, and indeed the basic effect is very similar. Setting  $\phi_{c,\varepsilon}(x) := (1 + c(c - 1)\varepsilon)^{c_*/c} (x)_+^{c_*}$ , we have

$$\bar{C}_\mathcal{P}(h) = \inf_{\theta \in \mathbb{R}} \left[ \theta + (\mathbf{E}_\mu \phi_{c,\varepsilon}(L(h) - \theta))^{1/c_*} \right] \quad (52)$$

with  $\phi_{c,\varepsilon}$  satisfying the key monotonicity and convexity properties required when making OCE risks; the only difference is the  $(\cdot)^{1/c_*}$  wrapping the expected value. In the end, both Cressie-Read DRO risks and OCE risks give us control over how much priority to put on “difficult” examples, where higher difficulty corresponds to larger losses incurred by a given  $h \in \mathcal{H}$ . Both classes are also risk-averse in the sense that they always upper-bound the expected loss  $\mathbf{E}_\mu \mathbf{L}(h)$ . High tail sensitivity means that empirical estimates will have to deal with outliers, making efficient estimation and learning under finite samples a challenge; as a natural compromise, Zhai et al. (2021a) propose a variant of DRO risks which emphasizes values of sufficient difficulty, but ignores everything beyond that threshold. ■

*Example 15* (Variance-based regularization). DRO risks received attention early on in the context of reducing both the expected loss and its variance in order to obtain faster convergence rate guarantees. Key initial works are the pre-prints (Duchi et al., 2016; Duchi and Namkoong, 2016) and the conference paper of Namkoong and Duchi (2016), later distilled into the journal papers (Duchi and Namkoong, 2019; Duchi et al., 2021). Of particular importance is the fact the (empirical) minimizer of DRO risk (44) using an appropriate choice of  $f$ -divergence (e.g.,  $\chi^2$ -divergence) can be shown to enjoy high-probability bounds on the expected loss taking the form

$$\inf_{h \in \mathcal{H}} \left[ \mathbf{E}_\mu \mathbf{L}(h) + \mathcal{O} \left( \sqrt{\frac{\text{var}_\mu \mathbf{L}(h)}{n}} \right) \right] + \mathcal{O} \left( \frac{1}{n} \right)$$

where  $\mathcal{O}(\cdot)$  hides constants and factors depending on the confidence level for simplicity; see Duchi and Namkoong (2019) and the references therein for more details. ■

*Example 16* (Sub-population awareness). A well-explored application of DRO risks is the design of learning algorithms which have comparable performance across certain “sub-populations.” Formally, if  $Z \sim \mu$  is our random data, we can formulate  $k$  sub-populations as sets  $Z_1, \dots, Z_k$  that  $Z$  may fall into. For some intuitive examples, these sets could represent data belonging to distinct ethnic, age, or gender categories, where  $k$  is the number of categories. Recalling §2.4 and the idea of fairness-aware learning, a natural goal is to reduce the gap in performance across sub-populations that tends to arise under naive optimization of the expected loss; this is called “representation disparity” by Hashimoto et al. (2018). One direct way of doing this is to consider a worst-case objective over sub-populations, such as

$$\max_{j \in [k]} \mathbf{E}_\mu [\mathbf{L}(h) \mid Z \in Z_j]. \tag{53}$$

The objective in (53) can be bounded using DRO risk in a straightforward manner. Letting  $\nu_j$  denote the conditional distribution of  $Z \mid Z \in Z_j$  for each  $j \in [k]$ , clearly for any  $Z$  such that  $\mu\{Z \in Z\} = 0$ , we have  $\nu_j\{Z \in Z\} = 0$ , meaning  $\nu_j \ll \mu$  holds, and we can find a constant  $0 < r_{\max} < \infty$  such that

$$\max_{j \in [k]} \frac{d\nu_j}{d\mu} \leq r_{\max} \tag{54}$$

with probability 1. It thus follows that

$$\text{Div}_f(\nu_j; \mu) = \mathbf{E}_\mu f \left( \frac{d\nu_j}{d\mu} \right) \leq f(r_{\max}) \tag{55}$$

and thus setting  $\varepsilon_{\max} := f(r_{\max})$ , we have  $\nu_j \in \mathcal{P}_{\varepsilon_{\max}}(\mu; \text{Div}_f)$  for each  $j \in [k]$ . We may thus conclude that

$$\max_{j \in [k]} \mathbf{E}_{\mu} [\mathbf{L}(h) | Z \in Z_j] = \max_{j \in [k]} \mathbf{E}_{\nu_j} \mathbf{L}(h) \leq \bar{C}_{\mathcal{P}}(h) \tag{56}$$

with  $\mathcal{P} = \mathcal{P}_{\varepsilon_{\max}}(\mu; \text{Div}_f)$  as above. This useful and general fact is stated concisely by Zhai et al. (2021a, Prop. 1), though it appears earlier for the special case of  $f = f_c$  following (50) with  $c = 2$  (known as  $\chi^2$ -divergence) in Hashimoto et al. (2018, Prop. 2). This fact is useful since it suggests that minimizing the DRO risk may be a fruitful strategy even in the “oblivious” setting where the sub-populations are unknown. ■

### 3.4 Miscellaneous

In the preceding §3.1–§3.3 we covered the three best-known classes of learning criteria in the machine learning literature. In this section, we briefly introduce some other classes which have appeared more recently in the literature.

#### 3.4.1 CRITERIA BASED ON BIDIRECTIONAL DISPERSION

In our description earlier of OCE and DRO risks (with uncertainty set as in (45)), we saw how these risks can be used to emphasize the tails in the upward direction, meaning that for any criterion  $C_{\mu}$  from either of these classes, we have

$$\mathbf{E}_{\mu} \mathbf{L} \leq C_{\mu}(\mathbf{L}), \quad \forall \mathbf{L} \in \mathcal{L}_{\mu}(\mathcal{H}) \tag{57}$$

regardless of the properties of  $\mu$  and  $\mathcal{H}$ . We also saw an “inverted” version of OCE risks (cf. Remark 11) which can be used to emphasize the tails on the downside, resulting in

$$\mathbf{E}_{\mu} \mathbf{L} \geq C_{\mu}(\mathbf{L}), \quad \forall \mathbf{L} \in \mathcal{L}_{\mu}(\mathcal{H}). \tag{58}$$

In any case, all OCE risks (inverted variants included) and DRO risks are “uni-directional” in the sense that either (57) or (58) holds; these criteria can only be “pulled” by heavy tails in one direction, never two. More precisely, we cannot have two distinct distributions  $\mathbf{L}, \mathbf{L}' \in \mathcal{L}_{\mu}(\mathcal{H})$  where  $\mathbf{E}_{\mu} \mathbf{L} \geq C_{\mu}(\mathbf{L})$  and  $\mathbf{E}_{\mu} \mathbf{L}' < C_{\mu}(\mathbf{L}')$  hold simultaneously. Of course, certain order-based criteria readily satisfy this. Considering the median for example, if  $Q_{0.5}(\mathbf{L}) > \mathbf{E}_{\mu} \mathbf{L}$ , then trivially  $Q_{0.5}(-\mathbf{L}) < \mathbf{E}_{\mu}(-\mathbf{L})$ . On the other hand, as mentioned in §2.1, while quantiles, like any M-parameter, are easy to compute, they are difficult to optimize as a function of the loss distribution. A potential compromise can be achieved using learning criteria of the general form

$$C_{\mu}(\mathbf{L}; \theta, \eta, \rho) := \eta\theta + \mathbf{E}_{\mu} \rho(\mathbf{L} - \theta) \tag{59}$$

where  $\eta \in \mathbb{R}$  (negative values are allowed),  $\theta \in \mathbb{R}$ , and  $\rho : \mathbb{R} \rightarrow [0, \infty)$  is assumed to be such that the induced set of M-parameters is non-empty and satisfies  $\arg \min_{\theta \in \mathbb{R}} \mathbf{E}_{\mu} \rho(\mathbf{L} - \theta) \subset \mathbb{R}$ . The criterion  $C_{\mu}(\mathbf{L}; \theta, \eta, \rho)$  is called the *threshold risk* (or *T-risk*) by Holland (2023). A natural sub-class is the “minimal T-risk” given by

$$\underline{C}_{\mu}(\mathbf{L}; \eta, \rho) := \inf_{\theta \in \mathbb{R}} C_{\mu}(\mathbf{L}; \theta, \eta, \rho) \tag{60}$$



which requires that  $\rho$  grow without bound at a linear rate (or faster); otherwise trivially  $\underline{C}_\mu(\mathbf{L}; \eta, \rho) = -\infty$ . This is a compromise in the sense that the criterion itself never yields the M-parameters induced by  $\rho$  exactly, only a “regularized” version biased slightly above or below, depending on  $\text{sign}(\eta)$ . On the other hand, it is computationally much easier to handle than true M-parameters, with minimization of  $C_\mu(\mathbf{L}; \theta, \eta, \rho)$  in  $(\mathbf{L}, \theta)$  quite easy to implement. The minimal T-risk (60) has obvious similarities to OCE risks from §3.2, but the main difference is that  $\rho(\cdot)$  can be *non-monotonic*. This means that convexity guarantees for  $h \mapsto C_\mu(\mathbf{L}(h); \theta, \eta, \rho)$  are not in general possible, but it gives us the flexibility to set  $\eta$  to both positive and negative values for a fixed  $\rho$ , and allows for bidirectional sensitivity to loss tails that is not possible with OCE or DRO risks.<sup>23</sup> This positions the minimal T-risk in a somewhat peculiar middle-ground; it does not appear in the survey of Ruszczyński and Shapiro (2006b), nor on any corner of the “quadrangles” in the rich examples of Rockafellar and Uryasev (2013).

*Example 17* (Quantile-based thresholds). Setting  $\rho(\cdot) = |\cdot|$  and manipulating  $\eta \in (-1, +1)$ , we can set the T-risk threshold to any quantile of  $\mathbf{L}$ . More precisely, if we set  $\eta = 1 - 2\beta$  for any quantile level  $0 < \beta < 1$ , we have

$$\underline{C}_\mu(\mathbf{L}; \eta, \rho) = (1 - 2\beta) \mathbf{Q}_\beta(\mathbf{L}) + \mathbf{E}_\mu \rho|\mathbf{L} - \mathbf{Q}_\beta(\mathbf{L})|.$$

For more on properties defined using convex programs, see Koltchinskii (1997). ■

*Example 18* (Weakly convex T-risk). While the objective  $h \mapsto C_\mu(\mathbf{L}(h); \theta, \eta, \rho)$  will not in general be convex when  $\rho$  is non-monotonic, a form of weak convexity can be guaranteed when the losses are sufficiently smooth (in  $h$ ) using convex  $\rho$  with sub-quadratic growth in the limit. One useful example is

$$\rho(x) = x \operatorname{atan}(x) - \frac{\log(1 + x^2)}{2} \tag{61}$$

whose strict convexity and Lipschitz continuity are amenable to weak convexity arguments, and coupled with a re-scaling mechanism lets us flexibly interpolate the minimizing threshold between  $\mathbf{E}_\mu \mathbf{L}$  and  $\mathbf{Q}_{0.5}(\mathbf{L})$  (Holland, 2022). ■

### 3.4.2 CRITERIA INSPIRED BY HUMAN PSYCHOLOGY

A unique and interesting direction is designing learning criteria which attempt to mimic human tendencies for risk aversion (or preference) when making decisions under uncertainty. In terms of the machine learning literature, the most salient example is an application of the metrics used in the *cumulative prospect theory (CPT)* of Tversky and Kahneman (1992). Roughly speaking, “prospect” refers to subjective evaluations of uncertain outcomes. To make our explanation as clear as possible, let us denote this uncertain outcome using an arbitrary random variable  $\mathbf{X} \sim \mu$ . Let  $V(\cdot)$  be a “value” function, and let  $F_\mu$  denote the distribution function of  $\mathbf{X}$ , namely  $F_\mu(x) := \mathbf{P}\{\mathbf{X} \leq x\}$ . The *CPT score* assigned to  $\mathbf{X}$  is defined by taking the expectation using a *modified* distribution function  $\tilde{F}_\mu$ , namely

$$\text{CPT}_\mu(\mathbf{X}) := \int_{\mathbb{R}} V(x) \tilde{F}_\mu(dx). \tag{62}$$

23. Early work on this class for a rather specialized choice of  $\rho$  dates back to a 2020 pre-print by Holland (2020), later expanded by Holland (2022). See Holland (2023) for more general results.

Thus, the CPT score is characterized by the value function  $V$  and the transformation  $F_\mu \mapsto \tilde{F}_\mu$ . Setting  $c$  as a threshold (often  $c = 0$ ), using indicators  $I\{X \geq c\}$  and  $I\{X < c\}$  the CPT score typically treats these two mutually exclusive events differently. To do this, we construct value functions  $V_+$  and  $V_-$  and differentiable weighting functions  $w_+$  and  $w_-$  for each of these cases, setting

$$\begin{aligned} V(x) &= I\{x \geq c\}V_+(x) + I\{x < c\}V_-(x) \\ \tilde{F}(x) &= I\{x \geq c\}w_+(F(x)) + I\{x < c\}w_-(F(x)). \end{aligned}$$

Putting these pieces together, the resulting CPT score takes the more explicit form

$$\text{CPT}_\mu(\mathbf{X}) = \int_{-\infty}^c V_-(x) w'_-(F(x)) F(dx) + \int_c^\infty V_+(x) w'_+(F(x)) F(dx) \quad (63)$$

where  $w'_+$  and  $w'_-$  are the derivatives of the weight functions. The threshold of  $c = 0$  is often used, but this is of course arbitrary; it can be set to any meaningful value given the context. The important concepts here are as follows: (a) “gains” and “losses” relative to some threshold are treated differently in human psychology; (b) small probabilities tend to be underweighted, large probabilities tend to be overweighted. See Rieger and Wang (2006) for a useful background reference. Taking this from CPT scores to a natural notion of *CPT risk* just amounts to setting  $\mathbf{X} = L(h)$ , with our usual shorthand notation  $\text{CPT}(h) := \text{CPT}(L(h))$  for each  $h \in \mathcal{H}$ . While the general-purpose CPT score itself is well-known, literature on the CPT risk in machine learning remains quite sparse. Concentration bounds for estimators of the CPT risk (pointwise in  $h$ ) are given by Bhat and Prashanth (2019). To the best of our knowledge, it appears that Liu et al. (2019) are the first to consider learning with CPT risks, under the catchy name “human-aligned risks,” although their studies are limited to an empirical analysis of learning using ERM implementations of certain special sub-classes of CPT risks. Under some assumptions, statistical learning guarantees are possible for specialized sub-classes of CPT risk using the machinery of Liu et al. (2022).

*Example 19* (Links to order-based risks). Recalling our formulation of order-based risks in §3.1, there is clear conceptual overlap with the CPT risk just introduced. Using a weight function

$$W_{\text{CPT}}(u) := I\{u \geq F_\mu(c)\}w'_+(u) + I\{u < F_\mu(c)\}w'_-(u)$$

we see that the criterion  $C_\mu(L; W)$  in (28) can be used to represent a particular CPT risk, i.e., leaving the value function as the identity  $V(x) = x$ , setting  $W = W_{\text{CPT}}$  we have

$$\text{CPT}_\mu(L) = C_\mu(L; W_{\text{CPT}}).$$

The general value function is not present in our formulation in §3.1, but it should be noted that arbitrary (value) functions of order statistics are reflected in the textbook formulation of L-estimates given by Huber and Ronchetti (2009, Ch. 3). ■

## 4. Historical Background

To complement the preceding survey of criteria used in machine learning, we take a deeper look at the historical context within which the “classical paradigm” described in §1.1 developed.

#### 4.1 Origins of Expected Loss Minimization

The problem of *decision-making under uncertainty* has interested scholars for centuries. For an uncertain outcome which can be expressed numerically, the traditional “mathematical expectation” is the sum of all possible outcomes, weighted by the probability of each outcome.<sup>24</sup> This notion was clearly described in letters from the Bernoullis and Gabriel Cramer dating back to the early 18th century (Bassett, 1987). With the expected value defined as a probability-weighted sum, then a problem of fundamental importance is *how to evaluate outcomes*, namely the process of assigning numerical values to real-world outcomes, which are appropriate within the context of the decisions that need to be made. The importance of this is highlighted by the “St. Petersburg paradox,” dating back to Nikolaus and Daniel Bernoulli (Menger, 1967). Consider the following game: a fair coin is tossed until it turns up heads, and one receives a payment of  $2^K$  dollars (or ducats), where  $K$  is a random variable representing the number of tosses it took for heads to turn up. Note that  $K$  is bounded below by 1, and unbounded above. The expected payout is  $\mathbf{E}[2^K] = 2(1/2) + 2^2(1/2^2) + \dots = \infty$ . While the expected payout is infinite, it seems irrational to pay a large amount to play this game, and thus arises the need to refine how we evaluate outcomes when making decisions. Essentially, this amounts to transforming the raw payouts in such a way that better reflects human preferences, say by introducing a function  $f$  such that  $\mathbf{E}[f(2^K)] < \infty$ . More generally, given any uncertain outcome  $X$ , the transformed value  $f(X)$  is often called the *utility* of that outcome.<sup>25</sup> The basic decision-making principle derived from this line of thought is simple: act in a way that *maximizes the expected utility*.<sup>26</sup>

Over the course of the 20th century, the theoretical foundations of this principle were substantially strengthened. Of particular historical significance is the seminal theory of games developed by von Neumann and Morgenstern (1953), where expected utility  $\mathbf{E}[f(X)]$  appears in the context of axioms related to preferences over uncertain outcomes (or “lotteries”) that we expect any “rational” agent to satisfy. These axioms are typically stated with respect to a class of uncertain outcomes, here a set  $\mathcal{X}$  of random variables. The famous “expected utility theorem” states that any agent satisfying these axioms (under  $\mathcal{X}$ ) *must* assign its preferences based on expected utility, that is, there exists a utility function  $f$  such that for any  $X_1, X_2 \in \mathcal{X}$ , we have  $X_1 \preceq X_2$  (i.e., “ $X_2$  is no worse than  $X_1$ ”) if and only if  $\mathbf{E}[f(X_1)] \leq \mathbf{E}[f(X_2)]$ .<sup>27</sup> As such, when the agent can choose between uncertain outcomes,

24. Historically, a broader notion of “expectation” can in many ways be seen as a more basic (and intuitive) concept than that of probability. Expectations are central to fair pricing in games of chance described by Christian Huygens in 1657, see Hacking (2006, Ch. 11) for detailed background, and Pollard (2002, Ch. 1) for similar ideas in the context of measure theoretic foundations.

25. The expected utility was historically referred to as the “moral expectation” (in contrast to the mathematical expectation) by Daniel Bernoulli, who suggested using the utility function  $f(x) = \log(x)$  as a solution to the St. Petersburg problem (Jordan, 1924). From our modern viewpoint, it is clear that this only solves one special case of the problem, since given any  $f$  which is unbounded above, one can construct an analogous payout scheme such that the expected utility is infinite. The first work known to have pointed out this problem was a 1934 work in German by Karl Menger (Menger, 1967); see Bassett (1987) for an interesting discussion on the history of this issue.

26. For more history on expected utility and its role in the theory of the social sciences (especially economics), see the works of Jordan (1924), Uzawa (1956), Samuelson (1977), and Bassett (1987).

27. For a modern proof of this fact for discrete random variables, see Mas-Colell et al. (1995, Props. 6.B.3, 6.E.1). Regarding the terminology used by von Neumann and Morgenstern, they use the term “utility” to describe the real-world concept to which we attach preferences, and “numerical utility” to refer to

it will choose the one with the largest expected utility. Due to the significant influence of this work, the decision-making principle of maximizing the expected utility is essentially synonymous with the “rational agent” model pervasive in the social sciences (Tversky and Kahneman, 1992; Mas-Colell et al., 1995).

Since the usual form of expected utility is a probability-weighted sum, one naturally runs into the question of where these probabilities come from, and what they mean.<sup>28</sup> When betting on the outcome of a coin toss, under most conditions it is reasonable to say that the equal probability of heads and tails is a fact of nature, and the probabilities  $\mathbf{P}\{X = \text{“heads”}\} = \mathbf{P}\{X = \text{“tails”}\} = 1/2$  are in this sense *objectively* determined. On the other hand, consider the question of “will the stock price of Company ABCD rise next quarter?” This is a typical example of uncertainty in the real world, and while the possible outcomes are crystal-clear (the price will either rise, or it will not), it is very difficult to conceive of any kind of objective probability that can be assigned to such an uncertain outcome. Depending on knowledge and past experience (among countless other factors), the probability that different parties assign to this event will assuredly differ, making probability an inherently *subjective* notion. Indeed, if we polled a handful of stock traders, employees of Company ABCD, and individuals with no interest or knowledge of the company, we would not be surprised if all these parties had very different expectations.<sup>29</sup>

Historically, this dichotomy of *objective and subjective probability* has received a great deal of attention, and underlies much scholarly activity in many disciplines. For example, objective probability can be seen as underlying the *principle of repeated sampling*, which is central to the statistical methodology of Fisher (Young and Smith, 2005). Indeed, it is hard to justify a procedure of taking controlled random samples to estimate probability unless that probability is “well-defined” in an objective sense. In contrast, Bayesian statistical methods give the statistician the freedom to reflect opinion and knowledge (or a lack thereof) in the inferential procedure, through a prior distribution over the parameter set, which allows for a subjective interpretation of the probabilistic model.<sup>30</sup> Given identical data sets, different statisticians can arrive at different conclusions. Arguably the most influential proponent of subjective probability is that of Leonard J. Savage; in his seminal work (Savage, 1954), he builds a theory of probability and statistics which is centered on “personal” probability. In the context of Bayesian inference, this approach can be implemented by emphasizing a subjective prior (Young and Smith, 2005, Sec. 3.6). Recalling the expected utility theorem described earlier, Savage proves an analogous theorem, but manages to do

---

functions that satisfy the properties of what we call the expected utility (von Neumann and Morgenstern, 1953, Sec. 3.5.1). The requirements that they place on numerical utilities (e.g., their axioms (3:1:a) and (3:1:b)) can be interpreted as conditions on maps from  $\mathcal{X}$  to  $\mathbb{R}$ , without explicitly involving a function  $f$ . Some authors refer to  $f$  as a *Bernoulli utility function*, and  $\mathbf{E}[f(X)]$  as the *von Neumann-Morgenstern utility* (Mas-Colell et al., 1995, Ch. 6, footnote 12).

28. Another important question is that of how to best represent probability in a mathematical form. See Pollard (2002, Ch. 1) and the references therein for a start on how the modern answer to this question was developed.

29. Clearly, many decisions under uncertainty will inevitably involve subjective probabilities. Recent books by Taleb (2008) and Kahneman (2011) aimed at broad audiences give a lucid discussion of the potential ramifications that this subjectivity has on our interaction with complex systems, in particular those of a socioeconomic nature.

30. For a textbook introduction to statistical inference that effectively describes the development of the field within the context of all the major schools of thought, Young and Smith (2005) is a useful reference.

it without assuming given probabilities (used in computing an “objective” expected utility), and instead shows that the rational agent axioms imply a form of *subjective* expected utility maximization.<sup>31</sup> This view of probability and decision-making is prevalent in modern game theory (Osborne and Rubinstein, 1994, Sec. 1.4), and the subjective expected utility theorem is a foundational result in the broader context of decision theory (Steele and Stefánsson, 2020).

The final important historical development we should highlight is that of viewing *statistical inference as a decision-making problem*, in contrast with the classical notion of inference being a rather mechanical summary of data (Young and Smith, 2005, Ch. 1–2). Of particular importance is the work of Abraham Wald, which laid the foundations for statistical decision theory (Wald, 1939, 1945, 1949). One key theme is that from the statistician’s perspective, different types of errors may naturally have a different importance, or weight. Starting with his seminal 1939 work (Wald, 1939), Wald introduces a general-purpose loss function (what he calls a “weight function”) that quantifies this decision-dependent importance, and captures a wide class of statistical problems by formulating them in terms of *minimizing the expected loss*, or what he calls the “risk.”<sup>32</sup> Underlying this “risk,” of course, is a random loss. With Wald’s framework as context, a straightforward generalization beyond typical statistical models and tasks leads us directly to the random test loss  $\ell(h; \mathbf{Z})$  that appears throughout §1–§3, depending on the candidate  $h \in \mathcal{H}$ . In particular, note that the class of random losses  $\{\ell(h; \mathbf{Z}) : h \in \mathcal{H}\}$  is in direct analogy with the class of “lotteries”  $\mathcal{X} = \{X_1, X_2, \dots\}$  discussed earlier in the context of expected utility. Furthermore, an important characteristic of the decision-theoretic approach to statistics is an emphasis upon the optimality of *general inferential procedures*, stated in terms that do not depend on any particular dataset (Young and Smith, 2005, Ch. 1). In short, the qualities of data-driven *algorithms* become an important object of analysis, which makes for a smooth methodological transition to machine learning problems.

## 4.2 Expected Loss and Learning Problems

Here we provide some additional justification for our position (cf. §1.1) that emphasizing performance on average is the dominant paradigm in machine learning. A natural starting point is the Perceptron of Frank Rosenblatt in the late 1950s (Rosenblatt, 1958), and the 1962 proof of the convergence of the Perceptron learning algorithm (here, PLA) for linearly separable data.<sup>33</sup> This important result is recognized as having effectively “started learning theory” (Vapnik, 1999, p. 5). Making this a bit more concrete, given a dataset of  $n$  points  $\mathbf{Z}_n = (Z_1, \dots, Z_n)$ , assumed to be input-output pairs  $Z_i = (X_i, Y_i)$  with binary labels  $Y_i \in \{-1, +1\}$ , and denoting the zero-one training losses as  $L_i(h) := \mathbb{I}\{h(X_i) \neq Y_i\}$  for  $i \in [n]$ , the convergence theorem tells us that whenever perfect accuracy is possible, then

31. See Mas-Colell et al. (1995, Sec. 6.F) for a high-level description and useful references.

32. Let us compare Wald’s terminology with ours (cf. §1.2). The “loss” of Wald (1939) depends on an unknown true model parameter (denoted by  $\theta$  in his notation), so his “loss” corresponds more closely to a criterion  $C_\mu$  here, and his “risk” then amounts to  $\mathbf{E}[C_\mu(H_n)]$ , with expectation taken over  $\mathbf{Z}_n$ .

33. See Novikoff (1963) and Minsky and Papert (1969, Ch. 11) for background on this result.

the PLA output  $\mathbf{Z}_n \mapsto \mathbf{H}_{\text{PLA}}$  will achieve it, namely that

$$\frac{1}{n} \sum_{i=1}^n L_i(\mathbf{H}_{\text{PLA}}) = 0. \quad (64)$$

In the context of learning, the question is then how the PLA *generalizes* to new data. Given an independent test point  $\mathbf{Z} = (\mathbf{X}, \mathbf{Y})$  and writing  $L(h) = \mathbf{I}\{h(\mathbf{X}) \neq \mathbf{Y}\}$ , off-sample misclassification is characterized by  $L(\mathbf{H}_{\text{PLA}})$ . Since this is just a Bernoulli random variable, its distribution is characterized by its mean. With the key fact of (64) in mind, the question of generalization can then be reduced to (uniform) control of the deviations between “relative frequencies of events and their probability,” which is precisely the topic of the seminal 1971 paper of Vapnik and Chervonenkis (1971).

The approach of the original “VC theory” goes well beyond hyperplanes and PLA. Assuming  $\mathbf{Z}$  and  $\mathbf{Z}_n$  are independently sampled from  $\mu$ , for any model  $\mathcal{H}$  of binary classifiers, write  $\widehat{C}_n(h) = (1/n) \sum_{i=1}^n L_i(h)$  and  $C_\mu(h) = \mathbf{P}\{h(\mathbf{X}) \neq \mathbf{Y}\}$  for the empirical and true zero-one risk, respectively, and note that for the output  $\mathbf{H}_n$  of any learning algorithm (not just PLA), we trivially have

$$|\widehat{C}_n(\mathbf{H}_n) - C_\mu(\mathbf{H}_n)| \leq \sup_{h \in \mathcal{H}} |\widehat{C}_n(h) - C_\mu(h)|. \quad (65)$$

The results of (Vapnik and Chervonenkis, 1971) enable us to characterize the convergence properties of the right-hand side of (65) based on properties of the model  $\mathcal{H}$  that are of a combinatorial nature. Having this control is particularly useful for evaluating the performance of empirical risk minimizers, namely any algorithm which minimizes  $\widehat{C}_n$ , such as PLA under linearly separable data. Making this explicit, if we constrain the learning algorithm such that  $\mathbf{H}_n$  minimizes  $\widehat{C}_n$ , since for any fixed reference point  $h^* \in \mathcal{H}$  we have  $\widehat{C}_n(\mathbf{H}_n) \leq \widehat{C}_n(h^*)$ , it follows that

$$\begin{aligned} C_\mu(\mathbf{H}_n) - C_\mu(h^*) &= C_\mu(\mathbf{H}_n) - \widehat{C}_n(\mathbf{H}_n) + \widehat{C}_n(\mathbf{H}_n) - \widehat{C}_n(h^*) + \widehat{C}_n(h^*) - C_\mu(h^*) \\ &\leq 2 \sup_{h \in \mathcal{H}} |\widehat{C}_n(h) - C_\mu(h)|. \end{aligned} \quad (66)$$

Assuming  $C_\mu$  takes its minimum on  $\mathcal{H}$ , an obvious reference point  $h^* \in \mathcal{H}$  is a minimizer of the expected test loss, i.e.,  $C_\mu(h^*) = C_\mu^* := \inf\{C_\mu(h) : h \in \mathcal{H}\}$ . This  $C_\mu^*$  represents the best off-sample performance we can hope to achieve, and thus it is natural to ask that the learning algorithm be able to get arbitrarily close to this benchmark with sufficiently high probability, given enough data. More precisely, we can make the following requirement:

$$\forall \varepsilon > 0, \quad \lim_{n \rightarrow \infty} \mathbf{P} \left\{ C_\mu(\mathbf{H}_n) - C_\mu^* > \varepsilon \right\} = 0. \quad (67)$$

This is just convergence of  $C_\mu(\mathbf{H}_n) \rightarrow C_\mu^*$  in probability, also known as *consistency* of the learning algorithm. There is no conceptual reason to limit this analytical framework to the zero-one loss; *any* random loss  $L(h) := \ell(h; \mathbf{Z})$  with finite expectation  $C_\mu(h) = \mathbf{E}_\mu L(h)$  can in principle be considered. Minimizing the expected value of a general-purpose loss over a hypothesis class  $\mathcal{H}$  is the “general setting” of learning formulated by Vapnik (1999, Ch. 1).

The requirement (67) can be stated in a different but equivalent fashion as follows: for any desired accuracy  $\varepsilon > 0$  and confidence  $0 < \delta < 1$ , there must exist  $N(\varepsilon, \delta) \in \mathbb{N}$  such that for any  $n \geq N(\varepsilon, \delta)$ , we have

$$\mathbf{P} \left\{ C_\mu(\mathbf{H}_n) - C_\mu^* > \varepsilon \right\} \leq \delta. \quad (68)$$

The condition (68) emphasizes *precision* and *confidence* requirements explicitly through the parameters  $\varepsilon$  and  $\delta$  respectively. The highly influential work of Valiant (1984b) considers a learning model which emphasizes precisely these points, with a particular interest in the computational complexity of achieving the  $(\varepsilon, \delta)$ -condition in (68). This became known as the *PAC-learning* model.<sup>34</sup> In effect, the work of Valiant brought the consistency requirement to the community of computer scientists working on machine learning.<sup>35</sup> In this community, as is mentioned by Kearns et al. (1994), it should be noted that the emphasis on the expected value of a generic loss (i.e., the general setting of Vapnik) was largely due to the learning model studied by Haussler (1992), in which the classical risk  $\mathbf{E}_\mu \mathbf{L}(h)$  plays a central role.

To conclude this section, given the historical context established in the preceding paragraphs, we give a brief overview highlighting how “learning as expected loss minimization” is emphasized in many influential and widely-read textbooks. Kearns and Vazirani (1994) give an elementary introduction to the PAC-learning framework and basics of VC theory, with emphasis placed on computational tractability and tasks of interest in computer science. It goes without saying that all the books of Vapnik (Vapnik, 1982, 1998, 1999) are chiefly concerned with the problem of expected loss minimization. While restricted to the pattern recognition problem (binary classification), the highly influential text of Devroye et al. (1996) characterizes the learning problem in terms of  $\mathbf{E}_\mu [\mathbf{L}(\mathbf{H}_n) | \mathbf{Z}_n]$ , formulated using the consistency condition (67). Expected loss is also central to the learning setup in the textbook of Ripley (1996). This setup is carried on by Anthony and Bartlett (1999), in their text on neural network learning. A more modern textbook dealing with statistical and computational learning theory is Mohri et al. (2012), and the expected loss plays a foundational role. Expected loss is used as a central criterion in the introductory text of Shalev-Shwartz and Ben-David (2014), which also emphasizes links between stochastic optimization and machine learning. It is clear that these theory-oriented textbooks are centered around the expected loss, but this notion also appears frequently in more applied machine learning textbooks aimed at introducing key methods and tools to a general audience, albeit in more

34. The original idea of (Valiant, 1984b) was fleshed out further in subsequent works by Valiant (1984a, 1985). The naming of this model given in the main text (PAC: probably approximately correct) appears to date back to Angluin (1988) and Angluin and Laird (1988). The textbook of Kearns and Vazirani (1994) provides good additional background.

35. While the PAC-learning model was the first to receive widespread attention, clearly the general setting of Vapnik supersedes the PAC-learning model; the  $(\varepsilon, \delta)$ -condition (68) appears from the start of a book by Vapnik (1982, Sec. 1.6), a 1982 English translation of a Russian book published in 1979. However, the work of Vapnik was not well-known by computer scientists in the West until the late 1980s, in particular due to the influential work of Blumer et al. (1989), who introduced and built upon the seminal paper of Vapnik and Chervonenkis (1971). As an aside, according to Vapnik (2006, Ch. 2), L. Valiant did not know of his work when authoring his influential 1984 article (Valiant, 1984b). On the other hand, many probability theorists in the West were familiar with the work of Vapnik and Chervonenkis; see for example Dudley (1978) or Pollard (1984).

of a supplementary role. To give a few examples, we have MacKay (2003, Ch. 36), Bishop (2006, Ch. 1), Murphy (2012, Ch. 5–6), Goodfellow et al. (2016, Ch. 5), among others. Given this bibliographic evidence, along with the historical context provided in section 4.1 earlier, it is evident that the notion of expected loss as a metric for off-sample generalization has been pervasive in both the theory and practice of modern machine learning.

## 5. A Look Forward

In the preceding §1–§4, it has been our aim to provide the reader with a unified introductory look at the key classes of criteria that have been studied and applied in the machine learning literature, with the traditional mean-driven paradigm providing important historical context for recent developments. Within this paper, the base loss function  $\ell(\cdot)$ , the decision set  $\mathcal{H}$ , and the desirability set  $\mathcal{C}$  were all left abstract, with focus placed on the learning criterion  $C_\mu(\mathbf{L})$ . Moving forward, a natural direction of interest is to ask how these four elements can interact with each other. Consider the following broad points of inquiry:

- Starting with  $\ell$  and  $C_\mu$ , when can we reduce the problem (7) (under  $\ell$ ) into problem (4) under a different loss function?
- Can we identify conditions (on  $\ell$ ,  $\mathcal{H}$ , and  $\mu$ ) under which traditional norm-based regularization can be emulated using an appropriately designed  $C_\mu$ ?
- How does the theory of surrogate function design change under intractable loss functions (e.g., zero-one) when  $C_\mu$  does not depend linearly on  $\mu$ ? (e.g.,  $C_\mu(\mathbf{L}) = \text{var}_\mu \mathbf{L}$ )

Beyond these points of technical interest, there is the glaring methodological problem that in many cases, it simply is not clear which criterion is “best” for a given problem, since this is often subjective, with precise numerical thresholds difficult to pin down, and will differ from case to case. Solutions to this problem (e.g., metric elicitation (Ali et al., 2022)) are still in their infancy, but there is no question that the development of principled, transparent, and accessible procedures for criteria design, ideally tied in with the back-end algorithmic details, is an important frontier for machine learning research.

## Acknowledgments

This work was supported by JST ACT-X Grant Number JPMJAX2000 and JST PRESTO Grant Number JPMJPR21C6. We also would like to thank Professor Takashi Washio for constructive feedback on an early version of this paper.

## References

- Acerbi, C. (2002). Spectral measures of risk: A coherent representation of subjective risk aversion. *Journal of Banking & Finance*, 26(7):1505–1518.
- Acerbi, C. and Tasche, D. (2002a). Expected shortfall: a natural coherent alternative to value at risk. *Economic Notes*, 31(2):379–388.



- Acerbi, C. and Tasche, D. (2002b). On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7):1487–1503.
- Agarwal, R., Schwarzer, M., Castro, P. S., Courville, A. C., and Bellemare, M. (2021). Deep reinforcement learning at the edge of the statistical precipice. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, pages 29304–29320.
- Ali, S., Upadhyay, S., Hiranandani, G., Glassman, E. L., and Koyejo, O. (2022). Metric elicitation; moving from theory to practice. *arXiv preprint arXiv:2212.03495*.
- Angluin, D. (1988). Queries and concept learning. *Machine learning*, 2(4):319–342.
- Angluin, D. and Laird, P. (1988). Learning from noisy examples. *Machine Learning*, 2(4):343–370.
- Anthony, M. and Bartlett, P. L. (1999). *Neural Network Learning: Theoretical Foundations*. Cambridge University Press.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.
- Ash, R. B. and Doléans-Dade, C. A. (2000). *Probability and Measure Theory*. Academic Press, 2nd edition.
- Bassett, Jr, G. W. (1987). The St. Petersburg paradox and bounded utility. *History of Political Economy*, 19(4):517–523.
- Ben-Tal, A. and Teboulle, M. (1986). Expected utility, penalty functions, and duality in stochastic nonlinear programming. *Management Science*, 32(11):1445–1466.
- Ben-Tal, A. and Teboulle, M. (2007). An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance*, 17(3):449–476.
- Bhat, S. P. and Prashanth, L. A. (2019). Concentration of risk measures: A Wasserstein distance approach. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*.
- Bickel, P. J. and Lehmann, E. L. (1975). Descriptive statistics for nonparametric models ii. location. *The Annals of Statistics*, 3(5):1045–1069.
- Bickel, P. J. and Lehmann, E. L. (1981). A minimax property of the sample mean in finite populations. *The Annals of Statistics*, 9(5):1119–1122.
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.
- Blumer, A., Ehrenfeucht, A., Haussler, D., and Warmuth, M. K. (1989). Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36(4):929–965.
- Boucheron, S., Lugosi, G., and Massart, P. (2013). *Concentration inequalities: a nonasymptotic theory of independence*. Oxford University Press.

- Breiman, L. (1999). Prediction games and arcing algorithms. *Neural Computation*, 11(7):1493–1517.
- Brownlees, C., Joly, E., and Lugosi, G. (2015). Empirical risk minimization for heavy-tailed losses. *Annals of Statistics*, 43(6):2507–2536.
- Cardoso, A. R. and Xu, H. (2019). Risk-averse stochastic convex bandit. In *22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 89 of *Proceedings of Machine Learning Research*, pages 39–47.
- Catoni, O. (2012). Challenging the empirical mean and empirical variance: a deviation study. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 48(4):1148–1185.
- Chow, Y., Ghavamzadeh, M., Janson, L., and Pavone, M. (2017). Risk-constrained reinforcement learning with percentile risk criteria. *Journal of Machine Learning Research*, 18(1):6070–6120.
- Chow, Y., Tamar, A., Mannor, S., and Pavone, M. (2016). Risk-sensitive and robust decision-making: a CVaR optimization approach. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 1522–1530.
- Cucker, F. and Smale, S. (2002). On the mathematical foundations of learning. *Bulletin (New Series) of the American Mathematical Society*, 39(1):1–49.
- Curi, S., Levy, K. Y., Jegelka, S., and Krause, A. (2020). Adaptive sampling for stochastic risk-averse learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, pages 1036–1047.
- Daniell, P. J. (1920). Observations weighted according to order. *American Journal of Mathematics*, 42(4):222–236.
- Devroye, L., Györfi, L., and Lugosi, G. (1996). *A Probabilistic Theory of Pattern Recognition*. Springer.
- Devroye, L., Lerasle, M., Lugosi, G., and Oliveira, R. I. (2016). Sub-Gaussian mean estimators. *The Annals of Statistics*, 44(6):2695–2725.
- Duchi, J. and Namkoong, H. (2019). Variance-based regularization with convex objectives. *Journal of Machine Learning Research*, 20(68):1–55.
- Duchi, J. C., Glynn, P. W., and Namkoong, H. (2016). Statistics of robust optimization: A generalized empirical likelihood approach. *arXiv preprint arXiv:1610.03425v3*.
- Duchi, J. C., Glynn, P. W., and Namkoong, H. (2021). Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 46(3):835–1234.
- Duchi, J. C. and Namkoong, H. (2016). Variance-based regularization with convex objectives. *arXiv preprint arXiv:1610.02581v3*.

- Duchi, J. C. and Namkoong, H. (2018). Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750v6*.
- Duchi, J. C. and Namkoong, H. (2021). Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics*, 49(3):1378–1406.
- Dudley, R. M. (1978). Central limit theorems for empirical measures. *Annals of Probability*, 6(6):899–929.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. (2012). Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 214–226.
- Fan, Y., Lyu, S., Ying, Y., and Hu, B. (2017a). Learning with average top-k loss. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*.
- Fan, Y., Lyu, S., Ying, Y., and Hu, B.-G. (2017b). Learning with average top-k loss. *arXiv preprint arXiv:1705.08826*.
- Finocchiaro, J., Frongillo, R., and Waggoner, B. (2021). Unifying lower bounds on prediction dimension of consistent convex surrogates. *arXiv preprint arXiv:2102.08218*.
- Föllmer, H. and Knispel, T. (2011). Entropic risk measures: Coherence vs. convexity, model ambiguity and robust large deviations. *Stochastics and Dynamics*, 11(02n03):333–351.
- Fröhlich, C. and Williamson, R. C. (2023). Tailoring to the tails: Risk measures for fine-grained tail sensitivity. *Transactions on Machine Learning Research*.
- García, J. and Fernández, F. (2015). A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- Gotoh, J.-y., Kim, M. J., and Lim, A. E. (2018). Robust empirical optimization is almost the same as mean–variance optimization. *Operations Research Letters*, 46(4):448–452.
- Grünwald, P. D. and Dawid, A. P. (2004). Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *The Annals of Statistics*, pages 1367–1433.
- Hacking, I. (2006). *The emergence of probability: a philosophical study of early ideas about probability, induction and statistical inference*. Cambridge University Press, 2nd edition.
- Hashimoto, T. B., Srivastava, M., Namkoong, H., and Liang, P. (2018). Fairness without demographics in repeated loss minimization. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 1929–1938.
- Hausler, D. (1992). Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100(1):78–150.

- Holland, M. J. (2020). Non-monotone risk functions for learning. *arXiv preprint arXiv:2012.02424v1*.
- Holland, M. J. (2022). Learning with risks based on M-location. *Machine Learning*, 111:4679–4718.
- Holland, M. J. (2023). Flexible risk design using bi-directional dispersion. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 206 of *Proceedings of Machine Learning Research*, pages 1586–1623.
- Holland, M. J. and Haress, E. M. (2021). Learning with risk-averse feedback under potentially heavy tails. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 130 of *Proceedings of Machine Learning Research*.
- Holland, M. J. and Haress, E. M. (2022). Spectral risk-based learning using unbounded losses. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 151 of *Proceedings of Machine Learning Research*, pages 1871–1886.
- Holland, M. J. and Ikeda, K. (2017). Robust regression using biased objectives. *Machine Learning*, 106(9):1643–1679.
- Hu, S., Wang, X., and Lyu, S. (2022). Rank-based decomposable losses in machine learning: A survey. *arXiv preprint arXiv:2207.08768v1*.
- Huber, P. J. (1964). Robust estimation of a location parameter. *Annals of Mathematical Statistics*, 35(1):73–101.
- Huber, P. J. and Ronchetti, E. M. (2009). *Robust Statistics*. John Wiley & Sons, 2nd edition.
- Jordan, C. (1924). On Daniel Bernoulli’s “moral expectation” and on a new conception of expectation. *American Mathematical Monthly*, 31(4):183–190.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kashima, H. (2007). Risk-sensitive learning via minimization of empirical conditional value-at-risk. *IEICE Transactions on Information and Systems*, 90(12):2043–2052.
- Kearns, M. J., Schapire, R. E., and Sellie, L. M. (1994). Toward efficient agnostic learning. *Machine Learning*, 17(2-3):115–141.
- Kearns, M. J. and Vazirani, U. V. (1994). *An Introduction to Computational Learning Theory*. MIT Press.
- Khan, M. E. and Rue, H. (2021). The bayesian learning rule. *arXiv preprint arXiv:2107.04562*.

- Khim, J., Liu, L., Prasad, A., and Ravikumar, P. (2020). Uniform convergence of rank-weighted learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pages 5254–5263.
- Koltchinskii, V. I. (1997).  $M$ -estimation, convexity and quantiles. *The Annals of Statistics*, pages 435–477.
- Lambert, N. S., Pennock, D. M., and Shoham, Y. (2008). Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 129–138.
- Lee, J., Park, S., and Shin, J. (2020). Learning bounds for risk-sensitive learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, pages 13867–13879.
- Lerasle, M. and Oliveira, R. I. (2011). Robust empirical mean estimators. *arXiv preprint arXiv:1112.3914*.
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. (2020). Tilted empirical risk minimization. *arXiv preprint arXiv:2007.01162*.
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. (2021a). On tilted losses in machine learning: Theory and applications. *arXiv preprint arXiv:2109.06141v2*.
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. (2021b). Tilted empirical risk minimization. In *The 9th International Conference on Learning Representations (ICLR)*.
- Liu, L., Huang, A., Lipton, Z., and Azizzadenesheli, K. (2022). Supervised learning with general risk functionals. In *39th International Conference on Machine Learning (ICML)*, volume 162 of *Proceedings of Machine Learning Research*, pages 12570–12592.
- Liu, L., Prasad, A., and Ravikumar, P. K. (2019). On human-aligned risk minimization. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*.
- Lugosi, G. and Mendelson, S. (2019a). Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics*, 19(5):1145–1190.
- Lugosi, G. and Mendelson, S. (2019b). Robust multivariate mean estimation: the optimality of trimmed mean. *arXiv preprint arXiv:1907.11391v1*.
- MacKay, D. J. C. (2003). *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1):77–91.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- Maurer, A., Parletta, D. A., Paudice, A., and Pontil, M. (2020). Robust unsupervised learning via L-statistic minimization. *arXiv preprint arXiv:2012.07399v3*.

- Maurer, A. and Pontil, M. (2009). Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the 22nd Conference on Learning Theory (COLT)*.
- Menger, K. (1967). The role of uncertainty in economics. In *Essays in Mathematical Economics, in Honor of Oskar Morgenstern*, pages 211–232. Princeton University Press. English translation of 1934 original written in German.
- Menon, A. K., Jayasumana, S., Rawat, A. S., Jain, H., Veit, A., and Kumar, S. (2021). Long-tail learning via logit adjustment. In *The 9th International Conference on Learning Representations (ICLR)*.
- Minsky, M. and Papert, S. A. (1969). *Perceptrons: An Introduction to Computational Geometry*. MIT Press. Reprinted in 2017 by MIT Press.
- Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2012). *Foundations of Machine Learning*. MIT Press.
- Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press.
- Namkoong, H. and Duchi, J. C. (2016). Stochastic gradient methods for distributionally robust optimization with  $f$ -divergences. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, volume 29, pages 2208–2216.
- Novikoff, A. B. J. (1963). On convergence proofs for perceptrons. Technical report, Stanford Research Institute, Menlo Park California.
- Ogryczak, W. and Tamir, A. (2003). Minimizing the sum of the  $k$  largest functions in linear time. *Information Processing Letters*, 85(3):117–122.
- Oja, H. (2012). Descriptive statistics for nonparametric models. The impact of some Erich Lehmann’s papers. In *Selected Works of E. L. Lehmann*, Selected Works in Probability and Statistics, pages 451–457. Springer.
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press.
- Pandey, A. K., Prashanth, L. A., and Bhat, S. P. (2019). Estimation of spectral risk measures. *arXiv preprint arXiv:1912.10398*.
- Pee, E. Y. and Royset, J. O. (2011). On solving large-scale finite minimax problems using exponential smoothing. *Journal of Optimization Theory and Applications*, 148(2):390–421.
- Pollard, D. (1984). *Convergence of Stochastic Processes*. Springer-Verlag.
- Pollard, D. (2002). *A User’s Guide to Measure Theoretic Probability*. Cambridge University Press.
- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1–2):122–136.

- Quiñonero Candela, J., Sugiyama, M., Schwaighofer, A., and Lawrence, N. D. (2009). *Dataset shift in machine learning*. MIT Press.
- Reyzin, L. and Schapire, R. E. (2006). How boosting the margin can also boost classifier complexity. In *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, pages 753–760.
- Rieger, M. O. and Wang, M. (2006). Cumulative prospect theory and the St. Petersburg paradox. *Economic Theory*, 28(3):665–679.
- Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge University Press.
- Robey, A., Chamon, L. F., Pappas, G. J., and Hassani, H. (2022). Probabilistically robust learning: Balancing average-and worst-case performance. *arXiv preprint arXiv:2202.01136*.
- Rockafellar, R. T. and Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42.
- Rockafellar, R. T. and Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471.
- Rockafellar, R. T. and Uryasev, S. (2013). The fundamental risk quadrangle in risk management, optimization and statistical estimation. *Surveys in Operations Research and Management Science*, 18(1-2):33–53.
- Rockafellar, R. T., Uryasev, S., and Zabarankin, M. (2006). Generalized deviations in risk analysis. *Finance and Stochastics*, 10(1):51–74.
- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408.
- Royset, J. O. (2022). Risk-adaptive approaches to learning and decision making: A survey. *arXiv preprint arXiv:2212.00856*.
- Ruszczynski, A. and Shapiro, A. (2006a). Optimization of convex risk functions. *Mathematics of Operations Research*, 31(3):433–452.
- Ruszczynski, A. and Shapiro, A. (2006b). Optimization of risk measures. In *Probabilistic and Randomized Methods for Design under Uncertainty*, pages 119–157. Springer.
- Samadi, S., Tantipongpipat, U., Morgenstern, J. H., Singh, M., and Vempala, S. (2018). The price of fair pca: One extra dimension. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*.
- Samuelson, P. A. (1977). St. Petersburg paradoxes: Defanged, dissected, and historically described. *Journal of Economic Literature*, 15(1):24–55.
- Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons. Dover 1972 edition of the 1954 original.

- Savage, L. J. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801.
- Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- Shalev-Shwartz, S. and Wexler, Y. (2016). Minimizing the maximal loss: How and why. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 793–801.
- Shapiro, A. (2013). On Kusuoka representation of law invariant risk measures. *Mathematics of Operations Research*, 38(1):142–152.
- Shapiro, A. (2017). Distributionally robust stochastic programming. *SIAM Journal on Optimization*, 27(4):2258–2275.
- Soma, T. and Yoshida, Y. (2020). Statistical learning with conditional value at risk. *arXiv preprint arXiv:2002.05826*.
- Steele, K. and Stefánsson, H. O. (2020). Decision Theory. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2020 edition.
- Stuart, A. and Ord, J. K. (1994). *Kendall’s Advanced Theory of Statistics Volume 1: Distribution Theory*. Hodder Arnold, 6th edition.
- Sutter, T., Krause, A., and Kuhn, D. (2021). Robust generalization despite distribution shift via minimum discriminating information. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, pages 29754–29767.
- Takeda, A. and Sugiyama, M. (2008).  $\nu$ -support vector machine as conditional value-at-risk minimization. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, pages 1056–1063.
- Taleb, N. N. (2008). *The Black Swan: The Impact of the Highly Improbable*. Penguin, 2nd edition.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323.
- Uzawa, H. (1956). Petersburg no mondai wo megutte [On the Petersburg problem]. *Kagaku Kisoron Kenkyuu*, 2(4):342–348.
- Valiant, L. G. (1984a). Deductive learning. *Philosophical Transactions of the Royal Society of London, Series A*, 312(1522):441–446.
- Valiant, L. G. (1984b). A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142.
- Valiant, L. G. (1985). Learning disjunction of conjunctions. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 560–566. Morgan Kaufmann.



- Vapnik, V. (1982). *Estimation of Dependences Based on Empirical Data*. Springer Series in Statistics. Springer-Verlag.
- Vapnik, V. (2006). Empirical inference science. In *Estimation of Dependences Based on Empirical Data: 2nd Edition*. Springer.
- Vapnik, V. N. (1998). *Statistical Learning Theory*. Wiley.
- Vapnik, V. N. (1999). *The Nature of Statistical Learning Theory*. Statistics for Engineering and Information Science. Springer, 2nd edition.
- Vapnik, V. N. and Chervonenkis, A. Y. (1971). On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280.
- von Neumann, J. and Morgenstern, O. (1953). *Theory of Games and Economic Behavior*. Princeton University Press, 3rd edition. Sixtieth-anniversary edition published in 2004.
- Wald, A. (1939). Contributions to the theory of statistical estimation and testing hypotheses. *Annals of Mathematical Statistics*, 10(4):299–326.
- Wald, A. (1945). Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, pages 265–280.
- Wald, A. (1949). Statistical decision functions. *Annals of Mathematical Statistics*, 20(2):165–205.
- Williamson, R. C. and Menon, A. K. (2019). Fairness risk measures. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 6786–6797.
- Young, G. A. and Smith, R. L. (2005). *Essentials of Statistical Inference*, volume 16 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press.
- Zhai, R., Dan, C., Kolter, J. Z., and Ravikumar, P. (2021a). DORO: Distributional and outlier robust optimization. In *38th International Conference on Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pages 12345–12355.
- Zhai, R., Dan, C., Suggala, A., Kolter, Z., and Ravikumar, P. (2021b). Boosted CVaR classification. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, pages 21860–21871.