Optimal Binary Classifier Aggregation for General Losses

Akshay Balsubramani University of California, San Diego abalsubr@ucsd.edu Yoav Freund

University of California, San Diego yfreund@ucsd.edu

Abstract

We address the problem of aggregating an ensemble of predictors with known loss bounds in a semi-supervised binary classification setting, to minimize prediction loss incurred on the unlabeled data. We find the minimax optimal predictions for a very general class of loss functions including all convex and many non-convex losses, extending a recent analysis of the problem for misclassification error. The result is a family of semi-supervised ensemble aggregation algorithms which are as efficient as linear learning by convex optimization, but are minimax optimal without any relaxations. Their decision rules take a form familiar in decision theory – applying sigmoid functions to a notion of ensemble margin – without the assumptions typically made in margin-based learning.

1 Introduction

Consider a binary classification problem, in which we are given an ensemble of individual classifiers to aggregate into the most accurate predictor possible for data falling into two classes. Our predictions are measured on a large test set of unlabeled data, on which we know the ensemble classifiers' predictions but not the true test labels. Without using the unlabeled data, the prototypical supervised solution is empirical risk minimization (ERM): measure the errors of the ensemble classifiers with labeled data, and then simply predict according to the best classifier. But can we learn a better predictor by using unlabeled data as well?

This problem is central to semi-supervised learning. The authors of this paper recently derived the worst-case-optimal solution for it when performance is measured with classification error ([1]). However, this zero-one loss is inappropriate for other common binary classification tasks, such as estimating label probabilities, and handling false positives and false negatives differently. Such goals motivate the use of different evaluation losses like log loss and cost-weighted misclassification loss.

In this paper, we generalize the setup of [1] to these loss functions and a large class of others. Like the earlier work, the choice of loss function completely specifies the minimax optimal ensemble aggregation algorithm in our setting, which is efficient and scalable.

The algorithm learns weights over the ensemble classifiers by minimizing a convex function. The optimal prediction on each example in the test set is a sigmoid-like function of a linear combination of the ensemble predictions, using the learned weighting. Due to the minimax structure, this decision rule depends solely upon the loss function and upon the structure of the ensemble predictions on data, with no parameter or model choices.

1.1 Preliminaries

Our setting generalizes that of [1], in which we are given an ensemble $\mathcal{H} = \{h_1, \dots, h_p\}$ and unlabeled (test) examples x_1, \dots, x_n on which to predict. The ensemble's predictions on the unlabeled

data are written as a matrix **F**:

$$\mathbf{F} = \begin{pmatrix} h_1(x_1) & h_1(x_2) & \cdots & h_1(x_n) \\ \vdots & \vdots & \ddots & \vdots \\ h_p(x_1) & h_p(x_2) & \cdots & h_p(x_n) \end{pmatrix}$$
(1)

We use vector notation for the rows and columns of \mathbf{F} : $\mathbf{h}_i = (h_i(x_1), \cdots, h_i(x_n))^{\top}$ and $\mathbf{x}_j = (h_1(x_j), \cdots, h_p(x_j))^{\top}$. Each example $j \in [n]$ has a binary label $y_j \in \{-1, 1\}$, but the test labels are allowed to be randomized, represented by values in [-1, 1] instead of just the two values $\{-1, 1\}$; e.g. $z_i = \frac{1}{2}$ indicates $y_i = +1$ w.p. $\frac{3}{4}$ and -1 w.p. $\frac{1}{4}$. So the labels on the test data can be represented by $\mathbf{z} = (z_1; \ldots; z_n) \in [-1, 1]^n$, and are unknown to the predictor, which predicts $\mathbf{g} = (g_1; \ldots; g_n) \in [-1, 1]^n$.

1.2 Loss Functions

We incur loss on test example j according to its true label y_j . If $y_j=1$, then the loss of predicting $g_j\in [-1,1]$ on it is some function $\ell_+(g_j)$; and if $y_j=-1$, then the loss is $\ell_-(g_j)$. To illustrate, if the loss is the expected classification error, then $g_j\in [-1,1]$ can be interpreted as a randomized binary prediction in the same way as z_j , so that $\ell_+(g_j)=\frac{1}{2}\left(1-g_j\right)$ and $\ell_-(g_j)=\frac{1}{2}\left(1+g_j\right)$.

We call ℓ_{\pm} the partial losses here, following earlier work (e.g. [16]). Since the true label can only be ± 1 , the partial losses fully specify the decision-theoretic problem we face, and changing them is tantamount to altering the prediction task.

What could such partial losses conceivably look like in general? Observe that they intuitively measure discrepancy to the true label ± 1 . Consequently, it is natural for e.g. $\ell_+(g)$ to be decreasing, as g increases toward the notional true label +1. This suggests that both partial losses $\ell_+(\cdot)$ and $\ell_-(\cdot)$ would be monotonic, which we assume hereafter in this paper (throughout we use *increasing* to mean "monotonically nondecreasing" and vice versa).

Assumption 1. Over the interval (-1,1), $\ell_+(\cdot)$ is decreasing and $\ell_-(\cdot)$ is increasing, and both are twice differentiable.

We view Assumption 1 as very mild, as motivated above. Notably, convexity or symmetry of the partial losses are not required. In this paper, "general losses" refer to loss functions whose partial losses satisfy Assumption 1, to contrast them with convex losses or other subclasses.

The expected loss incurred w.r.t. the *randomized* true labels z_j is a linear combination of the partial losses:

$$\ell(z_j, g_j) := \left(\frac{1 + z_j}{2}\right) \ell_+(g_j) + \left(\frac{1 - z_j}{2}\right) \ell_-(g_j) \tag{2}$$

Decision theory and learning theory have thoroughly investigated the nature of the loss ℓ and its partial losses, particularly how to estimate the "conditional label probability" z_j using $\ell(z_j, g_j)$. A natural operation to do this is to minimize the loss over g_j , and a loss ℓ such that $\arg\min_{g \in [-1,1]} \ell(z_j, g) = \lim_{g \in [-1,1]} \ell(z_j, g)$

 z_i (for all $z_i \in [-1, 1]$) is called a proper loss ([17, 16]).

1.3 Minimax Formulation

As in [1], we formulate the ensemble aggregation problem as a two-player zero-sum game between a predictor and an adversary. In this game, the first player is the predictor, playing predictions over the test set $\mathbf{g} \in [-1,1]^n$. The adversary then sets the true labels $\mathbf{z} \in [-1,1]^n$.

The key idea is that any ensemble constituent $i \in [p]$ known to have low loss on the test data gives us information about the unknown \mathbf{z} , as \mathbf{z} is constrained to be "close" to the test predictions \mathbf{h}_i . Each hypothesis in the ensemble represents such a constraint, and \mathbf{z} is in the intersection of all these constraint sets, which interact in ways that depend on the ensemble predictions \mathbf{F} .

Accordingly, for now assume the predictor knows a vector of label correlations b such that

$$\forall i \in [p]: \qquad \frac{1}{n} \sum_{j=1}^{n} h_i(x_j) z_j \ge b_i \tag{3}$$

i.e. $\frac{1}{n}\mathbf{Fz} \geq \mathbf{b}$. When the ensemble is composed of binary classifiers which predict in [-1,1], these p inequalities represent upper bounds on individual classifier error rates. These can be estimated from the training set w.h.p. when the training and test data are i.i.d. using uniform convergence, exactly as in the prototypical supervised ERM procedure discussed in the introduction ([5]). So in our game-theoretic formulation, the adversary plays under ensemble constraints defined by \mathbf{b} .

The predictor's goal is to *minimize the worst-case expected loss of* **g** *on the test data* (w.r.t. the randomized labeling **z**), using the loss function as defined earlier in Equation (2):

$$\ell(\mathbf{z}, \mathbf{g}) := \frac{1}{n} \sum_{j=1}^{n} \ell(z_j, g_j)$$

This goal can be written as the following optimization problem, a two-player zero-sum game:

$$V := \min_{\mathbf{g} \in [-1,1]^n} \max_{\mathbf{z} \in [-1,1]^n, \frac{1}{n} \mathbf{F} \mathbf{z} \ge \mathbf{b}} \ell(\mathbf{z}, \mathbf{g})$$

$$(4)$$

$$= \min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n} \mathbf{F} \mathbf{z} \ge \mathbf{b}}} \frac{1}{n} \sum_{j=1}^n \left[\left(\frac{1+z_j}{2} \right) \ell_+(g_j) + \left(\frac{1-z_j}{2} \right) \ell_-(g_j) \right]$$
(5)

In this paper, we solve the learning problem faced by the predictor, finding an optimal strategy \mathbf{g}^* realizing the minimum in (4) for any given "general loss" ℓ . This strategy guarantees the best possible worst-case performance on the unlabeled dataset, with an upper bound of V on the loss. Indeed, for all \mathbf{z}_0 and \mathbf{g}_0 obeying the constraints, Equation (4) implies the tight inequalities

$$\min_{\mathbf{g} \in [-1,1]^n} \ell(\mathbf{z}_0, \mathbf{g}) \stackrel{(a)}{\leq} V \leq \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \ell(\mathbf{z}, \mathbf{g}_0)$$
(6)

and g^* attains the equality in (a), with a worst-case loss as good as *any* aggregated predictor.

In our formulation of the problem, the constraints on the adversary take a central role. As discussed in previous work with this formulation ([1, 2]), these constraints encode the information we have about the true labels. Without them, the adversary would find it optimal to trivially guarantee error (arbitrarily close to) $\frac{1}{2}$ by simply setting all labels uniformly at random ($\mathbf{z} = \mathbf{0}^n$). It is clear that adding information through more constraints will never raise the error bound V. ¹

Nothing has yet been assumed about $\ell(\mathbf{z}, \mathbf{g})$ other than Assumption 1. Our main results will require only this, holding for general losses. This brings us to this paper's contributions:

- 1. We give the exact minimax $\mathbf{g}^* \in [-1,1]^n$ for general losses (Section 2.1). The optimal prediction on each example j is a sigmoid function of a fixed linear combination of the ensemble's p predictions on it, so \mathbf{g}^* is a non-convex function of the ensemble predictions. By (6), this incurs the lowest worst-case loss of any predictor constructed with the ensemble information \mathbf{F} and \mathbf{b} .
- 2. We derive an efficient algorithm for learning g*, by solving a *p*-dimensional convex optimization problem. This applies to a broad class of losses (cf. Lem. 2), including any with convex partial losses. Sec. 2 develops and discusses the results.
- 3. We extend the optimal g* and efficient learning algorithm for it, as above, to a large variety of more general ensembles and prediction scenarios (Sec. 3), including constraints arising from general loss bounds, and ensembles of "specialists" and heterogeneous features.

2 Results for Binary Classification

Based on the loss, define the function $\Gamma: [-1,1] \mapsto \mathbb{R}$ as $\Gamma(g) := \ell_{-}(g) - \ell_{+}(g)$. (We also write the vector $\Gamma(\mathbf{g})$ componentwise with $[\Gamma(\mathbf{g})]_j = \Gamma(g_j)$ for convenience, so that $\Gamma(\mathbf{h}_i) \in \mathbb{R}^n$ and $\Gamma(\mathbf{x}_j) \in \mathbb{R}^p$.) Observe that by Assumption 1, $\Gamma(g)$ is increasing on its domain; so we can discuss its inverse $\Gamma^{-1}(m)$, which is typically sigmoid-shaped, as will be illustrated.

With these we will set up the solution to the game (4), which relies on a convex function.

However, it may pose difficulties in estimating b by applying uniform convergence over a larger \mathcal{H} ([2]).

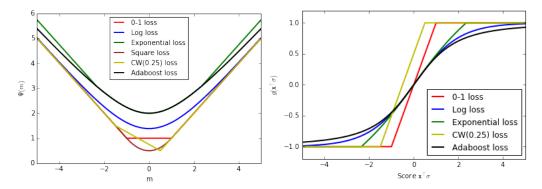


Figure 1: At left are plots of potential wells. At right are optimal prediction functions g, as a function of score. Both are shown for various losses, as listed in Section 2.3.

Definition 1 (Potential Well). Define the potential well

$$\Psi(m) := \begin{cases} -m + 2\ell_{-}(-1) & \text{if } m \leq \Gamma(-1) \\ \ell_{+}(\Gamma^{-1}(m)) + \ell_{-}(\Gamma^{-1}(m)) & \text{if } m \in (\Gamma(-1), \Gamma(1)) \\ m + 2\ell_{+}(1) & \text{if } m \geq \Gamma(1) \end{cases}$$

Lemma 2. The potential well $\Psi(m)$ is continuous and 1-Lipschitz. It is also convex under any of the following conditions:

- (A) The partial losses $\ell_{\pm}(\cdot)$ are convex over (-1,1).
- (B) The loss function $\ell(\cdot, \cdot)$ is a proper loss. (C) $\ell'_{-}(x)\ell''_{+}(x) \geq \ell''_{-}(x)\ell'_{+}(x)$ for all $x \in (-1, 1)$.

Condition (C) is also necessary for convexity of Ψ , under Assumption 1.

So the potential wells for different losses are shaped similarly, as seen in Figure 1. Lemma 2 tells us that the potential well is easy to optimize under any of the given conditions. Note that these conditions encompass convex surrogate losses commonly used in ERM, including all such "marginbased" losses (convex univariate functions of $z_i g_i$), introduced primarily for their favorable computational properties.

An easily optimized potential well benefits us, because the learning problem basically consists of optimizing it over the unlabeled data, as we will soon make explicit. The function that will actually be optimized is in terms of the dual parameters, so we call it the slack function.

Definition 3 (Slack Function). Let $\sigma \geq \mathbf{0}^p$ be a weight vector over \mathcal{H} (not necessarily a distribution). The vector of scores is $\mathbf{F}^{\top} \sigma = (\mathbf{x}_1^{\top} \sigma, \dots, \mathbf{x}_n^{\top} \sigma)$, whose elements' magnitudes are the margins. The prediction slack function is

$$\gamma(\sigma, \mathbf{b}) := \gamma(\sigma) := -\mathbf{b}^{\top} \sigma + \frac{1}{n} \sum_{j=1}^{n} \Psi(\mathbf{x}_{j}^{\top} \sigma)$$
 (7)

An optimal weight vector σ^* is any minimizer of the slack function: $\sigma^* \in \underset{\sigma \geq \mathbf{0}^p}{\arg\min} \left[\gamma(\sigma) \right]$.

2.1 **Solution of the Game**

These are used to describe the minimax equilibrium of the game (4), in our main result.

Theorem 4. The minimax value of the game (4) is

$$\min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \ell(\mathbf{z}, \mathbf{g}) = V = \frac{1}{2}\gamma(\sigma^*) = \frac{1}{2}\min_{\sigma \geq \mathbf{0}^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \sum_{j=1}^n \Psi(\mathbf{x}_j^\top \sigma) \right]$$

The minimax optimal predictions are defined as follows: for all $j \in [n]$,

$$g_j^* := g_j(\sigma^*) = \begin{cases} -1 & \text{if } \mathbf{x}_j^\top \sigma^* \le \Gamma(-1) \\ \Gamma^{-1}(\mathbf{x}_j^\top \sigma^*) & \text{if } \mathbf{x}_j^\top \sigma^* \in (\Gamma(-1), \Gamma(1)) \\ 1 & \text{if } \mathbf{x}_j^\top \sigma^* \ge \Gamma(1) \end{cases}$$
(8)

 g_i^* is always an increasing sigmoid, as shown in Figure 1.

We can also redo the proof of Theorem 4 when $\mathbf{g} \in [-1,1]^n$ is not left as a free variable set in the game, but instead is preset to $\mathbf{g}(\sigma)$ as in (8) for some (possibly suboptimal) weight vector σ .

Observation 5. For any weight vector $\sigma_0 \geq \mathbf{0}^p$, the worst-case loss after playing $\mathbf{g}(\sigma_0)$ is

$$\max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \ell(\mathbf{z}, \mathbf{g}(\sigma_0)) \leq \frac{1}{2}\gamma(\sigma_0)$$

The proof is a simplified version of that of Theorem 4; there is no minimum over g to deal with, and the minimum over $\sigma \geq \mathbf{0}^p$ in Equation (13) is upper-bounded by using σ_0 . This result is an expression of weak duality in our setting, and generalizes Observation 4 of [1].

2.2 Ensemble Aggregation Algorithm

Theorem 4 defines a prescription for aggregating the given ensemble predictions on the test set.

Learning: Minimize the slack function $\gamma(\sigma)$, finding the minimizer σ^* that achieves V.

This is a convex optimization under broad conditions (Lemma 2), and when the test examples are i.i.d. the Ψ term is a sum of n i.i.d. functions. Therefore, it is readily amenable to standard first-order optimization methods which require only O(1) test examples at once. In practice, learning employs such methods to *approximately* minimize γ , finding some σ_A such that $\gamma(\sigma_A) \leq \gamma(\sigma^*) + \epsilon$ for some small ϵ . Standard convex optimization methods are guaranteed to do this for binary classifier ensembles, because the slack function is Lipschitz (Lemma 2) and $\|\mathbf{b}\|_{\infty} \leq 1$.

Prediction: Predict $g(\sigma^*)$ on any test example, as indicated in (8).

This decouples the prediction task over each test example separately, which requires O(p) time and memory like p-dimensional linear prediction. After finding an ϵ -approximate minimizer σ_A in the learning step as above, Observation 5 tells us that the prediction $\mathbf{g}(\sigma_A)$ has loss $\leq V + \frac{\epsilon}{2}$.

In particular, note that there is no algorithmic dependence on n in either step in a statistical learning setting. So though our formulation is transductive, it is no less tractable than a stochastic optimization setting in which i.i.d. data arrive one at a time, and applies to this common situation.

2.3 Examples of Different Losses

To further illuminate Theorem 4, we detail a few special cases in which ℓ_{\pm} are explicitly defined. These losses may be found throughout the literature (see e.g. [16]). The key functions Ψ and g^* are listed for these losses in Appendix A, and in many cases in Figure 1. The nonlinearities used for g^* are sigmoids, arising solely from the intrinsic minimax structure of the classification game.

- **0-1 Loss**: Here g_j is taken to be a randomized binary prediction; this case was developed in [1], the work we generalize in this paper.
- Log Loss, Square Loss
- Cost-Weighted Misclassification (Quantile) Loss: This is defined with a parameter $c \in [0, 1]$ representing the relative cost of false positives vs. false negatives, making the Bayes-optimal classifier the c-quantile of the conditional probability distribution ([19]).
- Exponential Loss, Logistic Loss
- **Hellinger Loss**: This is typically given for $p,y \in [0,1]$ as $\frac{1}{2}\left(\left(\sqrt{p}-\sqrt{y}\right)^2+\left(\sqrt{1-p}-\sqrt{1-y}\right)^2\right)$. Our formulation is equivalent when the prediction and label are rescaled to [-1,1].

- "AdaBoost Loss": If the goal of AdaBoost ([18]) is interpreted as class probability estimation, the implied loss is proper and given in [6, 16].
- Absolute Loss and Hinge Loss: The absolute loss can be defined by $\ell_{\mp}^{abs}(g_j) = 1 \pm g_j$, and the hinge loss also has $\ell_{\mp}^{abs}(g_j) = 1 \pm g_j$ since the kink in the hinge loss only lies at $g_j = \mp 1$. These partial losses are the same as for 0-1 loss up to scaling, and therefore all our results for Ψ and \mathbf{g}^* are as well. So these losses are not shown in Appendix A.
- **Sigmoid Loss**: This is an example of a sigmoid-shaped margin loss, a nonconvex smooth surrogate for 0-1 loss. Similar losses have arisen in a variety of binary classification contexts, from robustness (e.g. [9]) to active learning ([10]) and structured prediction ([14]).

2.4 Related Work and Technical Discussion

There are two notable ways in which the result of Theorem 4 is particularly advantageous and general. First, the fact that $\ell(z,g)$ can be non-convex in g, yet solvable by convex optimization, is a major departure from previous work. Second, the solution has a convenient dependence on n (as in [1]), simply averaging a function over the unlabeled data, which is not only mathematically convenient but also makes stochastic O(1)-space optimization practical. This is surprisingly powerful, because the original minimax problem is *jointly* over the entire dataset, avoiding further independence or decoupling assumptions.

Both these favorable properties stem from the structure of the binary classification problem, as we can describe by examining the optimization problem constructed within the proof of Thm. 4 (Appendix C.1). In it, the constraints which do not explicitly appear with Lagrange parameters are all box, or L_{∞} norm, constraints. These decouple over the n test examples, so the problem can be reduced to the one-dimensional optimization at the heart of Eq. (14), which is solved ad hoc. So we are able to obtain minimax results for these non-convex problems – the g_i are "clipped" sigmoid functions because of the bounding effect of the [-1,1] box constraints intrinsic to binary classification. We introduce Lagrange parameters σ only for the p remaining constraints in the problem, which do not decouple as above, applying globally over the p test examples. However, these constraints only depend on p as an average over examples (which is how they arise in dual form in Equation (16) of the proof), and the loss function itself is also an average (Equation (12)). This makes the box constraint decoupling possible, and leads to the favorable dependence on p, making an efficient solution possible to a potentially flagrantly non-convex problem.

To summarize, the technique of optimizing only "halfway into" the dual allows us to readily manipulate the minimax problem exactly without using an approximation like weak duality, despite the lack of convexity in g. This technique was used implicitly for a different purpose in the "drifting game" analysis of boosting ([18], Sec. 13.4.1). Existing boosting work is loosely related to our approach in being a transductive game invoked to analyze ensemble aggregation, but it does not consider unlabeled data and draws its power instead from being a repeated game ([18]).

The predecessor to this work ([1]) addresses a problem, 0-1 loss minimization, that is known to be NP-hard when solved directly ([11]). Using the unlabeled data is essential to surmounting this. It gives the dual problem an independently interesting interpretation, so the learning problem is on the always-convex Lagrange dual function and is therefore tractable.

This paper's transductive formulation involves no surrogates or relaxations of the loss, in sharp contrast to most previous work. This allows us to bypass the consistency and agnostic-learning discussions ([22, 3]) common to ERM methods that use convex risk minimization. Convergence analyses of those methods make heavy use of convexity of the losses and are generally done presupposing a linear weighting over \mathcal{H} ([21]), whereas here such structure emerges directly from Lagrange duality and involves no convexity to derive the worst-case-optimal predictions.

The conditions in Assumption 1 are notably general. Differentiability of the partial losses is convenient, but not necessary, and only used because first-order conditions are a convenient way to establish convexity of the potential well in Lemma 2. It is never used elsewhere, including in the minimax arguments used to prove Theorem 4. These manipulations are structured to be valid even if ℓ_{\pm} are non-monotonic; but in this case, g_j^* could turn out to be multi-valued and hence not a genuine function of the example's score.

We emphasize that our result on the minimax equilibrium (Theorem 4) holds for general losses – the slack function may not be convex unless the further conditions of Lemma 2 are met, but

the interpretation of the optimum in terms of margins and sigmoid functions remains. All this emerges from the inherent decision-theoretic structure of the problem (the proof of Appendix C.1). It manifests in the fact that the function $g(\mathbf{x}_j^\top \sigma)$ is always increasing in $\mathbf{x}_j^\top \sigma$ for general losses, because the function Γ is increasing. This monotonicity typically needs to be assumed in a generalized linear model (GLM; [15]) and related settings. Γ appears loosely analogous to the link function used by GLMs, with its inverse being used for prediction.

The optimal decision rules emerging from our framework are *artificial neurons* of the ensemble inputs. Helmbold et al. introduce the notion of a "matching loss" ([13]) for learning the parameters of a (fully supervised) artificial neuron with an arbitrary increasing transfer function, effectively taking the opposite tack of this paper in using a neuron's transfer function to derive a loss to minimize in order to learn the neuron's weights by convex optimization. Our assumptions on the loss, particularly condition (C) of Lemma 2, have arisen independently in earlier online learning work by some of the same authors ([12]); this may suggest connections between our techniques. We also note that our family of general losses was defined independently by [19] in the ERM setting (dubbed "F-losses") – in which condition (C) of Lemma 2 also has significance ([19], Prop. 2) – but has seemingly not been revisited thereafter. Further fleshing out the above connections would be interesting future work.

3 Extensions

We detail a number of generalizations to the basic prediction scenario of Sec. 2. These extensions are not mutually exclusive, and apply in conjunction with each other, but we list them separately for clarity. They illustrate the versatility of our minimax framework, particularly Sec. 3.4.

3.1 Weighted Test Sets and Covariate Shift

Though our results here deal with binary classification of an unweighted test set, the formulation deals with a weighted set with only a slight modification of the slack function:

Theorem 6. For any vector $\mathbf{r} > \mathbf{0}^n$,

$$\min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \sum_{j=1}^n r_j \ell(z_j, g_j) = \frac{1}{2} \min_{\sigma \geq \mathbf{0}^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \sum_{j=1}^n r_j \Psi\left(\frac{\mathbf{x}_j^\top \sigma}{r_j}\right) \right]$$

Writing $\sigma_{\mathbf{r}}^*$ as the minimizer of the RHS above, the optimal predictions $\mathbf{g}^* = \mathbf{g}(\sigma_{\mathbf{r}}^*)$, as in Theorem 4.

Such weighted classification can be parlayed into algorithms for general supervised learning problems via learning reductions ([4]). Allowing weights on the test set for the evaluation is tantamount to accounting for known covariate shift in our setting; it would be interesting, though outside our scope, to investigate scenarios with more uncertainty.

3.2 General Loss Constraints on the Ensemble

So far in the paper, we have considered the constraints on ensemble classifiers as derived from their label correlations (i.e. 0-1 losses), as in (3). However, this view can be extended significantly with the same analysis, because any general loss $\ell(z,g)$ is linear in z (Eq. (2)), which allows our development to go through essentially intact.

In summary, a classifier can be incorporated into our framework for aggregation if we have a generalization loss bound on it, for any "general loss." This permits an enormous variety of constraint sets, as each classifier considered can have constraints corresponding to any number of loss bounds on it, even multiple loss bounds using different losses. For instance, h_1 can yield a constraint corresponding to a zero-one loss bound, h_2 can yield one constraint corresponding to a square loss bound and another corresponding to a zero-one loss bound, and so on. Appendix B details this idea, extending Theorem 4 to general loss constraints.

3.3 Uniform Convergence Bounds for b

In our basic setup, b has been taken as a lower bound on ensemble classifier label correlations. But as mentioned earlier, the error in estimating b is in fact often quantified by two-sided uniform

convergence (L_{∞}) bounds on b. Constraining z in this fashion amounts to L_1 regularization of the dual vector σ .

Proposition 7. For any $\epsilon \geq 0$,

$$\min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \left\|\frac{1}{n}\mathbf{F}\mathbf{z} - \mathbf{b}\right\|_{\infty} \le \epsilon}} \ell(\mathbf{z}, \mathbf{g}) = \min_{\sigma \in \mathbb{R}^p} \left[-\mathbf{b}^{\top} \sigma + \frac{1}{n} \sum_{j=1}^n \Psi(\mathbf{x}_j^{\top} \sigma) + \epsilon \|\sigma\|_1 \right]$$

As in Thm. 4, the optimal $\mathbf{g}^* = \mathbf{g}(\sigma_{\infty}^*)$, where σ_{∞}^* is the minimizer of the right-hand side above.

Here we optimize over all vectors σ (not just nonnegative ones) in an L_1 -regularized problem, convenient in practice because we can cross-validate over the parameter ϵ . To our knowledge, this particular analysis has been addressed in prior work only for the special case of the entropy loss on the probability simplex, discussed further in [8].

Prop. 7 is a corollary of a more general result using differently scaled label correlation deviations within the ensemble, i.e. $\left|\frac{1}{n}\mathbf{F}\mathbf{z} - \mathbf{b}\right| \leq \mathbf{c}$ for a general $\mathbf{c} \geq \mathbf{0}^n$. This turns out to be equivalent to regularizing the minimization over σ by its c-weighted L_1 norm $\mathbf{c}^{\top} |\sigma|$ (Thm. 11), used to penalize the ensemble nonuniformly ([7]). This situation is motivated by uniform convergence of heterogeneous ensembles comprised of e.g. "specialist" predictors, for which a union bound ([5]) results in $\left|\frac{1}{n}\mathbf{F}\mathbf{z} - \mathbf{b}\right|$ with varying coordinates. Such ensembles are discussed next.

3.4 Heterogenous Ensembles of Specialist Classifiers and Features

All the results and algorithms in this paper apply in full generality to ensembles of "specialist" classifiers that only predict on some subset of the test examples. This is done by merely calculating the constraints over only these examples, and changing **F** and **b** accordingly ([2]).

To summarize this from [2], suppose a classifier $i \in [p]$ decides to abstain on an example x_j $(j \in [n])$ with probability $1 - v_i(x)$, and otherwise predict $h_i(x)$. Our only assumption on $\{v_i(x_1), \ldots, v_i(x_n)\}$ is the reasonable one that $\sum_{j=1}^n v_i(x_j) > 0$, so classifier i is not a useless specialist that abstains everywhere.

The information about **z** contributed by classifier *i* is now not its overall correlation with **z** on the entire test set, but rather the correlation with **z** restricted to the test examples on which it predicts. On the test set, this is written as $\frac{1}{n}$ S**z**, where the matrix **S** is formed by reweighting each row of **F**:

$$\mathbf{S} := n \begin{pmatrix} \rho_1(x_1)h_1(x_1) & \rho_1(x_2)h_1(x_2) & \cdots & \rho_1(x_n)h_1(x_n) \\ \rho_2(x_1)h_2(x_1) & \rho_2(x_2)h_2(x_2) & \cdots & \rho_2(x_n)h_2(x_n) \\ \vdots & \vdots & \ddots & \vdots \\ \rho_p(x_1)h_p(x_1) & \rho_p(x_2)h_p(x_2) & \cdots & \rho_p(x_n)h_p(x_n) \end{pmatrix} , \quad \rho_i(x_j) := \frac{v_i(x_j)}{\sum_{k=1}^n v_i(x_k)}$$

(S = F when the entire ensemble consists of non-specialists, recovering our initial setup.) Therefore, the ensemble constraints (3) become $\frac{1}{n}Sz \ge b_S$, where b_S gives the label correlations of each classifier restricted to the examples on which it predicts. Though this rescaling results in entries of S having different ranges and magnitudes > 1, our results and proofs remain entirely intact.

Indeed, despite the title, this paper applies far more generally than to an ensemble of binary classifiers, because our proofs make no assumptions at all about the structure of ${\bf F}$. Each predictor in the ensemble can be thought of as a feature; it has so far been convenient to think of it as binary, following the perspective of binary classifier aggregation, but it could as well be e.g. real-valued, and the features can have very different scales (as in ${\bf S}$ above). An unlabeled example ${\bf x}$ is simply a vector of features, so arbitrarily abstaining specialists are equivalent to "missing features," which this framework handles seamlessly due to the given unlabeled data. Our development *applies generally to semi-supervised binary classification*.

Acknowledgements

AB is grateful to Chris "Ceej" Tosh for feedback that made the manuscript clearer. This work was supported by the NSF (grant IIS-1162581).

References

- [1] Akshay Balsubramani and Yoav Freund. Optimally combining classifiers using unlabeled data. In *Conference on Learning Theory*, pages 211–225, 2015.
- [2] Akshay Balsubramani and Yoav Freund. Scalable semi-supervised classifier aggregation. In Advances in Neural Information Processing Systems, pages 1351–1359, 2015.
- [3] Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. Journal of the American Statistical Association, 101(473):138–156, 2006.
- [4] Alina Beygelzimer, John Langford, and Bianca Zadrozny. Machine learning techniques reductions between prediction quality metrics. In *Performance Modeling and Engineering*, pages 3–28. Springer, 2008
- [5] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Occam's razor. Information processing letters, 24(6):377–380, 1987.
- [6] Andreas Buja, Werner Stuetzle, and Yi Shen. Loss functions for binary class probability estimation and classification: Structure and applications. 2005.
- [7] Emmanuel J Candes, Michael B Wakin, and Stephen P Boyd. Enhancing sparsity by reweighted 11 minimization. *Journal of Fourier analysis and applications*, 14(5-6):877–905, 2008.
- [8] Miroslav Dudik, Steven J Phillips, and Robert E Schapire. Performance guarantees for regularized maximum entropy density estimation. In *Learning Theory*, pages 472–486. Springer, 2004.
- [9] Yoav Freund. An adaptive version of the boost by majority algorithm. *Machine learning*, 43(3):293–318, 2001.
- [10] Andrew Guillory, Erick Chastain, and Jeff Bilmes. Active learning as non-convex optimization. *Journal of Machine Learning Research Proceedings Track*, 5:201–208, 2009.
- [11] Venkatesan Guruswami and Prasad Raghavendra. Hardness of learning halfspaces with noise. *SIAM Journal on Computing*, 39(2):742–765, 2009.
- [12] David Haussler, Jyrki Kivinen, and Manfred K Warmuth. Sequential prediction of individual sequences under general loss functions. *IEEE Transactions on Information Theory*, 44(5):1906–1925, 1998.
- [13] David P Helmbold, Jyrki Kivinen, and Manfred K Warmuth. Relative loss bounds for single neurons. *IEEE Transactions on Neural Networks*, 10(6):1291–1304, 1999.
- [14] Joseph Keshet and David A McAllester. Generalization bounds and consistency for latent structural probit and ramp loss. In Advances in Neural Information Processing Systems, pages 2205–2212, 2011.
- [15] Peter McCullagh and John A Nelder. Generalized linear models, volume 37. CRC press, 1989.
- [16] Mark D Reid and Robert C Williamson. Composite binary losses. *The Journal of Machine Learning Research*, 11:2387–2422, 2010.
- [17] Leonard J Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- [18] Robert E. Schapire and Yoav Freund. Boosting: Foundations and Algorithms. The MIT Press, 2012.
- [19] Yi Shen. Loss functions for binary classification and class probability estimation. PhD thesis, University of Pennsylvania, 2005.
- [20] Maurice Sion. On general minimax theorems. Pacific J. Math., 8(1):171-176, 1958.
- [21] Matus Telgarsky, Miroslav Dudik, and Robert Schapire. Convex risk minimization and conditional probability estimation. In Conference on Learning Theory, 2015.
- [22] Tong Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, pages 56–85, 2004.

A Examples of Optimal Decision Rules for Various Losses

Loss	Partial Losses	$\Gamma(g)$	$\Psi(m)$	$g_i(\sigma)$
0-1	$\ell_{-}(g) = \frac{1}{2} (1+g)$	9	$\max(1, m)$	$\operatorname{clip}(\mathbf{x}_i^{ op}\sigma)$
	$\ell_+(g) = \frac{1}{2} (1-g)$			
Log	$\ell_{-}(g) = \ln\left(\frac{2}{1-g}\right)$	$\ln\left(\frac{1+g}{1-g}\right)$	$\ln(1+e^m) + \ln(1+e^{-m})$	$ \begin{array}{c c} 1 & -c & -x \\ -x & -c & -x \\ -x & -c & -x \end{array} $
	$\ell_+(g) = \ln\left(\frac{2}{1+g}\right)$			$1+e^{-it}$
Square	$\ell_{-}(g) = \left(\frac{1+g}{2}\right)^2$	В	$\begin{cases} -m & m \le -1\\ \frac{1}{1}(m^2 \pm 1) & m \in (-1, 1) \end{cases}$	$\operatorname{clip}(\mathbf{x}_i^{ op}\sigma)$
	$\ell_+(g) = \left(\frac{1-g}{2}\right)^2$			
CW (param. c)	$\ell_{-}(g) = c\left(1+g\right)$	g + 2c - 1		$\operatorname{clip}(\mathbf{x}_i^{ op}\sigma+1-2c)$
	$\ell_{+}(g) = (1-c)(1-g)$		$\begin{cases} (2c-1)m + 4c(1-c) & m \in (2c-2,2c) \\ m & m \ge 2c \end{cases}$	
Exponential	$\ell_{-}(g) = e^g$	$e^{g} - e^{-g}$	l u	$\operatorname{clip}\left(\ln\left(\frac{1}{2}\mathbf{x}_i^\top\sigma+\sqrt{1+\frac{1}{4}\left(\mathbf{x}_i^\top\sigma\right)^2}\right)\right)$
	$\ell_+(g) = e^{-g}$		$\begin{cases} \sqrt{4+m^2} & m \in (-e+\frac{1}{e}, e-\frac{1}{e}) \\ m+2/e & m \geq e-\frac{1}{e} \end{cases}$	
Logistic	$\ell_{-}(g) = \ln{(1 + e^g)}$	9		$\operatorname{clip}(\mathbf{x}_i^{ op}\sigma)$
	$\ell_{+}(g) = \ln\left(1 + e^{-g}\right)$		$\begin{cases} \ln(1 + e^{-r}) + \ln(1 + e^{-rr}) & m \in (-1, 1) \\ m + 2 \ln(1 + 1/e) & m \ge 1 \end{cases}$	
Hellinger	$\ell(g) = 1 - \sqrt{\frac{1-g}{2}}$	$\sqrt{\frac{1+g}{2}} - \sqrt{\frac{1-g}{2}}$	$\begin{bmatrix} -m \\ \hline \end{bmatrix}$ $m \le -1$	$\left\{ (\mathbf{x}_i^\top \sigma) \sqrt{2 - (\mathbf{x}_i^\top \sigma)^2} \mathbf{x}_i^\top \sigma \leq 1 \right.$
	$\ell_+(g) = 1 - \sqrt{\frac{1+g}{2}}$		$\begin{cases} 2 - \sqrt{\frac{1 - m\sqrt{2 - m^2}}{2}} - \sqrt{\frac{1 + m\sqrt{2 - m^2}}{2}} & m \in (-1, 1) \\ m & m \ge 1 \end{cases}$	
"AdaBoost"	$\ell_{-}(g) = \sqrt{\frac{1+g}{1-g}}$	$\frac{2g}{\sqrt{1-g^2}}$	$\sqrt{\frac{\sqrt{m^2+4+m}}{\sqrt{m^2+4-m}}} + \sqrt{\frac{\sqrt{m^2+4-m}}{\sqrt{m^2+4+m}}}$	$\frac{\mathbf{x}_{i}^{\top} \sigma}{\sqrt{(\mathbf{x}^{\top} \sigma)^{2}+4}}$
	$\ell_+(g) = \sqrt{\frac{1-g}{1+g}}$			
Sigmoid	$\ell(g) = \frac{1}{1+e^{-g}}$	$\frac{e^{9}-1}{e^{9}+1}$	$\max\left(1, m +\frac{2}{1+\varepsilon}\right)$	$\operatorname{clip}\left(\ln\left(rac{1+m}{1-m} ight) ight)$
	$\ell_+(g) = \frac{1}{1 + eg}$			

Table 1: Some binary classification losses, as in Sec. 2.3. For convenience, we write $\operatorname{clip}(x) = \min(1, \max(-1, x))$.

B Constraints on General Losses for Binary Classification

In all other sections of the paper, we allow the evaluation function of the game to be a general loss, but assume that the constraints (our information about the ensemble) are in terms of zero-one loss as written in (3). However, here we relax that assumption, allowing each classifier h_i to constrain the test labels \mathbf{z} not with the zero-one loss of h_i 's predictions, but rather with some other general loss.

This is possible because when the true labels are binary, all the losses we consider are linear in z,

as seen in (5):
$$\ell(\mathbf{z}, \mathbf{g}) = \frac{1}{n} \sum_{j=1}^{n} \frac{1}{2} \left[\ell_{+}(g_{j}) + \ell_{-}(g_{j}) \right] - \frac{1}{2n} \mathbf{z}^{\top} [\Gamma(\mathbf{g})]$$
. Accordingly, recall that $\mathbf{h}_{i} \in \mathbb{R}$

 $[-1,1]^n$ is the vector of test predictions of classifier h_i . Suppose we have an upper bound on the generalization loss of h_i , i.e. $\ell(\mathbf{z}, \mathbf{h}_i) \leq \epsilon_i^{\ell}$. If we define $b_i^{\ell} := \frac{1}{n} \sum_{j=1}^n \left[\ell_+(h_i(x_j)) + \ell_-(h_i(x_j)) \right] - 2\epsilon_i^{\ell}$, then we can use the definition of $\ell(\mathbf{z}, \mathbf{g})$ to write

$$\ell(\mathbf{z}, \mathbf{h}_i) \le \epsilon_i^{\ell} \qquad \Longleftrightarrow \qquad \frac{1}{n} \mathbf{z}^{\top} [\Gamma(\mathbf{h}_i)] \ge b_i^{\ell}$$
 (9)

Now (9) is a linear constraint on **z**, just like each of the error constraints earlier considered in (3). We can derive an aggregation algorithm with constraints like (9), using essentially the same analysis as employed in Section 2 to solve the game (4). As mentioned in Sec. 3.2,

Matching Objective and Constraint Losses

Though the ensemble constraints can be completely heterogeneous, we focus on a special case of them in the rest of this section to glean intuition. Suppose when each classifier contributes exactly one constraint to the problem, and the losses used for these constraints are all the same as each other and as the loss ℓ used in the objective function. In other words, the minimax prediction problem we now consider is

$$V^{\ell} := \min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \forall i \in [p]: \ \ell(\mathbf{z}, \mathbf{h}_i) \le \epsilon_i^{\ell}}} \ell(\mathbf{z}, \mathbf{g}) = \min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \forall i \in [p]: \ \frac{1}{n} \mathbf{z}^{\top} [\Gamma(\mathbf{h}_i)] \ge b_i^{\ell}}} \ell(\mathbf{z}, \mathbf{g})$$
(10)

The matrix ${\bf F}$ and the slack function from (1) are therefore redefined:

$$F_{ij}^{\ell} := \Gamma(h_i(x_j)) = \ell_-(h_i(x_j)) - \ell_+(h_i(x_j))$$

$$\gamma^{\ell}(\sigma, \mathbf{b}^{\ell}) := \gamma^{\ell}(\sigma) := -[\mathbf{b}^{\ell}]^{\top} \sigma + \frac{1}{n} \sum_{j=1}^{n} \Psi\left([\Gamma(\mathbf{x}_{j})]^{\top} \sigma \right)$$

where $\mathbf{b}^\ell = (b_1^\ell, \dots, b_p^\ell)^\top$ and the vectors \mathbf{x}_j are now from the new unlabeled data matrix F_{ij}^ℓ . The game (10) is clearly of the same form as the earlier formulation (4). Therefore, its solution has the same structure as in Theorem 4, proved using that theorem's proof:

Theorem 8. The minimax value of the game (10) is $V := \frac{1}{2}\gamma^{\ell}(\sigma^{\ell*}) := \min_{\sigma \geq \mathbf{0}^p} \frac{1}{2}\gamma^{\ell}(\sigma)$. The minimax optimal predictions are defined as follows: for all $j \in [n]$,

$$g_j^* := g_j(\sigma^*) = \begin{cases} -1 & \text{if } [\Gamma(\mathbf{x}_j)]^\top \sigma^{\ell *} \leq \Gamma(-1) \\ \Gamma^{-1} \left([\Gamma(\mathbf{x}_j)]^\top \sigma^{\ell *} \right) & \text{if } [\Gamma(\mathbf{x}_j)]^\top \sigma^{\ell *} \in (\Gamma(-1), \Gamma(1)) \\ 1 & \text{if } [\Gamma(\mathbf{x}_j)]^\top \sigma^{\ell *} \geq \Gamma(1) \end{cases}$$

This provides a concise characterization of how to solve the semi-supervised binary classification game for general losses. Though on the face of it Theorem 8 is a much stronger result than even Theorem 4, we cannot overlook statistical issues. The loss bounds ϵ_i^ℓ on each classifier are estimated using a uniform convergence bound over the ensemble with loss ℓ , but the data now considered are not the ensemble predictions, but the predictions passed through function Γ . This can be impractical for losses like log loss, for which Γ is unbounded, and therefore uniform convergence to estimate b_i^ℓ in (9) is much less applicable than for 0-1 loss.

But such issues are outside our scope here, and our constrained minimax results hold in any case given b. They may be useful to obtain semi-supervised learnability results for different losses from tighter statistical characterizations, which we consider an interesting open problem.

Proofs and Supporting Results

Proof of Theorem 4

The main hurdle here is the constrained maximization over z. For this we use the following result, a basic application of Lagrange duality (from [1], but proved below for completeness).

Lemma 9. For any $\mathbf{a} \in \mathbb{R}^n$,

$$\max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \mathbf{z}^{\top} \mathbf{a} = \min_{\sigma \geq \mathbf{0}^p} \left[-\mathbf{b}^{\top} \sigma + \frac{1}{n} \left\| \mathbf{F}^{\top} \sigma + \mathbf{a} \right\|_1 \right]$$

With this lemma, we prove the main theorem of this paper.

Proof of Theorem 4. First note that $\ell(\mathbf{z}, \mathbf{g})$ is linear in \mathbf{z} ,

$$V = (5) = \frac{1}{2} \min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n} \mathbf{F} \mathbf{z} > \mathbf{b}}} \frac{1}{n} \sum_{j=1}^n \left[\ell_+(g_j) + \ell_-(g_j) + z_j \left(\ell_+(g_j) - \ell_-(g_j) \right) \right]$$

Here we can rewrite the constrained maximization over z using Lemma 9:

$$\max_{\mathbf{z} \in [-1,1]^n, \frac{1}{n} \mathbf{F}_{\mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \sum_{i=1}^n z_j \left(\ell_+(g_j) - \ell_-(g_j) \right) = \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n} \mathbf{F}_{\mathbf{z} \geq \mathbf{b}}}} -\frac{1}{n} \mathbf{z}^\top [\Gamma(\mathbf{g})]$$

$$= \min_{\sigma \geq \mathbf{0}^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \left\| \mathbf{F}^\top \sigma - \Gamma(\mathbf{g}) \right\|_1 \right]$$
(11)

Substituting (11) into (5) and simplifying,

$$V = \frac{1}{2} \min_{\mathbf{g} \in [-1,1]^{n}} \left[\frac{1}{n} \sum_{j=1}^{n} \left[\ell_{+}(g_{j}) + \ell_{-}(g_{j}) \right] + \max_{\mathbf{z} \in [-1,1]^{n}, \frac{1}{n} \sum_{j=1}^{n} z_{j} \left(\ell_{+}(g_{j}) - \ell_{-}(g_{j}) \right) \right]$$

$$= \frac{1}{2} \min_{\mathbf{g} \in [-1,1]^{n}} \left[\frac{1}{n} \sum_{j=1}^{n} \left[\ell_{+}(g_{j}) + \ell_{-}(g_{j}) \right] + \min_{\sigma \geq \mathbf{0}^{p}} \left[-\mathbf{b}^{\top} \sigma + \frac{1}{n} \left\| \mathbf{F}^{\top} \sigma - \Gamma(\mathbf{g}) \right\|_{1} \right] \right]$$

$$= \frac{1}{2} \min_{\sigma \geq \mathbf{0}^{p}} \left[-\mathbf{b}^{\top} \sigma + \min_{\mathbf{g} \in [-1,1]^{n}} \left[\frac{1}{n} \sum_{j=1}^{n} \left[\ell_{+}(g_{j}) + \ell_{-}(g_{j}) \right] + \frac{1}{n} \left\| \mathbf{F}^{\top} \sigma - \Gamma(\mathbf{g}) \right\|_{1} \right] \right]$$

$$= \frac{1}{2} \min_{\sigma \geq \mathbf{0}^{p}} \left[-\mathbf{b}^{\top} \sigma + \frac{1}{n} \sum_{j=1}^{n} \min_{g_{j} \in [-1,1]} \left[\ell_{+}(g_{j}) + \ell_{-}(g_{j}) + \left| \mathbf{x}_{j}^{\top} \sigma - \Gamma(g_{j}) \right| \right]$$

$$(13)$$

The absolute value breaks down into two cases, so the inner minimization's objective can be simplified:

$$\ell_{+}(g_{j}) + \ell_{-}(g_{j}) + \left| \mathbf{x}_{j}^{\top} \sigma - \Gamma(g_{j}) \right| = \begin{cases} 2\ell_{+}(g_{j}) + \mathbf{x}_{j}^{\top} \sigma & \text{if } \mathbf{x}_{j}^{\top} \sigma \geq \Gamma(g_{j}) \\ 2\ell_{-}(g_{j}) - \mathbf{x}_{j}^{\top} \sigma & \text{if } \mathbf{x}_{j}^{\top} \sigma < \Gamma(g_{j}) \end{cases}$$
(15)

Suppose g_j falls in the first case, so that $\mathbf{x}_j^{\top} \sigma \geq \Gamma(g_j)$. From Assumption 1, $2\ell_+(g_j) + \mathbf{x}_j^{\top} \sigma$ is decreasing in g_j , so it is minimized for the greatest $g_i^* \leq 1$ s.t. $\Gamma(g_i^*) \leq \mathbf{x}_i^\top \sigma$. Since $\Gamma(\cdot)$ is increasing, exactly one of two subcases holds:

1.
$$g_j^*$$
 is such that $\Gamma(g_j^*) = \mathbf{x}_j^{\top} \sigma$, in which case the minimand (15) is $\ell_+(g_j^*) + \ell_-(g_j^*)$

1.
$$g_j^*$$
 is such that $\Gamma(g_j^*) = \mathbf{x}_j^{\top} \sigma$, in which case the minimand (15) is $\ell_+(g_j^*) + \ell_-(g_j^*)$
2. $g_j^* = 1$ so that $\Gamma(g_j^*) = \Gamma(1) < \mathbf{x}_j^{\top} \sigma$, in which case the minimand (15) is $2\ell_+(1) + \mathbf{x}_j^{\top} \sigma$

A precisely analogous argument holds if g_j falls in the second case, where $\mathbf{x}_j^{\top} \sigma < \Gamma(g_j)$. Putting the cases together, we have shown the form of the summand Ψ , piecewise over its domain, so (14) is equal to $\frac{1}{2} \min_{\sigma \geq \mathbf{0}^p} [\gamma(\sigma)]$.

We have proved the dependence of g_j^* on $\mathbf{x}_j^{\top} \sigma^*$, where σ^* is the minimizer of the outer minimization of (14). This completes the proof.

Proof of Lemma 9. We have

$$\max_{\mathbf{z} \in [-1,1]^n, \ \mathbf{z}} \frac{1}{n} \mathbf{z}^{\top} \mathbf{a} = \frac{1}{n} \max_{\mathbf{z} \in [-1,1]^n} \min_{\sigma \ge \mathbf{0}^p} \left[\mathbf{z}^{\top} \mathbf{a} + \sigma^{\top} (\mathbf{F} \mathbf{z} - n \mathbf{b}) \right]
\stackrel{(a)}{=} \frac{1}{n} \min_{\sigma \ge \mathbf{0}^p} \max_{\mathbf{z} \in [-1,1]^n} \left[\mathbf{z}^{\top} (\mathbf{a} + \mathbf{F}^{\top} \sigma) - n \mathbf{b}^{\top} \sigma \right]
= \frac{1}{n} \min_{\sigma \ge \mathbf{0}^p} \left[\left\| \mathbf{a} + \mathbf{F}^{\top} \sigma \right\|_1 - n \mathbf{b}^{\top} \sigma \right] = \min_{\sigma \ge \mathbf{0}^p} \left[-\mathbf{b}^{\top} \sigma + \frac{1}{n} \left\| \mathbf{F}^{\top} \sigma + \mathbf{a} \right\|_1 \right]$$
(17)

where (a) is by the minimax theorem ([20]).

C.2 Other Proofs

Lemma 10. The function $\ell_+(\Gamma^{-1}(m)) + \ell_-(\Gamma^{-1}(m))$ is convex for $m \in (\Gamma(-1), \Gamma(1))$ under any of the conditions of Lemma 2.

Proof of Lemma 10. Define $F(m) = \ell_+(\Gamma^{-1}(m)) + \ell_-(\Gamma^{-1}(m))$. By basic properties of the derivative,

$$\frac{d\left[\Gamma^{-1}\right]}{dm} = \frac{1}{\Gamma'(\Gamma^{-1}(m))} = \frac{1}{\ell'_{-}(\Gamma^{-1}(m)) - \ell'_{+}(\Gamma^{-1}(m))} \ge 0 \tag{18}$$

where the last inequality follows by Assumption 1. Therefore, by the chain rule and (18),

$$F'(m) = \frac{\ell'_{-}(\Gamma^{-1}(m)) + \ell'_{+}(\Gamma^{-1}(m))}{\ell'_{-}(\Gamma^{-1}(m)) - \ell'_{+}(\Gamma^{-1}(m))}$$
(19)

From this, we calculate F''(m), writing $\ell'_{\pm}(\Gamma^{-1}(m))$ and $\ell''_{\pm}(\Gamma^{-1}(m))$ as simply ℓ'_{\pm} and ℓ''_{\pm} for clarity:

$$F''(m) = \underbrace{\frac{\frac{d[\Gamma^{-1}]}{dm}}{\left(\ell'_{-}(\Gamma^{-1}(m)) - \ell'_{+}(\Gamma^{-1}(m))\right)^{2}}}_{(a)} \left[\left(\ell'_{-} - \ell'_{+}\right) \left(\ell''_{-} + \ell''_{+}\right) - \left(\ell'_{-} + \ell'_{+}\right) \left(\ell''_{-} - \ell''_{+}\right) \right]$$

From (18), observe that the term $(a) = (\ell'_-(\Gamma^{-1}(m)) - \ell'_+(\Gamma^{-1}(m)))^{-3} \ge 0$. Therefore, it suffices to show that the other term is ≥ 0 . But this is equal to

$$(\ell'_{-} - \ell'_{+}) (\ell''_{-} + \ell''_{+}) - (\ell'_{-} + \ell'_{+}) (\ell''_{-} - \ell''_{+}) = 2(\ell'_{-} \ell''_{+} - \ell''_{-} \ell'_{+})$$
(20)

This proves that condition (C) of Lemma 2 is sufficient for convexity of F (and necessary also, under Assumption 1 on the partial losses).

We now address the other conditions of Lemma 2. (A) implies (C), because by Assumption 1, $\ell'_-, \ell''_+, \ell''_-$ are ≥ 0 and $\ell'_+ \leq 0$, so (20) is ≥ 0 as desired.

Finally we prove that (B) implies (C). If ℓ is proper, then it is well known (e.g. Thm. 1 of [16], and [6]) that for all $x \in (-1,1)$,

$$\frac{\ell'_{-}(x)}{1+x} = -\frac{\ell'_{+}(x)}{1-x}$$

(This is a simple and direct consequence of stationary conditions from the properness definition.)

Define the function $w(x) = \frac{\ell'_-(x)}{1+x} = -\frac{\ell'_+(x)}{1-x}$; we drop the argument and write it and its derivative as w and w' for clarity. By direct computation,

$$\begin{split} \ell'_-\ell''_+ - \ell''_-\ell'_+ &= \left[(1+x)w \left(w + (x-1)w' \right) \right] - \left[(w + (1+x)w')(x-1)w \right] \\ &= \left[(1+x)w^2 + (x^2-1)ww' \right] - \left[(x-1)w^2 + (x^2-1)ww' \right] = 2w^2 \geq 0 \\ \text{so (C) is true as desired.} \end{split}$$

Proof of Lemma 2. Continuity follows by checking $\Psi(m)$ at $m=\pm 1$.

For Lipschitzness, note that for $m \in (\Gamma(-1), \Gamma(1))$, by (19),

$$\Psi'(m) = \frac{\ell'_{-}(\Gamma^{-1}(m)) + \ell'_{+}(\Gamma^{-1}(m))}{\ell'_{-}(\Gamma^{-1}(m)) - \ell'_{+}(\Gamma^{-1}(m))}$$
(21)

$$= -1 + \frac{2\ell'_{-}(\Gamma^{-1}(m))}{\ell'_{-}(\Gamma^{-1}(m)) - \ell'_{+}(\Gamma^{-1}(m))}$$
 (22)

$$=1-\frac{2(-\ell'_{+}(\Gamma^{-1}(m)))}{\ell'_{-}(\Gamma^{-1}(m))-\ell'_{+}(\Gamma^{-1}(m))}$$
(23)

Using Assumption 1 on the partial losses, equations (22) and (23) respectively make clear that $\Psi'(m) \geq -1$ and $\Psi'(m) \leq 1$ on this interval. Since $\Psi'(m)$ is -1 for $m < \Gamma(-1)$ and 1 for $m > \Gamma(1)$, it is 1-Lipschitz.

As for convexity, since Ψ is linear outside the interval $(\Gamma(-1), \Gamma(1))$, it suffices to show that $\Psi(m)$ is convex inside this interval, which is shown in Lemma 10.

C.3 Results and Proofs from Section 3

Proof of Theorem 6. The proof is similar to that of Theorem 4, which it generalizes. First note that

$$\max_{\mathbf{z} \in [-1,1]^n, \frac{1}{n} \mathbf{F}_{\mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \sum_{i=1}^n r_j z_j \left(\ell_+(g_j) - \ell_-(g_j) \right) = \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n} \mathbf{F}_{\mathbf{z} \geq \mathbf{b}}}} -\frac{1}{n} \mathbf{z}^\top [\mathbf{r} \circ \Gamma(\mathbf{g})]$$

$$= \min_{\sigma \geq \mathbf{0}^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \left\| \mathbf{F}^\top \sigma - (\mathbf{r} \circ \Gamma(\mathbf{g})) \right\|_1 \right]$$
(24)

where the last equality uses Lemma 9.

Therefore, using (24) on the left-hand side of what we wish to prove,

$$V = \frac{1}{2} \min_{\mathbf{g} \in [-1,1]^{n}} \left[\frac{1}{n} \sum_{j=1}^{n} r_{j} \left[\ell_{+}(g_{j}) + \ell_{-}(g_{j}) \right] + \max_{\substack{\mathbf{z} \in [-1,1]^{n} \\ \frac{1}{n} \mathbf{F} \mathbf{z} \geq \mathbf{b}}}, \frac{1}{n} \sum_{i=1}^{n} r_{j} z_{j} \left(\ell_{+}(g_{j}) - \ell_{-}(g_{j}) \right) \right]$$

$$= \frac{1}{2} \min_{\mathbf{g} \in [-1,1]^{n}} \left[\frac{1}{n} \sum_{j=1}^{n} r_{j} \left[\ell_{+}(g_{j}) + \ell_{-}(g_{j}) \right] + \min_{\sigma \geq \mathbf{0}^{p}} \left[-\mathbf{b}^{\top} \sigma + \frac{1}{n} \sum_{j=1}^{n} \left| \mathbf{x}_{j}^{\top} \sigma - r_{j} \Gamma(g_{j}) \right| \right] \right]$$

$$= \frac{1}{2} \min_{\sigma \geq \mathbf{0}^{p}} \left[-\mathbf{b}^{\top} \sigma + \frac{1}{n} \sum_{j=1}^{n} \min_{g_{j} \in [-1,1]} \left(r_{j} \left[\ell_{+}(g_{j}) + \ell_{-}(g_{j}) \right] + \left| \mathbf{x}_{j}^{\top} \sigma - r_{j} \Gamma(g_{j}) \right| \right) \right]$$
(25)

As in the proof of Theorem 4, the inner minimization's objective can be simplified:

$$r_{j}(\ell_{+}(g_{j}) + \ell_{-}(g_{j})) + \left|\mathbf{x}_{j}^{\top}\sigma - r_{j}\Gamma(g_{j})\right| = \begin{cases} 2r_{j}\ell_{+}(g_{j}) + \mathbf{x}_{j}^{\top}\sigma & \text{if } \mathbf{x}_{j}^{\top}\sigma \geq r_{j}\Gamma(g_{j}) \\ 2r_{j}\ell_{-}(g_{j}) - \mathbf{x}_{j}^{\top}\sigma & \text{if } \mathbf{x}_{j}^{\top}\sigma < r_{j}\Gamma(g_{j}) \end{cases}$$
(26)

Suppose g_j falls in the first case, so that $\mathbf{x}_j^\top \sigma \geq r_j \Gamma(g_j)$. From Assumption 1, $2r_j \ell_+(g_j) + \mathbf{x}_j^\top \sigma$ is decreasing in g_j , so it is minimized for the greatest $g_j^* \leq 1$ s.t. $\Gamma(g_j^*) \leq \frac{\mathbf{x}_j^\top \sigma}{r_j}$. Since $\Gamma(\cdot)$ is increasing, exactly one of two subcases holds:

a)
$$g_j^*$$
 is such that $\Gamma(g_j^*) = \frac{\mathbf{x}_j^\top \sigma}{r_j}$, in which case the minimand (26) is $r_j(\ell_+(g_j^*) + \ell_-(g_j^*))$

b)
$$g_j^*=1$$
 so that $\Gamma(g_j^*)=\Gamma(1)<\frac{\mathbf{x}_j^{\intercal}\sigma}{r_j}$, in which case the minimand (26) is $2r_j\ell_+(1)+\mathbf{x}_j^{\intercal}\sigma$

A precisely analogous argument holds if g_j falls in the second case, where $\mathbf{x}_j^{\top} \sigma < \Gamma(g_j)$. So as before, we have proved the dependence of g_j^* on $\mathbf{x}_j^{\top} \sigma^*$, where σ^* is the minimizer of the outer minimization of (25). This completes the proof.

Theorem 11. For any $c \ge 0^n$,

$$\min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ |\frac{1}{n}\mathbf{F}\mathbf{z} - \mathbf{b}| \leq \mathbf{c}}} \ell(\mathbf{z}, \mathbf{g}) = \min_{\sigma \in \mathbb{R}^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \sum_{j=1}^n \Psi(\mathbf{x}_j^\top \sigma) + \mathbf{c}^\top |\sigma| \right]$$

Writing $\sigma_{\mathbf{c}}^*$ as the minimizer of the RHS above, the optimal predictions $\mathbf{g}^* = \mathbf{g}(\sigma_{\mathbf{c}}^*)$, as in Theorem 4.

Proof of Theorem 11. This is proved exactly like Theorem 4, but using Lemma 12 instead of Lemma 9 in that proof.

Lemma 12. For any $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{c} \geq \mathbf{0}^n$,

$$\max_{\mathbf{z} \in [-1,1]^n, \quad n} \frac{1}{n} \mathbf{z}^\top \mathbf{a} = \min_{\sigma \in \mathbb{R}^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \left\| \mathbf{F}^\top \sigma + \mathbf{a} \right\|_1 + \mathbf{c}^\top |\sigma| \right]$$

Proof.

$$\max_{\mathbf{z} \in [-1,1]^n, \atop \frac{1}{n}} \frac{1}{n} \mathbf{z}^{\top} \mathbf{a} = \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n} \mathbf{F} \mathbf{z} - \mathbf{b} \leq \mathbf{c}, \\ -\frac{1}{n} \mathbf{F} \mathbf{z} + \mathbf{b} \leq \mathbf{c}}} \\
= \frac{1}{n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n} \mathbf{F} \mathbf{z} + \mathbf{b} \leq \mathbf{c}}} \left[\mathbf{z}^{\top} \mathbf{a} + \lambda^{\top} (-\mathbf{F} \mathbf{z} + n \mathbf{b} + n \mathbf{c}) + \xi^{\top} (\mathbf{F} \mathbf{z} - n \mathbf{b} + n \mathbf{c}) \right] \\
= \frac{1}{n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \mathbf{z} \in [-1,1]^n}} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \mathbf{z} \in [-1,1]^n}} \left[\mathbf{z}^{\top} (\mathbf{a} + \mathbf{F}^{\top} (\xi - \lambda)) + \lambda^{\top} (n \mathbf{b} + n \mathbf{c}) + \xi^{\top} (-n \mathbf{b} + n \mathbf{c}) \right] \\
= \frac{1}{n} \min_{\substack{\lambda, \xi \geq 0^p \\ \lambda}} \left[\| \mathbf{a} + \mathbf{F}^{\top} (\xi - \lambda) \|_{1} - n \mathbf{b}^{\top} (\xi - \lambda) + n \mathbf{c}^{\top} (\xi + \lambda) \right]$$

where the interchanging of min and max is again justified by the minimax theorem ([20]), since the objective is linear in each variable and one of the constraint sets is closed.

Suppose for some $j \in [n]$ that $\xi_j > 0$ and $\lambda_j > 0$. Then subtracting $\min(\xi_j, \lambda_j)$ from both does not affect the value $[\xi - \lambda]_j$, but always decreases $[\xi + \lambda]_j$, and therefore always decreases the objective function. Therefore, we can w.l.o.g. assume that $\forall j \in [n] : \min(\xi_j, \lambda_j) = 0$. Defining $\sigma_j = \xi_j - \lambda_j$ for all j (so that $\xi_j = [\sigma_j]_+$ and $\lambda_j = [\sigma_j]_-$), the last equality above becomes

$$\frac{1}{n} \min_{\lambda, \xi \ge \mathbf{0}^p} \left[\left\| \mathbf{a} + \mathbf{F}^{\top}(\xi - \lambda) \right\|_1 - n\mathbf{b}^{\top}(\xi - \lambda) + n\mathbf{c}^{\top}(\xi + \lambda) \right]$$
$$= \frac{1}{n} \min_{\sigma \in \mathbb{R}^p} \left[\left\| \mathbf{a} + \mathbf{F}^{\top}\sigma \right\|_1 - n\mathbf{b}^{\top}\sigma + n\mathbf{c}^{\top}|\sigma| \right]$$