

WHEN OPTIMIZING f -DIVERGENCE IS ROBUST WITH LABEL NOISE

Jiaheng Wei and Yang Liu*

Department of Computer Science and Engineering
University of California, Santa Cruz
Santa Cruz, CA 95060, USA
`{jiahengwei, yangliu}@ucsc.edu`

ABSTRACT

We show when maximizing a properly defined f -divergence measure with respect to a classifier’s predictions and the supervised labels is robust with label noise. Leveraging its variational form, we derive a nice decoupling property for a family of f -divergence measures when label noise presents, where the divergence is shown to be a linear combination of the variational difference defined on the clean distribution and a bias term introduced due to the noise. The above derivation helps us analyze the robustness of different f -divergence functions. With established robustness, this family of f -divergence functions arises as useful metrics for the problem of learning with noisy labels, which do not require the specification of the labels’ noise rate. When they are possibly not robust, we propose fixes to make them so. In addition to the analytical results, we present thorough experimental evidence. Our code is available at <https://github.com/UCSC-REAL/Robust-f-divergence-measures>.

1 INTRODUCTION

A machine learning system continuously observes noisy training annotations and it remains a challenge to perform robust training in such scenarios. Earlier and classical approaches rely on estimation processes to understand the noise rate of the labels and then leverage this knowledge to perform label correction (Patrini et al., 2017; Lukasik et al., 2020), or loss correction (Natarajan et al., 2013; Liu & Tao, 2015; Patrini et al., 2017), or both, among many other more carefully designed approaches (please refer to our related work section for more detailed coverage). Recent works have started to propose robust loss functions or metrics that do not require the above estimation (Charoenphakdee et al., 2019; Xu et al., 2019; Liu & Guo, 2020; Cheng et al., 2021). Clear advantages of the latter approaches include their easiness in implementation, as well as their robustness to noisy estimates of the parameters. This work mainly contributes to the second line of studies and aimed to propose relevant loss functions and measures that are inherently robust with label noise.

We start with formulating the problem of maximizing an f -divergence defined between a classifier’s prediction and the labels:

$$h_f^* = \operatorname{argmax}_h D_f(P_{h \times Y} || Q_{h \times Y}), \quad (1)$$

where in above D_f is an f -divergence function, P and Q are the joint and product (marginal) distribution of the classifier h ’s predictions on a feature space X and label Y . Though optimizing the f -divergence measure is in general not the same as finding the Bayes optimal classifiers, we show these measures encourage a classifier that maximizes an extended definition of f -mutual information between the classifier’s prediction and the true label distribution. We will also provide analysis for when the maximizer of this f -divergence coincides with the Bayes optimal classifier.

Building on a careful treatment of its variational form, we then reveal a nice property that helps establish the robustness of the f -divergence specified in Eqn. (1): the variational difference term defined with noisy labels is an affine transformation of the clean variational difference, subject to

*Corresponding author. `yangliu@ucsc.edu`.

an addition of a bias term. Using this result, we analyze under which conditions maximizing an f -divergence measure would be robust to label noise. In particular, we demonstrate strong robustness results for Total Variation divergence, identify conditions under which several other divergences, including Jenson-Shannon divergence and Pearson χ^2 divergence, are robust. The resultant f -divergence functions offer ways to learn with noisy labels, without estimating the noise parameters. As mentioned above, this distinguishes our solutions from a major line of previous studies that would require such estimates. When the f -divergence functions are possibly not robust with label noise, our analysis also offers a new way to perform “loss correction”. We’d like to emphasize that instead of offering one method/loss/measure, our results effectively offer a family of functions that can be used to perform this noisy training task. Our contributions summarize as follows:

- We show a certain set of f -divergence measures that are robust with label noise (some under certain conditions). The corresponding f -divergence functions provide the community with robust learning measures that do not require the knowledge of the noise rates.
- When the f -divergence measures are possibly not robust with label noise, our analysis provides ways to correct the f -divergence functions to offer robustness. This process would require the estimation of the noise rates and our results contribute new ways to leverage existing estimation techniques to make the training more robust.
- We empirically verified the effectiveness of optimizing f -divergences when noisy labels present. We opensource our solutions at <https://github.com/UCSC-REAL/Robust-f-divergence-measures>.

1.1 RELATED WORKS

The now most popular approach of dealing with label noise is to first estimate the noise transition matrix and then use this knowledge to perform loss or sample correction (Scott et al., 2013; Natarajan et al., 2013; Patrini et al., 2017; Lu et al., 2018; Han et al., 2018; Tanaka et al., 2018; Yao et al., 2020; Zhu et al., 2021). In particular, the surrogate loss (Scott et al., 2013; Natarajan et al., 2013; Scott, 2015; Van Rooyen et al., 2015; Menon et al., 2015) uses the transition matrix to define unbiased estimates of the true losses. Other works include (Sukhbaatar & Fergus, 2014; Xiao et al., 2015), which consider building a neural network to facilitate the learning of noise rates or noise transition matrix. Symmetric loss has been studied and conditions have been identified for when there is no need to estimate noise rate (Manwani & Sastry, 2013; Ghosh et al., 2015; 2017; Van Rooyen et al., 2015; Charoenphakdee et al., 2019). Nonetheless, it remains a challenge to develop training approaches without requiring knowing the noise rates for more generic settings.

More recently, (Zhang & Sabuncu, 2018; Amid et al., 2019) proposed robust losses for neural networks. When noise rates are asymmetric (label class-dependent), (Xu et al., 2019) proposed an information-theoretic loss that is also robust to asymmetric noise rates. There are also some trials on modifying the regularization term to improve generalization ability with the existence of label noise (Jenni & Favaro, 2018; Yi & Wu, 2019), and on providing complementary negative labels (Kim et al., 2019). Peer loss (Liu & Guo, 2020) is a recently proposed loss function that does not require knowing noise rates.

f -divergence is a popular information theoretical measure, and has been widely used and studied. Most relevant to us, f -GAN was proposed in (Nowozin et al., 2016) to study f -divergence in training generative neural samplers. To our best knowledge, ours is the first to study the robustness of f -divergence measures in the context of improving the robustness of training with noisy labels.

2 LEARNING WITH NOISY LABELS USING f -DIVERGENCE

Our solution ties to the definition of f -divergence. The f -divergence between two distributions P and Q with probability density function p and q being measures for $Z \in \mathcal{Z}^1$ is defined as:

$$D_f(P\|Q) = \int_{\mathcal{Z}} q(Z)f\left(\frac{p(Z)}{q(Z)}\right) dZ . \quad (2)$$

¹We use Z instead of X as conventionally done for a good reason - we will be reserving X to explicitly denote the features.

$f(\cdot)$ is a convex function such that $f(1) = 0$. Examples include KL-divergence when $f(v) = v \log v$ and Total Variation (TV) divergence with $f(v) = \frac{1}{2}|v - 1|$. Other examples can be found in Table 1. Following from Fenchel’s convex duality, f -divergence admits the following variational form:

$$D_f(P\|Q) = \sup_{g: \mathcal{Z} \rightarrow \text{dom}(f^*)} \mathbb{E}_{Z \sim P}[g(Z)] - \mathbb{E}_{Z \sim Q}[f^*(g(Z))],$$

where f^* is the Fenchel duality of the function $f(\cdot)$, which is defined as $f^*(u) = \sup_{v \in \mathbb{R}}\{uv - f(v)\}$. We use $\text{dom}(f^*)$ to denote the domain of f^* .

We consider the classification problem of learning a classifier $h : \mathcal{X} \rightarrow \mathcal{Y}$ that maps features $X \in \mathcal{X}$ to labels $Y \in \mathcal{Y} := \{1, 2, \dots, K\}$, where in above $X \times Y$ denote the random variables for features and labels. $X \times Y$ jointly draw from a distribution \mathcal{D} . For a clear presentation, we will often focus on presenting the binary classification setting $\mathcal{Y} = \{-1, +1\}$, but most of our core results extend to multi-class classification problems, and we shall provide corresponding justifications.

Instead of having access to sampled training data from $X \times Y$, we consider a setting with noisy labels where the noisy label \tilde{Y} generates according to a transition matrix T defined between \tilde{Y} and the true label Y . The (i, j) element of T is defined as $T_{i,j} = \mathbb{P}(\tilde{Y} = j | Y = i)$ where $i, j \in \{1, \dots, K\}$. For the ease of presentation, when we present for the binary case, we adopt the following notation: $e_+ := \mathbb{P}(\tilde{Y} = -1 | Y = +1)$, $e_- := \mathbb{P}(\tilde{Y} = +1 | Y = -1)$, $e_+ + e_- < 1$. Suppose we have access to a noisy training dataset $\{(x_n, \tilde{y}_n)\}_{n=1}^N$, where \tilde{y}_n generates according to \tilde{Y} .

2.1 LEARNING USING D_f

We will start with presenting our idea of training a classifier using D_f with the clean training data. Then we will proceed to the case with noisy labels. For an arbitrary classifier h , let’s denote by $P_{h \times Y}$ the joint distribution of $h(X)$ and Y :

$$\text{Joint distribution: } P_{h \times Y} := \mathbb{P}(h(X) = y, Y = y'), \quad y, y' \in \mathcal{Y}.$$

And we use $Q_{h \times Y}$ to denote the product (marginal) distribution of $h(X)$ and Y :

$$\text{Product distribution: } Q_{h \times Y} := \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = y'), \quad y, y' \in \mathcal{Y}.$$

When it is clear from context we will also shorthand the above two distributions as P and Q . We formulate the problem of learning using f -divergence as follows: the goal of the learner is to find a classifier h that maximizes the following divergence measure between P and Q :

$$\text{Learning using } D_f: h_f^* = \underset{h}{\operatorname{argmax}} D_f(P_{h \times Y} || Q_{h \times Y}) \tag{3}$$

Effectively the goal is to find a classifier that maximizes the divergence between the joint distribution and the product distribution. Define a f -mutual information based on f -divergence: $M_f(h(X); Y) = D_f(P_{h \times Y} || Q_{h \times Y})$, equivalently the maximization in Eqn. (3) tries to find the classifier that maximizes the f -mutual information between a classifier’s output distribution and the true label distribution. A notable example is when $f(v) = v \log v$, the corresponding D_f and M_f become the famous KL divergence and the mutual information. It is important to note in general maximizing (f -) mutual information between the classifier’s predictions and labels does not promise the Bayes optimal classifier $h^* = \underset{h}{\operatorname{argmax}} \mathbb{P}(h(X) = Y)$. Nonetheless, maximizing it often returns a quality one. We provide further analysis in Section 2.2.

Variational representation As we mentioned earlier, f -divergence admits a variational form which further allows us to focus on maximizing the following variational difference:

$$h_f^* = \underset{h}{\operatorname{argmax}} \sup_g \mathbb{E}_{Z \sim P_{h \times Y}}[g(Z)] - \mathbb{E}_{Z \sim Q_{h \times Y}}[f^*(g(Z))],$$

where we use Z to shorthand the tuple $[h(X), Y]$. Denote the variational difference as follows:

$$\text{VD}_f(h, g) := \mathbb{E}_{Z \sim P_{h \times Y}}[g(Z)] - \mathbb{E}_{Z \sim Q_{h \times Y}}[f^*(g(Z))]. \tag{4}$$

Let g^* be the corresponding optimal variational function g for $\text{VD}_f(h, g)$. This variational form allows us to use a training dataset $\{(x_n, y_n)\}_{n=1}^N$ to perform the above maximization problem listed in Eqn. (3) (Nowozin et al., 2016). A list of f -divergence functions together with the optimal variational/conjugate functions g/f^* is summarized in Table 1.

Name	$D_f(P Q)$	g^*	dom_{f^*}	$f^*(u)$
Total Variation	$\int \frac{1}{2} p(z) - q(z) dz$	$\frac{1}{2} \text{sign} \frac{p(z)}{q(z)} - 1$	$u \in [-\frac{1}{2}, \frac{1}{2}]$	u
Jenson-Shannon	$\frac{1}{2} \int p(z) \log \frac{2p(z)}{p(z) + q(z)} + q(z) \log \frac{2q(z)}{p(z) + q(z)} dz$	$\log \frac{p(z) + q(z)}{2p(z)}$	$u < \log 2$	$-\log(2 - e^u)$
Pearson χ^2	$\int \frac{(q(z) - p(z))^2}{p(z)} dz$	$2 \left(\frac{p(z)}{q(z)} - 1 \right)$	\mathbb{R}	$\frac{1}{4} u^2 + u$
KL	$\int p(z) \log \frac{p(z)}{q(z)} dx$	$1 + \log \frac{p(z)}{q(z)}$	\mathbb{R}	e^{u-1}

Table 1: D_f s, optimal variational g (g^*), conjugate functions (f^*). A more complete table, including Jeffrey, Squared Hellinger, Neyman χ^2 , Reverse KL, is provided in the Appendix.

2.2 HOW GOOD IS h_f^* ?

As we mentioned earlier, maximizing our defined f -divergence measures (or maximizing the f -mutual information) between the classifier’s predictions and labels is not always returning the Bayes optimal classifier. However, for a binary classification problem, we prove below that with balanced dataset, maximizing Total Variation (TV) divergence returns the Bayes optimal classifier:

Theorem 1. *For TV, when $\mathbb{P}(Y = +1) = \mathbb{P}(Y = -1)$ (balanced), h_f^* is the Bayes optimal classifier.*

Remark 2. *The above theorem extends to the multi-class setting when we restrict attentions to confident classifiers. See Appendix for details.*

The above observation is not easily true for other f -divergence. Nonetheless, denote by $Y^*(X = x)$ the Bayes optimal label for an instance x : $Y^*(X = x) = \text{argmax}_y \mathbb{P}(Y = y|X = x)$. Denote by $P_{h \times Y^*}, Q_{h \times Y^*}$ the joint and product distribution P, Q defined w.r.t. $h(X)$ and Y^* . We prove:

Theorem 3. *When $\mathbb{P}(Y^* = +1) = \mathbb{P}(Y^* = -1)$ (balanced), maximizing $D_f(P_{h \times Y^*}||Q_{h \times Y^*})$ returns the Bayes optimal classifier, if $f(v)$ is monotonically increasing in $|v - 1|$ on $\text{dom}(f)$.*

For example, Pearson χ^2 ($f(v) = (v - 1)^2$) satisfies the monotonicity condition. In practice, when the label distribution $\mathbb{P}(Y|X = x)$ has small uncertainties, the ground truth labels are approximately equivalent to the Bayes optimal label. Therefore, the above theorem implies that maximizing $D_f(P_{h \times Y}||Q_{h \times Y})$ is also likely to return a high-quality classifier for other f -divergences.

2.3 LEARNING WITH NOISY LABELS

Consider an arbitrary classifier h . Denote by $\tilde{P}_{h \times \tilde{Y}}$ the joint distribution of $h(X)$ and \tilde{Y} :

$$\text{Joint noisy distribution: } \tilde{P}_{h \times \tilde{Y}} := \mathbb{P}(h(X) = y, \tilde{Y} = y'), y, y' \in \mathcal{Y}.$$

Similarly, we use $\tilde{Q}_{h \times \tilde{Y}}$ to denote the product (marginal) distribution of $h(X)$ and \tilde{Y} :

$$\text{Product noisy distribution: } \tilde{Q}_{h \times \tilde{Y}} := \mathbb{P}(h(X) = y) \cdot \mathbb{P}(\tilde{Y} = y'), y, y' \in \mathcal{Y}.$$

When it is clear from context, we shorthand using \tilde{P}, \tilde{Q} . We are interested in understanding the robustness in maximizing $D_f(\tilde{P}_{h \times \tilde{Y}}||\tilde{Q}_{h \times \tilde{Y}})$. Using training samples $\{x_n, \tilde{y}_n\}_{n=1}^N$, there exists algorithms to compute the gradient of D_f leveraging its variational form (Nowozin et al., 2016), such that one can apply gradient descent or ascent to optimize it. We provide details in Section 5.

3 VARIATIONAL DIFFERENCE WITH NOISY LABELS

For an arbitrary g , we define the variational difference term w.r.t. the noisy label as follows:

$$\widetilde{\text{VD}}_f(h, g) := \mathbb{E}_{\tilde{Z} \sim \tilde{P}_{h \times \tilde{Y}}} [g(\tilde{Z})] - \mathbb{E}_{\tilde{Z} \sim \tilde{Q}_{h \times \tilde{Y}}} [f^*(g(\tilde{Z}))] \quad (5)$$

where we use \tilde{Z} to denote $[h(X), \tilde{Y}]$. Denote by \tilde{g}^* the corresponding optimal variational function g for $\widetilde{\text{VD}}_f(h, g)$. In this section, we show that the variational difference term under noisy labels is closely related to the variational difference term defined on the clean distributions P, Q . Define the following quantity: $\Delta_f^y(h, g) := \mathbb{E}_X [g(h(X), y)] - \mathbb{E}_X [f^*(g(h(X), y))]$, for example

$\Delta_f^{+1}(h, g) := \mathbb{E}_X[g(h(X), +1)] - \mathbb{E}_X[f^*(g(h(X), +1))]$. For a binary classification problem, further denote by $\text{Bias}_f(h, g) := e_+ \cdot \Delta_f^{-1}(h, g) + e_- \cdot \Delta_f^{+1}(h, g)$. We derive the following fact:

Theorem 4. *For binary classification, the variational difference between the noisy distributions \tilde{P} and \tilde{Q} relates to the one defined on the clean distributions in the following way:*

$$\tilde{VD}_f(h, g) = (1 - e_+ - e_-) VD_f(h, g) + \text{Bias}_f(h, g) \quad (6)$$

The above decoupling result is inspiring: $\text{Bias}_f(h, g)$ can be viewed as the additional bias term introduced by label noise. If this term has negligible effect in the maximization problem, maximizing the noisy variational difference term will be equivalent to maximizing $(1 - e_+ - e_-) \cdot VD_f(h, g)$, and therefore the clean variational difference term. If the above is true, we have established the robustness of the corresponding f -divergence. This result also points out that when the effects from the bias term are non-negligible, finding ways to counter the additional bias term will help us retain the robustness of D_f measures. Next we show that Theorem 4 extends to the multi-class setting under two broad families of noise rate models, both covering the binary setting as a special case.

Multi-class extension of Theorem 4: uniform off-diagonal case We first consider the following transition matrix: uniform off-diagonal transition matrix, where $e_j = T_{i,j}, \forall i \neq j$, that is any other classes $i \neq j$ has the same chance of being flipped to class j . The diagonal entry $T_{i,i}$ (chance of a correct label) becomes $1 - \sum_{j \neq i} e_j$. We further require that $\sum_j e_j < 1$. Note that the binary noise rate model is easily a uniform off-diagonal transition matrix.

Theorem 5. [Multi-class] *For uniform off-diagonal noise transition model, the noisy variational difference term relates to the clean one in the following way:*

$$\tilde{VD}_f(h, g) = (1 - \sum_{j=1}^K e_j) \cdot VD_f(h, g) + \sum_{j=1}^K e_j \cdot \Delta_f^j(h, g) \quad (7)$$

If we define $\text{Bias}_f(h, g) := \sum_{j=1}^K e_j \cdot \Delta_f^j(h, g)$, we reproduced the results in Theorem 4: for binary case, relabel class $1 \rightarrow +1, 2 \rightarrow -1$. Then $e_1 := \mathbb{P}(\tilde{Y} = +1 | Y = -1) = e_-, e_2 := \mathbb{P}(\tilde{Y} = -1 | Y = +1) = e_+$. Another case of noise model we consider is *sparse noise*. Mathematically, assume K is an even number, sparse noise model specifies $\frac{K}{2}$ disjoint pairs of classes (i_c, j_c) where $c \in [\frac{K}{2}]$ and $i_c < j_c$. The labels flip between each pair. We provide details in the Appendix.

4 WHEN D_f IS ROBUST WITH LABEL NOISE

Denote by \mathcal{H} an arbitrary hypothesis space for training a candidate classifier h . We will focus on \mathcal{H} throughout this section, and with abusing notation a bit, let $h_f^* = \operatorname{argmax}_{h \in \mathcal{H}} D_f(P_{h \times Y} || Q_{h \times Y})$. We first define formally what we mean by robustness of $D_f(P_{h \times Y} || Q_{h \times Y})$.

Definition 1. $D_f(P_{h \times Y} || Q_{h \times Y})$ is \mathcal{H} -robust if $h_f^* = \operatorname{argmax}_{h \in \mathcal{H}} D_f(\tilde{P}_{h \times \tilde{Y}} || \tilde{Q}_{h \times \tilde{Y}})$.

The above definition is stating that the label noise does not disrupt the optimality of h_f^* when maximizing $D_f(\tilde{P}_{h \times \tilde{Y}} || \tilde{Q}_{h \times \tilde{Y}})$ instead of $D_f(P_{h \times Y} || Q_{h \times Y})$.

4.1 IMPACT OF THE BIAS TERMS

In this section, we take a closer look at the Bias terms and argue that they have diminishing effects as compared to the VD terms when label noise increases. Recall g^*, \tilde{g}^* are the corresponding optimal variational functions for $VD_f(h, g)$ and $\tilde{VD}_f(h, g)$.

Total variation (TV) For TV, since $f(v) = \frac{1}{2}|v - 1|$, $f^*(u) = u$, we immediately have $\forall y' \ g(h = y', y) - f^*(g(h = y', y)) = 0$ and therefore $\Delta_f^y(h, g) = \mathbb{E}_X[g(h(X), y)] - \mathbb{E}_X[f^*(g(h(X), y))] = 0, \forall y$, and further $\text{Bias}_f(h, g) \equiv 0$. This fact helps establish the robustness of TV divergence measure (Theorem 7).

Other divergences The above nice property generally does not hold for other f -divergence functions. Next we focus on the binary classification setting and prove the following lemma:

Lemma 1. *For f -divergence listed in Table 6 (Appendix), $\text{Bias}_f(h, \tilde{g}^*) = O((1 - e_+ - e_-)^2)$.*

Note the variational form will be used when optimizing $D_f(\tilde{P}_{h \times \tilde{Y}} \parallel \tilde{Q}_{h \times \tilde{Y}})$ (and therefore we will be using \tilde{g}^*). This lemma simplifies Eqn. 4 to $\widetilde{\text{VD}}_f(h, g) \propto \text{VD}_f(h, g) + O(1 - e_+ - e_-)$. Since $0 < 1 - e_+ - e_- \leq 1$, when the noise rate $e_+ + e_-$ is high, the effect of **Bias** term diminishes. When the **Bias** term becomes negligible, we will have $\widetilde{\text{VD}}_f(h, g) \propto \text{VD}_f(h, g)$ if $e_+ + e_- \rightarrow 1$, establishing the fact that optimizing $\widetilde{\text{VD}}_f(h, g)$ is approximately the same as optimizing $\text{VD}_f(h, g)$.

4.2 HOW ROBUST ARE D_f S?

We first prove the following result:

Theorem 6. *D_f is \mathcal{H} -robust when $\text{Bias}_f(h, g)$ satisfies either of the following conditions: (I) $\forall h \in \mathcal{H}, \text{Bias}_f(h, g) \equiv \text{const.}$; (II) $\forall h \in \mathcal{H}, h \neq h_f^*, \text{Bias}_f(h, \tilde{g}^*) \leq \text{Bias}_f(h_f^*, g^*)$.*

Theorem 6 gives sufficient conditions when the **Bias** term does not get in the way of reaching the optimality h_f^* . Intuitively, when $\text{Bias}_f(h_f^*, g^*)$ is an upper bound of $\text{Bias}_f(h, \tilde{g}^*)$, the **Bias** term will not interfere with the convergence of the **VD** term. Next we provide specific examples of f -divergence functions that would satisfy these conditions.

Total Variation (TV) is robust For TV, the fact that $\Delta_f^y(h, g) \equiv 0$ allows us to prove:

Theorem 7. *For TV divergence, $\text{Bias}_f(h, g) \equiv \text{const.}$ and $D_f(P_{h \times Y} \parallel Q_{h \times Y})$ is \mathcal{H} -robust with label noise for any arbitrary hypothesis space \mathcal{H} .*

This result establishes TV as a strong measure that does not require specifying the noise rates.

Divergences that are conditionally robust Other divergences functions do not enjoy the above nice property as TV has. The robustness of these functions need more careful analysis. Define the following measures that capture the degree a classifier fits to a particular label distribution:

Definition 2. *The fitness of h to $R \in \{Y, \tilde{Y}\}$ is defined as $\text{FIT}(h = y, R = y') := \frac{\mathbb{P}(h(X)=y|R=y')}{\mathbb{P}(h(X)=y)}$.*

FIT measures capture the degree of fit of the classifier to the corresponding label distribution. A high $\text{FIT}(h = y, \tilde{Y} = y)$ (same label) indicates a potential overfit to the noisy label. Denote by

$$\mathcal{H}^* := \{h \in \mathcal{H} : \min_y \text{FIT}(h = y, \tilde{Y} = y) \geq \max_y \text{FIT}(h_f^* = y, Y = y) \geq 1\} \cup \{h_f^*\}$$

The 1 in the “ ≥ 1 ” above corresponds to the **FIT** for a random classifier. \mathcal{H}^* contains the classifiers that are likely to overfit to the noisy labels. We argue, and also as observed in training, that \mathcal{H}^* is the set of classifiers the training should avoid converging to, especially when the training only sees noisy labels. Suppose $\mathbb{P}(Y = +1) = \mathbb{P}(Y = -1)$ (balanced clean labels) and $e_+ = e_-$ (symmetric noise rate), we have the following theorem for binary classification:

Theorem 8. *f -divergences listed in Table 6 (Appendix, except for Jeffrey) are \mathcal{H}^* -robust.*

4.3 MAKING D_f MEASURES ROBUST TO LABEL NOISE

For the general case, to further improve robustness of D_f measures, we will need to estimate the noise rates (e.g., e_+, e_-) and then subtract $\text{Bias}_f(g, h)$ from the noisy variational difference term to correct the bias introduced by the noisy labels. As a corollary of Theorem 4 we have:

Corollary 1. *Maximizing the following bias-corrected $\widetilde{\text{VD}}_f(h, g)$ defined over \tilde{P} and \tilde{Q} leads to h_f^**

$$h_f^* = \underset{h \in \mathcal{H}}{\operatorname{argmax}} \sup_g \mathbb{E}_{\tilde{Z} \sim \tilde{P}_{h \times \tilde{Y}}} [g(\tilde{Z})] - \mathbb{E}_{\tilde{Z} \sim \tilde{Q}_{h \times \tilde{Y}}} [f^*(g(\tilde{Z}))] - \text{Bias}_f(h, g).$$

By removing the **Bias** term, maximizing $\mathbb{E}_{\tilde{Z} \sim \tilde{P}}[g(\tilde{Z})] - \mathbb{E}_{\tilde{Z} \sim \tilde{Q}}[f^*(g(\tilde{Z}))]$ becomes the same with maximizing the divergence defined on the clean distribution $(1 - \sum_{j=1}^K e_j) \cdot \text{VD}_f(h, g)$. The Corollary follows trivially from this fact. The calculation of the **Bias** terms will require the inputs of

noise rates. Our work does not intend to particularly focus on noise rate estimation. But rather, we can leverage the existing results in performing efficient noise rate estimation. There are existing literature on estimating noise rates (noise transition matrix) which can be implemented without the need of ground truth labels. For interested readers, please refer to (Liu & Tao, 2015; Menon et al., 2015; Harish et al., 2016; Patrini et al., 2017; Arazo et al., 2019; Yao et al., 2020; Zhu et al., 2021). We will test the effectiveness of this bias correction step in Section 5.

5 EXPERIMENTS

In this section, we validate our analysis of D_f measures' robustness via a set of empirical evaluations on 5 datasets: MNIST (LeCun et al. (1998)), Fashion-MNIST (Xiao et al. (2017)), CIFAR-10 and CIFAR-100 (Krizhevsky et al. (2009)), and Clothing1M (Xiao et al. (2015)). Omitted experiment details are available in the appendix.

Baselines We compare our approach with five baseline methods: **Cross-Entropy (CE)**, **Backward (BLC)** and **Forward Loss Correction (FLC)** methods as introduced in (Patrini et al., 2017), the **determinant-based mutual information (DMI)** method introduced in (Xu et al., 2019) and **Peer-Loss (PL)** functions in (Liu & Guo, 2020). BLC and FLC methods require estimating the noise transition matrix. DMI and PL are approaches that do not require such estimation.

Noise model We test three types of noise transition models: uniform noise, sparse noise, and random noise. All details of the noise are in the Appendix. Here we briefly overview them. The uniform and sparse noise are as specified at the end of Section 3 for which our theoretical analyses mainly focus on. The noise rates of low-level uniform noise and sparse noise are both approximately 0.2 (the average probability of a label being wrong). The high-levels are about 0.55 and 0.4 respectively. In the random noise setting, each class randomly flips to one of 10 classes with probability p (Random p). For CIFAR-100, the noise rate of uniform noise is about 0.25. The sparse label noise is generated by randomly dividing 100 classes into 50 pairs, and the noise rate is about 0.4.

Optimizing $D_f(\tilde{P}_{h \times \tilde{Y}} \parallel \tilde{Q}_{h \times \tilde{Y}})$ using noisy samples With the noisy training dataset $\{(x_n, \tilde{y}_n)\}_{n=1}^N$, we optimize $D_f(\tilde{P}_{h \times \tilde{Y}} \parallel \tilde{Q}_{h \times \tilde{Y}})$ using gradient ascent of its variational form. Sketch is given in Algorithm 1. For the bias correction version of our algorithm, the gradient will simply include the $\nabla \text{Bias}_f(h, g)$. The variational function \tilde{g}^* can be updated progressively or can be fixed beforehand using an approximate activation function for each f (see e.g., (Nowozin et al., 2016)).

Algorithm 1 Maximizing D_f measures: one step gradient

- 1: **Inputs:** Training data $\{(x_n, \tilde{y}_n)\}_{n=1}^N$, f , variational function \tilde{g}^* , conjugate f^* , classifier h_t .
- 2: Randomly sample three mini-batches $\{(x_n, \tilde{y}_n)\}_{n=1}^B$, $\{(x_n^\dagger, \tilde{y}_n^\dagger)\}_{n=1}^B$, $\{(x_n^\diamond, \tilde{y}_n^\diamond)\}_{n=1}^B$ from $\{(x_n, \tilde{y}_n)\}_{n=1}^N$. $\{(x_n, \tilde{y}_n)\}_{n=1}^B$: simulate samples $\sim P$; $\{(x_n^\dagger, \tilde{y}_n^\dagger)\}_{n=1}^B$ to simulate \tilde{Q} .
- 3: Use $h_{t,x_n}[\tilde{y}_n]$ to denote model prediction on x_n for label \tilde{y}_n , $\tilde{\mathbb{E}}_{\{(x_n, \tilde{y}_n)\}_{n=1}^B}$, $\tilde{\mathbb{E}}_{\{(x_n^\dagger, \tilde{y}_n^\dagger)\}_{n=1}^B}$ to denote the empirical sample mean calculated using the mini-batch data.
- 4: At step t , update h_t by ascending its stochastic gradient with learning rate η_t :

$$h_{t+1} := h_t + \eta_t \cdot \nabla_{h_t} \left[\tilde{\mathbb{E}}_{\{(x_n, \tilde{y}_n)\}_{n=1}^B} [\tilde{g}^*(h_{t,x_n}[\tilde{y}_n])] - \tilde{\mathbb{E}}_{\{(x_n^\dagger, \tilde{y}_n^\dagger)\}_{n=1}^B} [f^*(\tilde{g}^*(h_{t,x_n^\dagger}[\tilde{y}_n^\dagger]))] \right].$$

Tips: In practice, we suggest (also implemented in our experiments) using the fixed form of \tilde{g}^* which appears as $g_f(v)$ in Table 6 (appendix).

5.1 HOW GOOD IS h_f^* ON CLEAN DATA

As a supplementary of Section 2.2, we validate the quality of h_f^* on clean dataset of MNIST, Fashion MNIST, CIFAR-10 and CIFAR-100. In experiments, since the estimation of product noisy distribution are unstable when trained on CIFAR-100 training dataset, we use CE as a warm-up (120 epochs) and then switch to train with D_f measures. For other datasets, we train with D_f measures without the warm-up stage. Results in Table 2 demonstrate that optimizing f -divergence on clean dataset returns a high-quality h_f^* by referring to the performance of CE. Even though D_f measures

can't outperform CE on clean dataset, we do observe that the gap between CE and D_f measures are negligible, for example, the largest gap of Total-Variation (TV) is only 0.81% among four datasets.

Dataset	CE	TV	Gap	J-S	Gap	KL	Gap
MNIST	99.39(99.38±0.01)	99.37(99.34±0.02)	-0.04	99.35(99.31±0.04)	-0.07	99.31(99.21±0.06)	-0.17
Fashion MNIST	90.44(90.34±0.12)	89.98(89.94±0.06)	-0.40	90.40(90.17±0.24)	-0.17	90.19(89.96±0.14)	-0.38
CIFAR-10	93.58(93.47±0.08)	92.80(92.66±0.13)	-0.81	92.35(92.23±0.07)	-1.24	90.55(90.38±0.15)	-3.09
CIFAR-100	73.47(73.39±0.05)	73.43(73.39±0.06)	0.00	73.47(73.26±0.17)	-0.13	73.33(73.16±0.10)	-0.23

Table 2: Experiment results comparison on clean datasets: We report the maximum accuracy of CE and each D_f measures along with (mean \pm standard deviation); Gap: mean performance comparison w.r.t. CE. Numbers highlighted in blue indicate the gap is less than 1%.

5.2 ROBUSTNESS OF D_f MEASURES

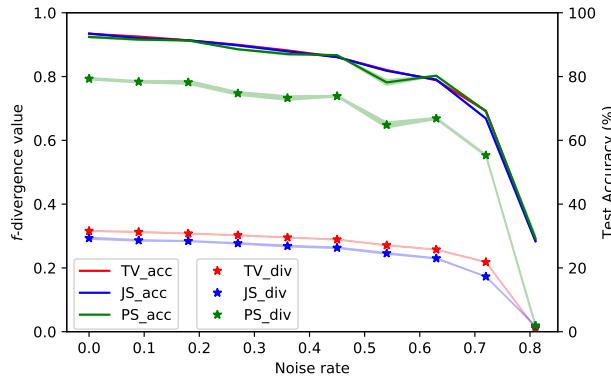


Figure 1: Robustness of TV, JS, PS divergences.

As a demonstration, we apply the uniform noise model to CIFAR-10 dataset to test the robustness of three D_f measures: Total-Variation (TV), Jenson-Shannon (JS) and Pearson (PS). We trained models with D_f measures using Algorithm 1 on 10 noise settings with an increasing noise rate from 0% to approximately 81%. The visualization of the D_f values and accuracy w.r.t. noise rates are shown in Figure 1. Both the D_f values and test accuracy are calculated on the reserved clean test data. We observe that almost all D_f measures are robust to noisy labels, especially when the percentage of noisy labels is not overwhelmingly large, e.g., $\leq 70\%$. Note that the curves for other f -divergences are almost the same as the curve of total variation (TV), which is proved to be robust theoretically. This partially validates the analytical evidences we provided for the robustness of other f -divergences in Section 4.1 and 4.2.

5.3 PERFORMANCE EVALUATION AND COMPARISON

From Table 3, several D_f measures arise as competitive solutions in a variety of noise scenarios. Among the proposed f -divergences, Total Variation (TV) has been consistently ranked as one of the top performing method. This aligns also with our analyses that TV is inherently robust. For most settings, the presented f -divergences outperformed the baselines we compare to, while they fell short to DMI (once) and Peer Loss (5 times) on several cases, particularly when the noise is sparse and high. The sparse high noise setting tends to be a challenging setting for all methods. We conjecture this is because sparse high noise setting creates a highly imbalanced dataset, model training is more likely to converge to a "sub-optimal" early in the training process. It is also possible that with sparse noise, the impact of Bias terms becomes non-negligible. We do observe better performances with very careful and intensive hyper-parameter tuning, but the results are not confident and we chose to not report it. Fully understanding the limitation of our approach in this setting remains an interesting on-going investigation.

In Table 4 (full details on MNIST and Fashion MNIST can be found in Appendix), we use noise transition estimation method in (Patrini et al. (2017)) to estimate the noise rate. The estimates help

us define the bias term and perform bias correction for D_f measures. We observe that while adding bias correction can further improve the performance of several divergence functions (Gap being positive), the improvement or difference is not significant. This partially justified our analysis of the bias term, especially when the noise is dense and high (uniform and random high).

Dataset	Noise	CE	BLC	FLC	DMI	PL	TV	J-S	KL
MNIST	Sparse, Low	97.21	95.23	97.37	97.76	98.59	99.23(99.11±0.08)	99.15(99.03±0.09)	99.21(99.15±0.05)
	Sparse, High	48.55	55.86	49.67	49.61	60.27	58.27(54.72±4.36)	58.93(55.80±1.93)	49.24(49.17±0.06)
	Uniform, Low	97.14	94.27	95.51	97.72	99.06	99.23(99.17±0.05)	99.1(99.08±0.04)	99.13(99.06±0.07)
	Uniform, High	93.25	85.92	87.75	95.50	97.77	98.09(97.96±0.13)	97.86(97.71±0.10)	98.14(97.88±0.18)
	Random (0.2)	98.26	97.46	97.61	98.82	99.25	99.26(99.19±0.05)	99.29(99.27±0.02)	99.26(99.19±0.06)
	Random (0.7)	97.00	93.52	87.74	95.47	98.52	98.81(98.73±0.06)	98.72(98.63±0.08)	98.76(98.65±0.10)
Fashion MNIST	Sparse, Low	84.36	86.02	88.15	85.65	88.32	89.74(89.34±0.33)	88.80(88.79±0.01)	89.77(89.42±0.34)
	Sparse, High	43.33	46.97	47.63	47.16	51.92	45.66(45.22±0.26)	47.46(46.39±0.70)	38.96(38.90±0.06)
	Uniform, Low	82.98	84.48	86.58	83.69	89.31	89.00(88.75±0.16)	88.58(88.46±0.18)	88.32(88.16±0.11)
	Uniform, High	79.52	78.10	82.41	77.94	84.69	85.58(85.07±0.31)	85.62(85.39±0.33)	85.69(85.43±0.30)
	Random (0.2)	85.47	83.40	77.61	86.21	89.78	90.22(90.09±0.19)	89.73(89.43±0.24)	89.24(89.05±0.14)
	Random (0.7)	82.05	78.41	73.42	80.89	87.22	86.69(86.49±0.16)	87.79(87.33±0.29)	87.06(87.00±0.06)
CIFAR-10	Sparse, Low	87.20	72.96	76.17	92.32	91.35	91.81(91.56±0.16)	91.49(91.43±0.08)	91.62(91.32±0.31)
	Sparse, High	61.81	56.30	66.12	27.94	69.70	63.96(62.25±1.00)	67.33(65.27±1.34)	46.55(46.43±0.08)
	Uniform, Low	85.68	72.73	77.12	90.39	91.70	92.10(92.01±0.09)	91.52(91.47±0.08)	92.26(92.08±0.12)
	Uniform, High	71.38	54.41	64.22	82.68	83.42	85.56(85.44±0.08)	84.49(84.35±0.13)	84.36(84.19±0.13)
	Random (0.5)	78.40	59.31	68.97	85.06	86.47	87.28(87.03±0.17)	86.92(86.80±0.10)	86.93(86.85±0.11)
	Random (0.7)	68.26	38.59	54.39	77.91	57.81	80.59(80.45±0.10)	80.50(80.27±0.15)	78.93(78.59±0.30)
CIFAR-100	Uniform	63.87	51.40	60.04	64.39	67.94	69.15(68.90±0.17)	69.13(68.80±0.21)	68.79(68.60±0.11)
	Sparse	40.45	36.57	43.39	40.53	44.25	42.45(38.06±2.82)	38.09(38.00±0.08)	37.74(37.63±0.08)
	Random (0.2)	65.84	61.21	61.52	66.23	62.92	70.43(70.22±0.13)	70.40(70.12±0.21)	70.28(70.06±0.14)
	Random (0.5)	56.92	22.21	55.88	56.06	49.62	62.14(61.89±0.18)	61.58(61.15±0.27)	61.68(61.49±0.13)

Table 3: Experiment results comparison (w/o bias correction): The best performance in each setting is highlighted in **blue**. We report the maximum accuracy of each D_f measures along with (mean \pm standard deviation). All f -divergences will be highlighted if their mean performances are better (or no worse) than all baselines we compare to. A supplementary table including Pearson χ^2 and Jeffrey (JF) is attached in Table 7 (Appendix).

Noise	J-S	Gap	PS	Gap	KL	Gap	JF	Gap
Sparse, Low	91.23(90.93±0.34)	-0.26	91.48(91.12±0.42)	+0.08	91.73(91.57±0.18)	+0.11	91.45(91.18±0.21)	-0.10
Sparse, High	46.45(46.31±0.14)	-20.88	46.31(45.90±0.44)	-0.05	46.59(46.52±0.05)	+0.04	46.25(45.77±0.50)	+0.04
Uniform, Low	92.16(92.09±0.09)	+0.64	92.25(92.13±0.09)	-0.12	90.92(90.84±0.10)	-1.34	92.19(92.10±0.08)	+0.02
Uniform, High	84.31(84.13±0.10)	-0.18	83.79(83.61±0.12)	+0.18	83.98(83.79±0.12)	-0.38	83.93(83.62±0.22)	+0.13

Table 4: D_f measures with bias correction on CIFAR-10: Numbers highlighted in **blue** indicate better than all baseline methods; Gap: relative performance w.r.t. their version w/o bias correction (Table 3); those in **red** indicate better than w/o bias correction.

Clothing1M Clothing1M is a large-scale clothes dataset with comprehensive annotations and can be categorized as a feature-dependent human-level noise dataset. Although this noise setting does not exactly follow our assumption, we are interested in testing the robustness of our f -divergence approaches. Experiment results in Table 5 demonstrate the robustness of the D_f measures. TV and KL divergences have outperformed other baseline methods.

Dataset	Noise	CE	BLC	FLC	DMI	PL	T-V	J-S	Pear	KL	Jeffrey
Clothing1M	Human Noise	68.94	69.13	69.84	72.46	72.60	73.09	72.32	72.22	72.65	72.46

Table 5: Experiment results comparison on Clothing1M dataset.

6 CONCLUSION

In this paper, we explore the robustness of a properly defined f -divergence measure when used to train a classifier in the presence of label noise. We identified a set of nice robustness properties for a family of f -divergence functions. We also experimentally verified our findings. Our work primarily contributed to the problem of learning with noisy labels without requiring the knowledge of noise rate. Beyond this noisy learning problem, the derivation and analysis might be useful for understanding the robustness of f -divergences for other learning tasks.

Acknowledgement This work is partially supported by the National Science Foundation (NSF) under grant IIS-2007951 and the Office of Naval Research under grant N00014-20-1-22.

REFERENCES

- Ehsan Amid, Manfred KK Warmuth, Rohan Anil, and Tomer Koren. Robust bi-tempered logistic loss based on bregman divergences. In *Advances in Neural Information Processing Systems*, pp. 14987–14996, 2019.
- Eric Arazo, Diego Ortego, Paul Albert, Noel E. O’Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. *arXiv preprint arXiv:1904.11238*, 2019.
- Nontawat Charoenphakdee, Jongyeong Lee, and Masashi Sugiyama. On symmetric losses for learning from corrupted labels. *arXiv preprint arXiv:1901.09314*, 2019.
- Hao Cheng, Zhaowei Zhu, Xingyu Li, Yifei Gong, Xing Sun, and Yang Liu. Learning with instance-dependent label noise: A sample sieve approach. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=2VXyy9mIyU3>.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *COLT 2010 - The 23rd Conference on Learning Theory*, 12:257–269, 2010. ISSN 15324435.
- Aritra Ghosh, Naresh Manwani, and PS Sastry. Making risk minimization tolerant to label noise. *Neurocomputing*, 160:93–107, 2015.
- Aritra Ghosh, Himanshu Kumar, and PS Sastry. Robust loss functions under label noise for deep neural networks. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in neural information processing systems*, pp. 8527–8537, 2018.
- Ramaswamy Harish, Clayton Scott, and Ambuj Tewari. Mixture proportion estimation via kernel embeddings of distributions. In *International conference on machine learning*, pp. 2052–2060, 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016.
- Simon Jenni and Paolo Favaro. Deep bilevel learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 618–633, 2018.
- Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. Nlnl: Negative learning for noisy labels. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 101–110, 2019.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38:447–461, 2015.
- Yang Liu and Hongyi Guo. Peer loss functions: Learning from noisy labels without knowing noise rates. *ICML*, 2020.
- Nan Lu, Gang Niu, Aditya K Menon, and Masashi Sugiyama. On the minimal supervision for training any binary classifier from only unlabeled data. *arXiv preprint arXiv:1808.10585*, 2018.
- Michal Lukasik, Srinadh Bhojanapalli, Aditya Krishna Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? *arXiv preprint arXiv:2003.02819*, 2020.

- Naresh Manwani and PS Sastry. Noise tolerance under risk minimization. *IEEE transactions on cybernetics*, 43(3):1146–1151, 2013.
- Aditya Menon, Brendan Van Rooyen, Cheng Soon Ong, and Bob Williamson. Learning from corrupted binary labels via class-probability estimation. In *International Conference on Machine Learning*, pp. 125–134, 2015.
- Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *Advances in neural information processing systems*, pp. 1196–1204, 2013.
- Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems*, pp. 271–279, 2016.
- Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1944–1952, 2017.
- Clayton Scott. A rate of convergence for mixture proportion estimation, with application to learning from noisy labels. In *AISTATS*, 2015.
- Clayton Scott, Gilles Blanchard, Gregory Handy, Sara Pozzi, and Marek Flaska. Classification with asymmetric label noise: Consistency and maximal denoising. In *COLT*, pp. 489–511, 2013.
- Sainbayar Sukhbaatar and Rob Fergus. Learning from noisy labels with deep neural networks. *arXiv preprint arXiv:1406.2080*, 2(3):4, 2014.
- Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5552–5560, 2018.
- Brendan Van Rooyen, Aditya Menon, and Robert C Williamson. Learning with symmetric label noise: The importance of being unhinged. In *Advances in Neural Information Processing Systems*, pp. 10–18, 2015.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2691–2699, 2015.
- Yilun Xu, Peng Cao, Yuqing Kong, and Yizhou Wang. L_dmi: An information-theoretic noise-robust loss function. *NeurIPS, arXiv:1909.03388*, 2019.
- Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Jiankang Deng, Gang Niu, and Masashi Sugiyama. Dual t: Reducing estimation error for transition matrix in label-noise learning. *arXiv preprint arXiv:2006.07805*, 2020.
- Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- Zhilu Zhang and Mert R. Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels, 2018.
- Zhaowei Zhu, Yiwen Song, and Yang Liu. Clusterability as an alternative to anchor points when learning with noisy labels. *arXiv preprint arXiv:2102.05291*, 2021.

A PROOF AND ADDITIONAL THEOREMS

A.1 SHORT NOTATIONS

Throughout the appendix, we will denote by $p := \mathbb{P}(Y = +1)$ and $\tilde{p} := \mathbb{P}(\tilde{Y} = +1) = p \cdot (1 - e_+) + (1 - p) \cdot e_-$. We will also shorthand \tilde{P}, \tilde{Q} for $\tilde{P}_{h \times \tilde{Y}}, \tilde{Q}_{h \times \tilde{Y}}$ and P, Q for $P_{h \times Y}, Q_{h \times Y}$.

A.2 FULL TABLES OF TABLE 1

In experiments, we adopt the output activation function $g_f(v)$ instead of optimal activation functions g^* by referring to the implementation of f-GAN (Nowozin et al., 2016).

Name	$g_f(v)$	g^*	dom_{f^*}	$f^*(u)$
Total Variation (✓)	$\frac{1}{2} \tanh(v)$	$\frac{1}{2} \text{sign} \frac{p(z)}{q(z)} - 1$	$u \in [-\frac{1}{2}, \frac{1}{2}]$	u
Jenson-Shannon (✗)	$\log \frac{2}{1 + e^{-v}}$	$\log \frac{2p(z)}{p(z) + q(z)}$	$u < \log 2$	$-\log(2 - e^u)$
Squared Hellinger (✗)	$1 - e^v$	$1 - \sqrt{\frac{q(z)}{p(z)}}$	$u < 1$	$\frac{u}{1-u}$
Pearson χ^2 (✓)	v	$2 \left(\frac{p(z)}{q(z)} - 1 \right)$	\mathbb{R}	$\frac{1}{4}u^2 + u$
Neyman χ^2 (✗)	$1 - e^v$	$1 - \left(\frac{q(z)}{p(z)} \right)^2$	$u < 1$	$2 - 2\sqrt{1-u}$
KL (✓)	v	$1 + \log \frac{p(z)}{q(z)}$	\mathbb{R}	e^{u-1}
Reverse KL (✗)	$-e^v$	$-\frac{q(z)}{p(z)}$	\mathbb{R}_-	$-1 - \log(-u)$
Jeffrey (✓)	v	$1 + \log \frac{p(z)}{q(z)} - \frac{q(z)}{p(z)}$	\mathbb{R}	$W(e^{1-u}) + \frac{1}{W(e^{1-u})} + u - 2$

Table 6: Exemplary output activation functions g_f (used for approximating g , see e.g. (Nowozin et al., 2016)), optimal activation functions, optimal conjugate functions (full table). W is the Lambert– W product log function. ‘✓’ indicates that the f –divergence function (in practice) is robust to label noise and ‘✗’ means non-robust.

A.3 MAIN RESULTS: PROOF OF THEOREM 4

Proof. First note

$$\begin{aligned}
& \mathbb{E}_{\tilde{Z} \sim \tilde{P}} [g(\tilde{Z})] \\
&= p \cdot \mathbb{E}_{\tilde{Z} \sim \tilde{P}|Y=+1} [g(\tilde{Z})] + (1-p) \cdot \mathbb{E}_{\tilde{Z} \sim \tilde{P}|Y=-1} [g(\tilde{Z})] \\
&= p \cdot \mathbb{E}_{X|Y=+1} [(1 - e_+) \cdot g(h(X), +1) + e_+ \cdot g(h(X), -1)] \\
&\quad + (1-p) \cdot \mathbb{E}_{X|Y=-1} [(1 - e_-) \cdot g(h(X), -1) + e_- \cdot g(h(X), +1)] \\
&= p \cdot \mathbb{E}_{X|Y=+1} [(1 - e_+ - e_-) \cdot g(h(X), +1) + e_+ \cdot g(h(X), -1) + e_- \cdot g(h(X), +1)] \\
&\quad + (1-p) \cdot \mathbb{E}_{X|Y=-1} [(1 - e_+ - e_-) \cdot g(h(X), -1) + e_+ \cdot g(h(X), -1) + e_- \cdot g(h(X), +1)] \\
&= (1 - e_+ - e_-) \cdot \mathbb{E}_{Z \sim P} [g(Z)] + \mathbb{E}_X [e_+ \cdot g(h(X), -1) + e_- \cdot g(h(X), +1)]
\end{aligned}$$

The second term in the variational difference derives as:

$$\begin{aligned}
& \mathbb{E}_{\tilde{Z} \sim \tilde{Q}} [f^*(g(\tilde{Z}))] \\
&= \tilde{p} \cdot \mathbb{E}_X [f^*(g(h(X), +1))] + (1 - \tilde{p}) \cdot \mathbb{E}_X [f^*(g(h(X), -1))] \\
&= p \cdot (1 - e_+ - e_-) \cdot \mathbb{E}_X [f^*(g(h(X), +1))] + e_- \cdot \mathbb{E}_X [f^*(g(h(X), +1))] \\
&\quad + (1-p) \cdot (1 - e_+ - e_-) \cdot \mathbb{E}_X [f^*(g(h(X), -1))] + e_+ \cdot \mathbb{E}_X [f^*(g(h(X), -1))] \\
&= (1 - e_+ - e_-) \cdot \mathbb{E}_{Z \sim Q} [f^*(g(Z))] + \mathbb{E}_X [e_- \cdot f^*(g(h(X), +1)) + e_+ \cdot f^*(g(h(X), -1))]
\end{aligned}$$

Combining $\mathbb{E}_{\tilde{Z} \sim \tilde{P}} [g(\tilde{Z})]$ and $\mathbb{E}_{\tilde{Z} \sim \tilde{Q}} [f^*(g(\tilde{Z}))]$: the leading terms combine into

$$(1 - e_+ - e_-) \cdot \mathbb{E}_{Z \sim P} [g(Z)] - (1 - e_+ - e_-) \cdot \mathbb{E}_{Z \sim Q} [f^*(g(Z))] = (1 - e_+ - e_-) \cdot \mathbf{VD}_f(h, g)$$

and the rest:

$$\begin{aligned} & \mathbb{E}_X[e_+ \cdot g(h(X), -1) + e_- \cdot g(h(X), +1)] - \mathbb{E}_X[e_- \cdot f^*(g(h(X), +1)) + e_+ \cdot f^*(g(h(X), -1))] \\ &= e_+ \cdot \Delta_f^{-1}(h, g) + e_- \cdot \Delta_f^{+1}(h, g) \\ &= \text{Bias}_f(h, g) \end{aligned}$$

we proved the claim. \square

A.4 PROOF OF THEOREM 5: MULTI-CLASS EXTENSION I OF THEOREM 4

Proof. Denote $p_i = \mathbb{P}(Y = i)$, $\tilde{p}_i := \mathbb{P}(\tilde{Y} = i)$, $\tilde{p}_i = (1 - \sum_{j \neq i} e_j) \cdot p_i + e_i \cdot \sum_{j \neq i} p_j$. We have:

$$\begin{aligned} & \mathbb{E}_{\tilde{Z} \sim \tilde{P}} [g(\tilde{Z})] \\ &= \sum_{i=1}^K p_i \cdot \mathbb{E}_{\tilde{Z} \sim \tilde{P}|Y=i} [g(\tilde{Z})] = \sum_{i=1}^K p_i \cdot \mathbb{E}_{X|Y=i} \left[\sum_{j=1}^K T_{i,j} \cdot g(h(X), \tilde{Y} = j) \right] \\ &= \sum_{i=1}^K p_i \cdot \mathbb{E}_{X|Y=i} \left[(1 - \sum_{j \neq i} e_j) \cdot g(h(X), \tilde{Y} = i) + \sum_{j \neq i} e_j \cdot g(h(X), \tilde{Y} = j) \right] \\ &= \sum_{i=1}^K p_i \cdot \mathbb{E}_{X|Y=i} \left[(1 - \sum_{j=1}^K e_j) \cdot g(h(X), \tilde{Y} = i) + \sum_{j=1}^K e_j \cdot g(h(X), \tilde{Y} = j) \right] \\ &= (1 - \sum_{j=1}^K e_j) \cdot \mathbb{E}_{Z \sim P}[g(Z)] + \sum_{j=1}^K e_j \cdot \mathbb{E}_X[g(h(X), j)] \\ \\ & \mathbb{E}_{\tilde{Z} \sim \tilde{Q}} [f^*(g(\tilde{Z}))] \\ &= \sum_{i=1}^K \tilde{p}_i \cdot \mathbb{E}_X [f^*(g(h(X), \tilde{Y} = i))] \\ &= \sum_{i=1}^K [(1 - \sum_{j=1}^K e_j) \cdot p_i + e_i \cdot \sum_{j=1}^K p_j] \cdot \mathbb{E}_X [f^*(g(h(X), \tilde{Y} = i))] \\ &= (1 - \sum_{j=1}^K e_j) \cdot \mathbb{E}_{Z \sim Q}[f^*(g(Z))] + \sum_{j=1}^K e_j \cdot \mathbb{E}_X [f^*(g(h(X), j))] \end{aligned}$$

Similar to the binary case, combining $\mathbb{E}_{\tilde{Z} \sim \tilde{P}} [g(\tilde{Z})]$ and $\mathbb{E}_{\tilde{Z} \sim \tilde{Q}} [f^*(g(\tilde{Z}))]$ we proved the claim. \square

A.5 MULTI-CLASS EXTENSION II OF THEOREM 4: SPARSE CASE

For sparse transition matrix, assume K is an even number, sparse noise model specifies $\frac{K}{2}$ disjoint pairs of classes (i_c, j_c) where $c \in [\frac{K}{2}]$ and $i_c < j_c$. The diagonal entry T_{i_c, i_c} becomes $1 - T_{i_c, j_c}$. Suppose $\forall c \in [\frac{K}{2}], T_{i_c, j_c} = e_{p_1}, T_{j_c, i_c} = e_{p_2}, e_{p_1} + e_{p_2} < 1$.

Theorem 9. [Multi-class extension II] In the scenario of sparse noise transition model, the variational difference between the noisy distributions \tilde{P} and \tilde{Q} relates to the one defined over the clean distributions in the following way:

$$\widetilde{VD}_f(h, g) = (1 - e_{p_1} - e_{p_2}) \cdot VD_f(h, g) + \sum_{(i_c, j_c)} \left[e_{p_1} \cdot \Delta_f^{j_c}(g, h) + e_{p_2} \cdot \Delta_f^{i_c}(g, h) \right].$$

Proof.

$$\begin{aligned}
& \mathbb{E}_{\tilde{Z} \sim \tilde{P}} [g(\tilde{Z})] \\
&= \sum_{i=1}^K p_i \cdot \mathbb{E}_{\tilde{Z} \sim \tilde{P}|Y=i} [g(\tilde{Z})] = \sum_{i=1}^K p_i \cdot \mathbb{E}_{X|Y=i} \left[\sum_{j=1}^K T_{i,j} \cdot g(h(X), \tilde{Y}=j) \right] \\
&= \sum_{i_c} p_{i_c} \cdot \mathbb{E}_{X|Y=i_c} \left[(1 - e_{p_1}) \cdot g(h(X), \tilde{Y}=i_c) + e_{p_1} \cdot g(h(X), \tilde{Y}=j_c) \right] \\
&\quad + \sum_{j_c} p_{j_c} \cdot \mathbb{E}_{X|Y=j_c} \left[(1 - e_{p_2}) \cdot g(h(X), \tilde{Y}=j_c) + e_{p_2} \cdot g(h(X), \tilde{Y}=i_c) \right] \\
&= \sum_{i=1}^K p_i \cdot \mathbb{E}_{X|Y=i} \left[(1 - e_{p_1} - e_{p_2}) \cdot g(h(X), \tilde{Y}=i) \right] \\
&\quad + \sum_{(i_c, j_c)} (e_{p_1} \cdot g(h(X), \tilde{Y}=j_c) + e_{p_2} \cdot g(h(X), \tilde{Y}=i_c)) \\
&= (1 - e_{p_1} - e_{p_2}) \cdot \mathbb{E}_{Z \sim P}[g(Z)] + \sum_{(i_c, j_c)} \mathbb{E}_X \left[e_{p_1} \cdot g(h(X), \tilde{Y}=j_c) + e_{p_2} \cdot g(h(X), \tilde{Y}=i_c) \right]
\end{aligned}$$

Similarly, we have:

$$\begin{aligned}
& \mathbb{E}_{\tilde{Z} \sim \tilde{Q}} [f^*(g(\tilde{Z}))] \\
&= \sum_{i=1}^K \tilde{p}_i \cdot \mathbb{E}_X \left[f^*(g(h(X), \tilde{Y}=i)) \right] \\
&= \sum_{i=1}^K \left[(1 - \sum_{j=1}^K e_j) \cdot p_i + e_i \cdot \sum_{j=1}^K p_j \right] \cdot \mathbb{E}_X \left[f^*(g(h(X), \tilde{Y}=i)) \right] \\
&= (1 - e_{p_1} - e_{p_2}) \cdot \mathbb{E}_{Z \sim Q}[f^*(g(Z))] \\
&\quad + \sum_{(i_c, j_c)} \mathbb{E}_X \left[e_{p_1} \cdot f^*(g(h(X), \tilde{Y}=j_c)) + e_{p_2} \cdot f^*(g(h(X), \tilde{Y}=i_c)) \right]
\end{aligned}$$

Combining $\mathbb{E}_{\tilde{Z} \sim \tilde{P}} [g(\tilde{Z})]$ and $\mathbb{E}_{\tilde{Z} \sim \tilde{Q}} [f^*(g(\tilde{Z}))]$ we proved the claim. \square

A.6 PROOF OF THEOREM 1: TOTAL-VARIATION GENERATES BAYES OPTIMAL

Proof. For total variation, we have

$$D_f(P_{h \times Y} \| Q_{h \times Y}) = \frac{1}{2} \sum_{y, y' \in \{-1, +1\}} |\mathbb{P}(h(X) = y, Y = y') - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = y')|$$

We again present the main proof for the binary classification setting.

First note the following fact that

$$\mathbb{P}(h(X) = y, Y = y) - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = y)$$

and

$$\mathbb{P}(h(X) = y, Y = -y) - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = -y)$$

have opposite signs. This is simply because

$$\mathbb{P}(h(X) = y, Y = y) - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = y) + \mathbb{P}(h(X) = y, Y = -y) - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = -y) = 0$$

Because of the above, there are four possible combinations of cases:

Case 1 $\mathbb{P}(h(X) = +1, Y = +1) - \mathbb{P}(h(X) = +1) \cdot \mathbb{P}(Y = +1) > 0, \mathbb{P}(h(X) = -1, Y = -1) - \mathbb{P}(h(X) = -1) \cdot \mathbb{P}(Y = -1) > 0$:

$$\begin{aligned} & \sum_{y,y' \in \{-1,+1\}} |\mathbb{P}(h(X) = y, Y = y') - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = y')| \\ &= \mathbb{P}(h(X) = +1, Y = +1) - \mathbb{P}(h(X) = +1) \cdot \mathbb{P}(Y = +1) \\ &\quad - \mathbb{P}(h(X) = +1, Y = -1) + \mathbb{P}(h(X) = +1) \cdot \mathbb{P}(Y = -1) \\ &\quad + \mathbb{P}(h(X) = -1, Y = -1) - \mathbb{P}(h(X) = -1) \cdot \mathbb{P}(Y = -1) \\ &\quad - \mathbb{P}(h(X) = -1, Y = +1) - \mathbb{P}(h(X) = -1) \cdot \mathbb{P}(Y = +1) \\ &= \mathbb{P}(h(X) = Y) - \mathbb{P}(h(X) \neq Y) \\ &= 2\mathbb{P}(h(X) = Y) - 1 \end{aligned}$$

Therefore, maximizing D_f total variation returns the Bayes optimal classifier h^* , and the optimal value arrives at $\mathbb{P}(h^*(X) = Y) - \frac{1}{2}$.

Case 2 $\mathbb{P}(h(X) = +1, Y = +1) - \mathbb{P}(h(X) = +1) \cdot \mathbb{P}(Y = +1) < 0, \mathbb{P}(h(X) = -1, Y = -1) - \mathbb{P}(h(X) = -1) \cdot \mathbb{P}(Y = -1) > 0$:

$$\begin{aligned} & \sum_{y,y' \in \{-1,+1\}} |\mathbb{P}(h(X) = y, Y = y') - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = y')| \\ &= -\mathbb{P}(h(X) = +1, Y = +1) + \mathbb{P}(h(X) = +1) \cdot \mathbb{P}(Y = +1) \\ &\quad + \mathbb{P}(h(X) = +1, Y = -1) - \mathbb{P}(h(X) = +1) \cdot \mathbb{P}(Y = -1) \\ &\quad + \mathbb{P}(h(X) = -1, Y = -1) - \mathbb{P}(h(X) = -1) \cdot \mathbb{P}(Y = -1) \\ &\quad - \mathbb{P}(h(X) = -1, Y = +1) - \mathbb{P}(h(X) = -1) \cdot \mathbb{P}(Y = +1) \\ &= \mathbb{P}(Y = -1) - \mathbb{P}(Y = +1) \\ &= 0 \end{aligned}$$

Case 3 $\mathbb{P}(h(X) = +1, Y = +1) - \mathbb{P}(h(X) = +1) \cdot \mathbb{P}(Y = +1) > 0, \mathbb{P}(h(X) = -1, Y = -1) - \mathbb{P}(h(X) = -1) \cdot \mathbb{P}(Y = -1) < 0$: This case is symmetrical to Case 2.

Case 4 This is symmetrical to Case 1:

$$D_f(P_{h \times Y} \| Q_{h \times Y}) = \mathbb{P}(h(X) \neq Y) - \frac{1}{2}$$

The optimal classifier is then the opposite of h^* , but

$$\mathbb{P}(h^*(X) = Y) - \frac{1}{2} > \mathbb{P}(h^*(X) \neq Y) - \frac{1}{2}$$

so the maximizer returns a smaller value compared to Case 1.

Multi-class extension We provide arguments for the multi-class generalization. First note that

$$\begin{aligned} & \sum_{y,y' \in \mathcal{Y}} |\mathbb{P}(h(X) = y, Y = y') - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = y')| \\ &= \sum_{y,y' \in \mathcal{Y}} |\mathbb{P}(h(X) = y | Y = y') - \mathbb{P}(h(X) = y)| \cdot \mathbb{P}(Y = y') \\ &= \frac{1}{K} \sum_{y'} \sum_y |\mathbb{P}(h(X) = y | Y = y') - \mathbb{P}(h(X) = y)| \end{aligned}$$

For any classifier h and for each y , one of the following terms

$$\mathbb{P}(h(X) = y, Y = 1) - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = 1), \dots, \mathbb{P}(h(X) = y, Y = K) - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = K)$$

must be non-negative as: $\sum_{y'} \mathbb{P}(h(X) = y, Y = y') - \mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y = y') = 0$.

Our following derivation focuses on confident classifiers:

Definition 3. We call a classifier confident if for each label class y , only one class $y_k \in \mathcal{Y}$ returns positive correlation:

$$\mathbb{P}(h(X) = y_k, Y = k) - \mathbb{P}(h(X) = y_k) \cdot \mathbb{P}(Y = k) \geq 0$$

while for all other $y' \neq y_k$, we have $\mathbb{P}(h(X) = y', Y = k) - \mathbb{P}(h(X) = y') \cdot \mathbb{P}(Y = k) \leq 0$

This above definition is saying the classifier h is “dominantly” confident in predicting one class for the each true label class.

For a given class k , if all other classes $k' \neq y_k$ are negative in $\mathbb{P}(h(X) = k', Y = k) - \mathbb{P}(h(X) = k') \cdot \mathbb{P}(Y = k)$, the total variation becomes:

$$\begin{aligned} & \sum_{k'} |\mathbb{P}(h(X) = k', Y = k) - \mathbb{P}(h(X) = k') \cdot \mathbb{P}(Y = k)| \\ &= \mathbb{P}(h(X) = y_k, Y = k) - \mathbb{P}(h(X) = y_k) \cdot \mathbb{P}(Y = k) \\ &+ \sum_{k' \neq y_k} (\mathbb{P}(h(X) = k') \cdot \mathbb{P}(Y = k) - \mathbb{P}(h(X) = k', Y = k)) \\ &= \mathbb{P}(h(X) = y_k, Y = k) - \mathbb{P}(h(X) = y_k) \cdot \mathbb{P}(Y = k) \\ &+ \mathbb{P}(Y = k)(1 - \mathbb{P}(h(X) = y_k)) - (1 - \mathbb{P}(h(X) = y_k, Y = k)) \\ &= 2(\mathbb{P}(h(X) = y_k, Y = k) - \mathbb{P}(h(X) = y_k) \cdot \mathbb{P}(Y = k)). \end{aligned}$$

Summing up, for a confident classifier, the total variation becomes (ignoring constant 2):

$$\begin{aligned} & \sum_k \mathbb{P}(h(X) = y_k, Y = k) - \mathbb{P}(h(X) = y_k) \cdot \mathbb{P}(Y = k) \\ &= \frac{1}{K} \sum_k \mathbb{P}(h(X) = y_k | Y = k) - \mathbb{P}(h(X) = y_k) \\ &= \frac{1}{K} \sum_k \mathbb{P}(h(X) = y_k | Y = k) - 1 \\ &= \frac{1}{K} \sum_k \frac{\mathbb{P}(Y = k | h(X) = y_k) \mathbb{P}(h(X) = y_k)}{\mathbb{P}(Y = k)} - 1 \\ &= \sum_k \int_X f_X(x) \cdot \mathbb{P}(Y = k | h(x) = y_k) \mathbb{P}(h(x) = y_k) dx - 1 \\ &= \sum_k \int_X f_X(x) \cdot \mathbb{P}(Y = k | X = x) \cdot 1(h(x) = y_k) \cdot \mathbb{P}(h(x) = y_k) dx - 1 \\ &= \int_X f_X(x) \cdot \sum_k \mathbb{P}(Y = k | X = x) \cdot 1(h(x) = y_k) \cdot \mathbb{P}(h(x) = y_k) dx - 1 \\ &\leq \int_X f_X(x) \cdot \sum_k \mathbb{P}(Y = k | X = x) \cdot 1(h^*(X) = k) \cdot \mathbb{P}(h^*(X) = k) dx - 1 \\ &= \sum_k \mathbb{P}(h^*(X) = k, Y = k) - \mathbb{P}(h^*(X) = k) \cdot \mathbb{P}(Y = k). \end{aligned}$$

where the last inequality is due to the fact that the Bayes optimal classifier selects the highest $\mathbb{P}(Y = k | X)$ for each x .

□

A.7 PROOF OF THEOREM 3

Proof. By definition of D_f :

$$D_f(P_{h \times Y^*} \| Q_{h \times Y^*}) = \sum_{y, y'} \mathbb{P}(h(X) = y, Y^* = y') \cdot f\left(\frac{\mathbb{P}(h(X) = y, Y^* = y')}{\mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y^* = y')}\right)$$

First we want to prove

$$\left| \frac{\mathbb{P}(h^*(X) = y, Y^* = y)}{\mathbb{P}(h^*(X) = y) \cdot \mathbb{P}(Y^* = y)} - 1 \right| > \left| \frac{\mathbb{P}(h(X) = y', Y^* = y)}{\mathbb{P}(h(X) = y') \cdot \mathbb{P}(Y^* = y)} - 1 \right|, \forall h \neq h^*$$

This is because:

$$\frac{\mathbb{P}(h^*(X) = y, Y^* = y)}{\mathbb{P}(h^*(X) = y) \cdot \mathbb{P}(Y^* = y)} = \frac{\mathbb{P}(h^*(X) = y | Y^* = y)}{\mathbb{P}(h^*(X) = y)} = \frac{1}{\mathbb{P}(Y^* = y)}$$

On the other hand

$$\begin{aligned} & \frac{\mathbb{P}(h(X) = y', Y^* = y)}{\mathbb{P}(h(X) = y') \cdot \mathbb{P}(Y^* = y)} \\ &= \frac{\mathbb{P}(h(X) = y' | Y^* = y)}{\mathbb{P}(h(X) = y' | Y^* = y) \mathbb{P}(Y^* = y) + \mathbb{P}(h(X) = y' | Y^* = -y) \mathbb{P}(Y^* = -y)} \\ &= \frac{1}{\mathbb{P}(Y^* = y) + \frac{\mathbb{P}(h(X) = y' | Y^* = -y)}{\mathbb{P}(h(X) = y' | Y^* = y)} \mathbb{P}(Y^* = -y)} \end{aligned}$$

When $\frac{\mathbb{P}(h(X) = y' | Y^* = -y)}{\mathbb{P}(h(X) = y' | Y^* = y)} < 1$, $\mathbb{P}(Y^* = y) + \frac{\mathbb{P}(h(X) = y' | Y^* = -y)}{\mathbb{P}(h(X) = y' | Y^* = y)} \mathbb{P}(Y^* = -y) < \mathbb{P}(Y^* = y) + \mathbb{P}(Y^* = -y) = 1$, therefore $\frac{\mathbb{P}(h(X) = y', Y^* = y)}{\mathbb{P}(h(X) = y') \cdot \mathbb{P}(Y^* = y)} > 1$. Further

$$\frac{\mathbb{P}(h(X) = y', Y^* = y)}{\mathbb{P}(h(X) = y') \cdot \mathbb{P}(Y^* = y)} < \frac{1}{\mathbb{P}(Y^* = y)} = \frac{\mathbb{P}(h^*(X) = y, Y^* = y)}{\mathbb{P}(h^*(X) = y) \cdot \mathbb{P}(Y^* = y)}$$

When $\frac{\mathbb{P}(h(X) = y' | Y^* = -y)}{\mathbb{P}(h(X) = y' | Y^* = y)} > 1$, denote $\alpha := \frac{\mathbb{P}(h(X) = y' | Y^* = -y)}{\mathbb{P}(h(X) = y' | Y^* = y)} > 1$. We have

$$\begin{aligned} & 1 - \frac{1}{\mathbb{P}(Y^* = y) + \alpha \cdot \mathbb{P}(Y^* = -y)} \\ &= \frac{(\alpha - 1)\mathbb{P}(Y^* = -y)}{\mathbb{P}(Y^* = y) + \alpha \cdot \mathbb{P}(Y^* = -y)} \\ &= \frac{\mathbb{P}(Y^* = -y)}{\mathbb{P}(Y^* = y)} \cdot \frac{\alpha - 1}{\alpha + 1} \\ &< \frac{\mathbb{P}(Y^* = -y)}{\mathbb{P}(Y^* = y)} \\ &= \frac{1}{\mathbb{P}(Y^* = y)} - 1. \end{aligned}$$

Therefore we proved

$$\left| \frac{\mathbb{P}(h^*(X) = y, Y^* = y)}{\mathbb{P}(h^*(X) = y) \cdot \mathbb{P}(Y^* = y)} - 1 \right| > \left| \frac{\mathbb{P}(h(X) = y', Y^* = y)}{\mathbb{P}(h(X) = y') \cdot \mathbb{P}(Y^* = y)} - 1 \right|, \forall h \neq h^*$$

Because $f(v)$ is monotonically increasing in $|v - 1|$, we proved that

$$\begin{aligned} & D_f(P_{h \times Y^*} \| Q_{h \times Y^*}) \\ &= \sum_{y, y'} \mathbb{P}(h(X) = y, Y^* = y') \cdot f\left(\frac{\mathbb{P}(h(X) = y, Y^* = y')}{\mathbb{P}(h(X) = y) \cdot \mathbb{P}(Y^* = y')}\right) \\ &< \sum_{y, y'} \mathbb{P}(h(X) = y, Y^* = y') f\left(\frac{\mathbb{P}(h^*(X) = y', Y^* = y')}{\mathbb{P}(h^*(X) = y') \cdot \mathbb{P}(Y^* = y')}\right) \\ &= \sum_y \mathbb{P}(Y^* = y) f\left(\frac{\mathbb{P}(h^*(X) = y, Y^* = y)}{\mathbb{P}(h^*(X) = y) \cdot \mathbb{P}(Y^* = y)}\right) \\ &= D_f(P_{h^* \times Y^*} \| Q_{h^* \times Y^*}) \end{aligned}$$

The last equality is because h^* always agrees with Y^* , so $\mathbb{P}(h^*(X) = y, Y^* = y) = \mathbb{P}(Y^* = y)$ and $\mathbb{P}(h^*(X) = -y, Y^* = y) = 0$.

□

A.8 PROOF OF THEOREM 6: \mathcal{H} -ROBUST

Proof. The proofs for the multi-class case under uniform diagonal and sparse noise setting are entirely symmetrical due to Theorem 5 and 9. We deliver the main idea for the binary case.

The proof for condition (I) is easy to see:

$$\begin{aligned}
& \underset{h \in \mathcal{H}}{\operatorname{argmax}} D_f(\tilde{P}_{h \times \tilde{Y}} \| \tilde{Q}_{h \times \tilde{Y}}) \\
&= \underset{h \in \mathcal{H}}{\operatorname{argmax}} \sup_g \mathbb{E}_{\tilde{Z} \sim \tilde{P}} [g(\tilde{Z})] - \mathbb{E}_{\tilde{Z} \sim \tilde{Q}} [f^*(g(\tilde{Z}))] \\
&= \underset{h \in \mathcal{H}}{\operatorname{argmax}} \sup_g (1 - e_+ - e_-) [\mathbb{E}_{Z \sim P} [g(Z)] - \mathbb{E}_{Z \sim Q} [f^*(g(Z))]] + \text{Bias}_f(h, g) \\
&= \underset{h \in \mathcal{H}}{\operatorname{argmax}} \sup_g \mathbb{E}_{Z \sim P} [g(Z)] - \mathbb{E}_{Z \sim Q} [f^*(g(Z))] \\
&= \underset{h \in \mathcal{H}}{\operatorname{argmax}} D_f(P_{h \times Y} \| Q_{h \times Y}) \\
&= h_f^*.
\end{aligned}$$

Now we prove the robustness of D_f under condition (II). Denote by h' the classifier that maximizes $D_f(\tilde{P}_{h' \times \tilde{Y}} \| \tilde{Q}_{h' \times \tilde{Y}})$ and

$$D_f(\tilde{P}_{h' \times \tilde{Y}} \| \tilde{Q}_{h' \times \tilde{Y}}) > D_f(\tilde{P}_{h_f^* \times \tilde{Y}} \| \tilde{Q}_{h_f^* \times \tilde{Y}})$$

But

$$\begin{aligned}
& D_f(\tilde{P}_{h' \times \tilde{Y}} \| \tilde{Q}_{h' \times \tilde{Y}}) \\
&= (1 - e_+ - e_-) [\mathbb{E}_{Z \sim P} [\tilde{g}^*([h'(X), Y])] - \mathbb{E}_{Z \sim Q} [f^*(\tilde{g}^*([h'(X), Y]))]] + \text{Bias}_f(h', \tilde{g}^*) \\
&\leq \max_{h \in \mathcal{H}} (1 - e_+ - e_-) \cdot \sup_g [\mathbb{E}_{Z \sim P} [g([h(X), Y])] - \mathbb{E}_{Z \sim Q} [f^*(g([h(X), Y]))]] + \text{Bias}_f(h_f^*, g^*) \\
&\quad (\text{Bias}_f(h, \tilde{g}^*) \leq \text{Bias}_f(h_f^*, g^*)) \\
&= \max_{h \in \mathcal{H}} (1 - e_+ - e_-) \cdot D_f(P_{h \times Y} \| Q_{h \times Y}) + \text{Bias}_f(h_f^*, g^*) \quad (\text{variational form of } D_f) \\
&= (1 - e_+ - e_-) \cdot D_f(P_{h_f^* \times Y} \| Q_{h_f^* \times Y}) + \text{Bias}_f(h_f^*, g^*) \\
&\leq \sup_g [\mathbb{E}_{Z=[h_f^*(X), Y] \sim P} [g(Z)] - \mathbb{E}_{Z=[h_f^*(X), Y] \sim Q} [f^*(g(Z))]] + \text{Bias}_f(h_f^*, g) \\
&= D_f(\tilde{P}_{h_f^* \times \tilde{Y}} \| \tilde{Q}_{h_f^* \times \tilde{Y}}),
\end{aligned}$$

which is a contradiction. \square

A.9 PROOF OF LEMMA 1: IMPACT OF BIAS TERM FOR DIFFERENT f -DIVERGENCES

Denote $p_y = \mathbb{P}(Y = y)$, $\tilde{p}_y := \mathbb{P}(\tilde{Y} = y)$, $\frac{\tilde{P}(y, y')}{\tilde{Q}(y, y')} := \frac{\mathbb{P}(h(x)=y, \tilde{Y}=y')}{\mathbb{P}(h(x)=y) \cdot \mathbb{P}(\tilde{Y}=y')}$. We first prove that $\left| \frac{\tilde{P}(y, y')}{\tilde{Q}(y, y')} - 1 \right|$ approaches to 0 as a function of $1 - e_+ - e_-$:

Proposition 10. When $e_+, e_- < 0.5$, $\left| \frac{\tilde{P}(y, y')}{\tilde{Q}(y, y')} - 1 \right| = O((1 - e_+ - e_-))$.

Proof.

$$\begin{aligned} \left| \frac{\tilde{P}(y, y')}{\tilde{Q}(y, y')} - 1 \right| &= \left| \frac{\mathbb{P}(h(x) = y | \tilde{Y} = y')}{\mathbb{P}(h(x) = y)} - 1 \right| \\ &= \left| \frac{\mathbb{P}(h(x) = y | \tilde{Y} = y') - \tilde{p}_{y'} \mathbb{P}(h(x) = y | \tilde{Y} = y') - (1 - \tilde{p}_{y'}) \mathbb{P}(h(x) = y | \tilde{Y} = -y')}{\mathbb{P}(h(x) = y)} \right| \\ &= \frac{1 - \tilde{p}_{y'}}{\mathbb{P}(h(x) = y)} \cdot |\mathbb{P}(h(x) = y | \tilde{Y} = y') - \mathbb{P}(h(x) = y | \tilde{Y} = -y')| \end{aligned}$$

$|\mathbb{P}(h(x) = y | \tilde{Y} = y') - \mathbb{P}(h(x) = y | \tilde{Y} = -y')|$ derives as

$$\begin{aligned} &|\mathbb{P}(h(x) = y | \tilde{Y} = y') - \mathbb{P}(h(x) = y | \tilde{Y} = -y')| \\ &= \left| p_{y'} \mathbb{P}(h(x) = y | Y = y') \cdot \mathbb{P}(Y = y' | \tilde{Y} = y') + (1 - p_{y'}) \mathbb{P}(h(x) = y | Y = -y') \cdot \mathbb{P}(Y = -y' | \tilde{Y} = y') \right. \\ &\quad \left. - p_{y'} \mathbb{P}(h(x) = y | Y = y) \cdot \mathbb{P}(Y = y' | \tilde{Y} = -y') \right. \\ &\quad \left. - (1 - p_{y'}) \mathbb{P}(h(x) = y | Y = -y) \cdot \mathbb{P}(Y = -y' | \tilde{Y} = -y') \right| \\ &= \left| p_{y'} \mathbb{P}(h(x) = y | Y = y') \cdot (\mathbb{P}(Y = y' | \tilde{Y} = y') - \mathbb{P}(Y = y' | \tilde{Y} = -y')) \right. \\ &\quad \left. + (1 - p_{y'}) \mathbb{P}(h(x) = y | Y = -y) \cdot (\mathbb{P}(Y = -y' | \tilde{Y} = y') - \mathbb{P}(Y = -y' | \tilde{Y} = -y')) \right| \\ &= \left| p_{y'} \mathbb{P}(h(x) = y | Y = y') - (1 - p_{y'}) \mathbb{P}(h(x) = y | Y = -y') \right| \cdot |\mathbb{P}(Y = y' | \tilde{Y} = y') - \mathbb{P}(Y = y' | \tilde{Y} = -y')| \end{aligned}$$

The last equation is satisfied because $\forall y$:

$$\begin{aligned} &\mathbb{P}(Y = y | \tilde{Y} = y) + \mathbb{P}(Y = -y | \tilde{Y} = y) = \mathbb{P}(Y = -y | \tilde{Y} = -y) + \mathbb{P}(Y = y | \tilde{Y} = -y) \\ \iff &\mathbb{P}(Y = y | \tilde{Y} = y) - \mathbb{P}(Y = y | \tilde{Y} = -y) = \mathbb{P}(Y = -y | \tilde{Y} = -y) - \mathbb{P}(Y = -y | \tilde{Y} = y) \end{aligned}$$

Now focus on $|\mathbb{P}(Y = y' | \tilde{Y} = y') - \mathbb{P}(Y = y' | \tilde{Y} = -y')|$:

$$\begin{aligned} &|\mathbb{P}(Y = y' | \tilde{Y} = y') - \mathbb{P}(Y = y' | \tilde{Y} = -y')| \\ &= \left| \frac{\mathbb{P}(\tilde{Y} = y' | Y = y') \cdot p_{y'}}{\mathbb{P}(\tilde{Y} = y')} - \frac{\mathbb{P}(\tilde{Y} = -y' | Y = y') \cdot p_{y'}}{\mathbb{P}(\tilde{Y} = -y')} \right| \\ &= \frac{p_{y'}}{\tilde{p}_{y'} \cdot (1 - \tilde{p}_{y'})} |(1 - e_{y'}) \cdot (1 - \tilde{p}_{y'}) - e_{y'} \cdot \tilde{p}_{y'}| \\ &= \frac{p_{y'}}{\tilde{p}_{y'} \cdot (1 - \tilde{p}_{y'})} |(1 - e_{y'}) \cdot (1 - p_{y'}(1 - e_{y'}) - (1 - p_{y'})e_{-y'}) - e_{y'} \cdot (p_{y'}(1 - e_{y'}) + (1 - p_{y'})e_{-y'})| \\ &= \frac{p_{y'}}{\tilde{p}_{y'} \cdot (1 - \tilde{p}_{y'})} |(1 - p_{y'})(1 - e_+ - e_-)| \\ &= \frac{p_{y'} \cdot (1 - p_{y'})}{\tilde{p}_{y'} \cdot (1 - \tilde{p}_{y'})} |1 - e_+ - e_-| \end{aligned}$$

Putting everything up together:

$$\left| \frac{\tilde{P}(y, y')}{\tilde{Q}(y, y')} - 1 \right| = p_{y'}(1 - p_{y'}) (p_{y'} \mathbb{P}(h(X) = y | Y = y') - (1 - p_{y'}) \mathbb{P}(h(X) = y | Y = -y')) \frac{|1 - e_+ - e_-|}{\tilde{p}_{y'}}$$

When $e_+, e_- < 0.5$, we have

$$\tilde{p}_{y'} = p_{y'}(1 - e_{y'}) + p_{-y'}e_{-y'} \geq 0.5 \min\{p, 1 - p\}.$$

Therefore

$$\begin{aligned} &p_{y'}(1 - p_{y'}) (p_{y'} \mathbb{P}(h(X) = y | Y = y') - (1 - p_{y'}) \mathbb{P}(h(X) = y | Y = -y')) \frac{|1 - e_+ - e_-|}{\tilde{p}_{y'}} \\ &\leq 2 \max\{p, 1 - p\} (p_{y'} \mathbb{P}(h(X) = y | Y = y') - (1 - p_{y'}) \mathbb{P}(h(X) = y | Y = -y')) \cdot |1 - e_+ - e_-| \end{aligned}$$

□

Next, we prove Lemma 1:

Proof. Shorthand $\tilde{P}(y, y') := \mathbb{P}(h(x) = y, \tilde{Y} = y')$, $\tilde{Q}(y, y') := \mathbb{P}(h(x) = y) \cdot \mathbb{P}(\tilde{Y} = y')$. Denote $x := \frac{\mathbb{P}(h(x) = y, \tilde{Y} = y')}{\mathbb{P}(h(x) = y) \cdot \mathbb{P}(\tilde{Y} = y')} - 1 = \frac{\tilde{P}(y, y')}{\tilde{Q}(y, y')} - 1$. Next we prove $\text{Bias}_f(h, \tilde{g}^*) = O(x^2)$ for different f -divergences.

Because $\text{Bias}_f(h, \tilde{g}^*) := \sum_{j=1}^K e_j \cdot \Delta_f^j(h, \tilde{g}^*) = \sum_{j=1}^K e_j \cdot \mathbb{E}_X[\tilde{g}^*(h(X), y) - f^*(\tilde{g}^*(h(X), y))]$, we analyze each of the term in expectation: $\tilde{g}^*(y, y') - f^*(\tilde{g}^*(y, y'))$.

Jenson-Shannon For Jenson-Shannon divergence, we have:

$$\begin{aligned} \tilde{g}^*(y, y') - f^*(\tilde{g}^*(y, y')) &= \tilde{g}^*(y, y') + \log(2 - e^{\tilde{g}^*(y, y')}) \\ &= \log \frac{2\tilde{P}(y, y')}{\tilde{P}(y, y') + \tilde{Q}(y, y')} + \log \left(2 - e^{\log \frac{2\tilde{P}(y, y')}{\tilde{P}(y, y') + \tilde{Q}(y, y')}} \right) \\ &= \log \frac{2\tilde{P}(y, y')}{\tilde{P}(y, y') + \tilde{Q}(y, y')} + \log \frac{2\tilde{Q}(y, y')}{\tilde{P}(y, y') + \tilde{Q}(y, y')} \\ &= \log \frac{2}{1 + \frac{\tilde{Q}(y, y')}{\tilde{P}(y, y')}} + \log \frac{2}{1 + \frac{\tilde{P}(y, y')}{\tilde{Q}(y, y')}} \\ &= \log \frac{2}{1 + \frac{1}{x+1}} + \log \frac{2}{1 + x + 1} \\ &= \log \frac{4}{2 + x + 1 + \frac{1}{x+1}} \end{aligned}$$

Using Taylor expansion we know

$$\log \frac{4}{2 + x + 1 + \frac{1}{x+1}} = 0 + x \cdot \left(-\frac{4 \cdot (1 - \frac{1}{(x+1)^2})}{2 + x + 1 + \frac{1}{x+1}} \right) \Big|_{x=0} + O(x^2) = O(x^2).$$

Squared Hellinger

$$\begin{aligned} \tilde{g}^*(y, y') - f^*(\tilde{g}^*(y, y')) &= \tilde{g}^*(y, y') - \frac{\tilde{g}^*(y, y')}{1 - \tilde{g}^*(y, y')} = \frac{\tilde{g}^{*2}(y', y)}{1 - \tilde{g}^*(y, y')} \\ &= \sqrt{\frac{\tilde{P}(h(x) = y, \tilde{Y} = y')}{\tilde{Q}(h(x) = y, \tilde{Y} = y')}} + \sqrt{\frac{\tilde{Q}(h(x) = y, \tilde{Y} = y')}{\tilde{P}(h(x) = y, \tilde{Y} = y')}} - 2 \\ &= \sqrt{1+x} + \sqrt{\frac{1}{1+x}} - 2. \end{aligned}$$

$$\sqrt{1+x} + \sqrt{\frac{1}{1+x}} - 2 = \left[1 + \frac{1}{2 \cdot \sqrt{1+x}} + 1 - \frac{1}{2}(1+x)^{-1.5} \right] \Big|_{x=0} - 2 + O(x^2) = O(x^2).$$

Pearson χ^2

$$\begin{aligned} \tilde{g}^*(y, y') - f^*(\tilde{g}^*(y, y')) &= \tilde{g}^*(y, y') - \tilde{g}^*(y, y') - \frac{1}{4}(\tilde{g}^*(y, y'))^2 \\ &= - \left(\frac{\tilde{P}(y, y')}{\tilde{Q}(y, y')} - 1 \right)^2 = -x^2 = O(x^2). \end{aligned}$$

Neyman χ^2

$$\begin{aligned} \tilde{g}^*(y, y') - f^*(\tilde{g}^*(y, y')) &= \tilde{g}^*(y, y') - 2 - 2\sqrt{1 - \tilde{g}^*(y, y')} \\ &= - \left(\frac{\tilde{Q}(y, y')}{\tilde{P}(y, y')} - 1 \right)^2 = - \left(\frac{1}{1+x} - 1 \right)^2 = O(x^2). \end{aligned}$$

KL

$$\begin{aligned} \tilde{g}^*(y, y') - f^*(\tilde{g}^*(y, y')) &= \tilde{g}^*(y, y') - e^{\tilde{g}^*(y, y') - 1} \\ &= 1 + \log \frac{\tilde{P}(h(x) = y, \tilde{Y} = y')}{\tilde{Q}(h(x) = y, \tilde{Y} = y')} - \frac{\tilde{P}(h(x) = y, \tilde{Y} = y')}{\tilde{Q}(h(x) = y, \tilde{Y} = y')} \end{aligned}$$

$$1 + \log(1 + x) - (1 + x) = \left(\frac{1}{1+x} - 1 \right) \Big|_{x=0} + O(x^2) = O(x^2).$$

Reverse KL

$$\begin{aligned} \tilde{g}^*(y, y') - f^*(\tilde{g}^*(y, y')) &= \tilde{g}^*(y, y') + 1 + \log(-\tilde{g}^*(y, y')) \\ &= 1 - \frac{\tilde{Q}(h(x) = y, \tilde{Y} = y')}{\tilde{P}(h(x) = y, \tilde{Y} = y')} + \log \frac{\tilde{Q}(h(x) = y, \tilde{Y} = y')}{\tilde{P}(h(x) = y, \tilde{Y} = y')} \end{aligned}$$

$$1 + \log(1 + x)^{-1} - (1 + x)^{-1} = 1 + \left[-\frac{1}{1+x} - (1 - \frac{1}{(1+x)^2}) \right] \Big|_{x=0} + O(x^2) = O(x^2)$$

Jeffrey

$$1 - \tilde{g}^*(y, y') = \frac{\tilde{Q}(h(x) = y, \tilde{Y} = y')}{\tilde{P}(h(x) = y, \tilde{Y} = y')} - \log \frac{\tilde{P}(h(x) = y, \tilde{Y} = y')}{\tilde{Q}(h(x) = y, \tilde{Y} = y')} = \frac{1}{1+x} - \log(1+x)$$

And

$$\begin{aligned} \tilde{g}^*(y, y') &= \frac{x}{1+x} + \log(1+x) \\ \tilde{g}^*(y, y') - f^*(\tilde{g}^*(y, y')) &= \tilde{g}^*(y, y') - W(e^{1-\tilde{g}^*(y, y')}) - \frac{1}{W(e^{1-\tilde{g}^*(y, y')})} - \tilde{g}^*(y, y') + 2 \\ &= 2 - W(e^{1-\tilde{g}^*(y, y')}) - \frac{1}{W(e^{1-\tilde{g}^*(y, y')})} \\ &= 2 - W(e^{1-\tilde{g}^*(y, y')}) \Big|_{x=0} - W'(e^{1-\tilde{g}^*(y, y')}) \Big|_{x=0} \cdot x - \frac{1}{W(e^{1-\tilde{g}^*(y, y')})} \Big|_{x=0} \\ &\quad - \left[\frac{1}{W(e^{1-\tilde{g}^*(y, y')})} \right]' \Big|_{x=0} \cdot x + O(x^2) \\ &= -W'(e^{1-\tilde{g}^*(y, y')}) \Big|_{x=0} \cdot x - \left[\frac{1}{W(e^{1-\tilde{g}^*(y, y')})} \right]' \Big|_{x=0} \cdot x + O(x^2) \\ &= O(x^2) \end{aligned}$$

□

A.10 PROOF OF THEOREM 7: \mathcal{H} -ROBUSTNESS OF TOTAL-VARIATION

Proof. We present the binary derivation but it extends easily to the multi-class case.

For TV, since $f(v) = \frac{1}{2}|v - 1|$, $f^*(u) = u$, we immediately have $\forall y' g^*(h = y', y) - f^*(g(h = y', y)) = 0$ and therefore

$$\Delta_f^y(h, g) = \mathbb{E}_X[g(h(X), y)] - \mathbb{E}_X[f^*(g(h(X), y))] \equiv 0, \forall y$$

and further for the binary case

$$\text{Bias}_f(h, g) := e_+ \cdot \Delta_f^{-1}(h, g) + e_- \cdot \Delta_f^{+1}(h, g) = 0$$

and for the multi-class case

$$\text{Bias}_f(h, g) = \sum_j e_j \Delta_f^j(h, g) = 0.$$

Therefore $\text{Bias}_f(h, g) \equiv 0$. We then know TV is \mathcal{H} -robust for an arbitrary \mathcal{H} using Theorem 6. TV's robustness can also be derived straightforwardly for binary classification:

$$\begin{aligned}
& \sup_g \mathbb{E}_{\tilde{Z} \sim \tilde{P}} [g(\tilde{Z})] - \mathbb{E}_{\tilde{Z} \sim \tilde{Q}} [f^*(g(\tilde{Z}))] \\
&= \sup_{|g| \leq 1/2} \left\{ (1 - e_+ - e_-) [\mathbb{E}_{Z \sim P} [g(Z)] - \mathbb{E}_{Z \sim Q} [f^*(g(Z))]] \right. \\
&\quad + \mathbb{E} [e_+ \cdot g(h(X), -1) + e_- \cdot g(h(X), +1)] \\
&\quad \left. - \mathbb{E} [e_+ \cdot f^*(g(h(X), -1)) + e_- \cdot f^*(g(h(X), +1))] \right\} \\
&= \sup_{|g| \leq 1/2} \left\{ (1 - e_+ - e_-) [\mathbb{E}_{Z \sim P} [g(Z)] - \mathbb{E}_{Z \sim Q} [f^*(g(Z))]] \right. \\
&\quad + \mathbb{E} [e_+ \cdot g(h(X), -1) + e_- \cdot g(h(X), +1)] \\
&\quad \left. - \mathbb{E} [e_+ \cdot g(h(X), -1) + e_- \cdot g(h(X), +1)] \right\} \\
&= \sup_{|g| \leq 1/2} (1 - e_+ - e_-) [\mathbb{E}_{Z \sim P} [g(Z)] - \mathbb{E}_{Z \sim Q} [f^*(g(Z))]] \\
&= (1 - e_+ - e_-) D_f(P_{h \times Y} \| Q_{h \times Y})
\end{aligned}$$

That is for total variation, minimizing the f -divergence between $h(X)$ and \tilde{Y} is the same as minimizing the f -divergence between $h(X)$ and the clean distribution Y . The above proof generalizes to multi-class easily. For example for the uniform diagonal noise, we have

$$\Delta_f^j(h, g) = \sum_{j=1}^K e_j \cdot [\mathbb{E}_X [g(h(X), \tilde{Y} = j)] - \mathbb{E}_X [f^*(g(h(X), \tilde{Y} = j))]] = 0.$$

Similar argument holds for sparse noise too. □

A.11 WHEN f -DIVERGENCE MEASURE IS \mathcal{H}^* -ROBUST?

Theorem 11. *For binary classification, suppose $\Delta_f^y(h, g)$ has the following form:*

$$\Delta_f^y(h, \tilde{g}^*) = w_h \cdot t(\text{FIT}(h = y, \tilde{Y} = y)) + (1 - w_h) \cdot t(\text{FIT}(h = -y, \tilde{Y} = y)) \quad (8)$$

$$\Delta_f^y(h_f^*, g^*) = w_{h_f^*} \cdot t(\text{FIT}(h_f^* = y, Y = y)) + (1 - w_{h_f^*}) \cdot t(\text{FIT}(h_f^* = -y, Y = y)) \quad (9)$$

where $w_h, w_{h_f^*} \in [0, 1]$. When $t(x)$ is monotonically decreasing as a function of $|x - 1|$ on both sides of $[1, \infty]$ and $(-\infty, 1)$, then $\text{Bias}_f(h_f^*, g^*) \geq \max_{h \in \mathcal{H}^*} \text{Bias}_f(h, \tilde{g}^*)$. Further, according to Theorem 6, the corresponding f -divergence measure is \mathcal{H}^* -robust.

Proof. For binary case, when $\mathbb{P}(Y = +1) = \mathbb{P}(Y = -1)$ and $e_+ = e_-$, we have $\mathbb{P}(\tilde{Y} = +1) = \mathbb{P}(\tilde{Y} = -1)$.

Denote $\mathcal{H}_-^* := \{h \in \mathcal{H} : \max_{y'} \text{FIT}(h(X) = -y', \tilde{Y} = y') \leq \min_y \text{FIT}(h_f^* = -y, Y = y)\}$, we first prove $\mathcal{H}_-^* \subseteq \mathcal{H}^*$. It is equivalent to prove, $\forall h \in \mathcal{H}_-^*, h \in \mathcal{H}^*$:

$$\begin{aligned}
\mathcal{H}_-^* &= \{h \in \mathcal{H} : \max_{y'} \text{FIT}(h(X) = -y', \tilde{Y} = y') \leq \min_y \text{FIT}(h_f^* = -y, Y = y)\} \\
&\subseteq \left\{ h \in \mathcal{H} : \max_{y'} \frac{\mathbb{P}(h(X) = -y' | \tilde{Y} = y')}{\mathbb{P}(h(X) = -y')} \leq \min_y \frac{\mathbb{P}(h_f^*(X) = -y | Y = y)}{\mathbb{P}(h_f^*(X) = -y)} \right\} \cup \{h_f^*\} \\
&\subseteq \left\{ h \in \mathcal{H} : \max_{y'} \frac{\mathbb{P}(h(X) = y' | \tilde{Y} = -y')}{\mathbb{P}(h(X) = y')} \leq \min_y \frac{\mathbb{P}(h_f^*(X) = y | Y = -y)}{\mathbb{P}(h_f^*(X) = y)} \right\} \cup \{h_f^*\} \\
&\subseteq \left\{ h \in \mathcal{H} : \max_{y'} \frac{\mathbb{P}(h(X) = y' | \tilde{Y} = -y') - \mathbb{P}(h(X) = y' | \tilde{Y} = y')}{\mathbb{P}(h(X) = y')} \right. \\
&\quad \left. \leq \min_y \frac{\mathbb{P}(h_f^*(X) = y | Y = -y) - \mathbb{P}(h_f^*(X) = y | Y = y)}{\mathbb{P}(h_f^*(X) = y)} \right\} \cup \{h_f^*\} \\
&\subseteq \left\{ h \in \mathcal{H} : 1 - \min_{y'} \frac{\mathbb{P}(h(X) = y' | \tilde{Y} = y')}{\mathbb{P}(h(X) = y')} \leq 1 - \max_y \frac{\mathbb{P}(h_f^*(X) = y | Y = y)}{\mathbb{P}(h_f^*(X) = y)} \right\} \cup \{h_f^*\} \\
&\subseteq \left\{ h \in \mathcal{H} : \min_{y'} \frac{\mathbb{P}(h(X) = y' | \tilde{Y} = y')}{\mathbb{P}(h(X) = y')} \geq \max_y \frac{\mathbb{P}(h_f^*(X) = y | Y = y)}{\mathbb{P}(h_f^*(X) = y)} \right\} \cup \{h_f^*\} \\
&\subseteq \mathcal{H}^* = \left\{ h \in \mathcal{H} : \min_{y'} \text{FIT}(h(X) = y', \tilde{Y} = y') \geq \max_y \text{FIT}(h_f^* = y, Y = y) \right\} \cup \{h_f^*\}
\end{aligned}$$

Since $t(x)$ is monotonically decreasing as a function of $|x - 1|$, for $h \in \mathcal{H}^*$,

$$\begin{aligned}
\text{Bias}_f(h_f^*, g^*) &= e_+ \cdot \sum_y \mathbb{P}(h_f^*(X) = y) \cdot t(\text{FIT}(h_f^*(X) = y, Y = +1)) \\
&\quad + e_- \cdot \sum_y \mathbb{P}(h_f^*(X) = y) \cdot t(\text{FIT}(h_f^*(X) = y, Y = -1)) \\
&= e_+ \cdot \sum_y \mathbb{P}(h_f^*(X) = y) \cdot t(\text{FIT}(h_f^*(X) = y, Y = y)) \cdot \mathbb{P}(Y = +1) \\
&\quad + e_+ \cdot \sum_y \mathbb{P}(h_f^*(X) = y) \cdot t(\text{FIT}(h_f^*(X) = y, Y = -y)) \cdot \mathbb{P}(Y = -1) \\
&\quad + e_- \cdot \sum_y \mathbb{P}(h_f^*(X) = y) \cdot t(\text{FIT}(h_f^*(X) = y, Y = y)) \cdot \mathbb{P}(Y = -1) \\
&\quad + e_- \cdot \sum_y \mathbb{P}(h_f^*(X) = y) \cdot t(\text{FIT}(h_f^*(X) = y, Y = -y)) \cdot \mathbb{P}(Y = +1) \\
&= \frac{e_+ + e_-}{2} \cdot \sum_y \mathbb{P}(h_f^*(X) = y) \cdot \left[t(\text{FIT}(h_f^*(X) = y, Y = y)) + t(\text{FIT}(h_f^*(X) = y, Y = -y)) \right] \\
&\geq \max_{h \in \mathcal{H}^*} \frac{e_+ + e_-}{2} \cdot \left[t(\max_y \text{FIT}(h_f^*(X) = y, Y = y)) + t(\min_y \text{FIT}(h_f^*(X) = y, Y = -y)) \right] \\
\max_{h \in \mathcal{H}^*} \text{Bias}_f(h, \tilde{g}^*) &= \max_{h \in \mathcal{H}^*} e_+ \cdot \sum_y \mathbb{P}(h(X) = y) \cdot t(\text{FIT}(h(X) = y, \tilde{Y} = +1)) \\
&\quad + e_- \cdot \sum_y \mathbb{P}(h(X) = y) \cdot t(\text{FIT}(h(X) = y, \tilde{Y} = -1)) \\
&= \max_{h \in \mathcal{H}^*} [e_+ \cdot \mathbb{P}(\tilde{Y} = +1) + e_- \cdot \mathbb{P}(\tilde{Y} = -1)] \cdot \sum_{\tilde{y}} \mathbb{P}(h(X) = \tilde{y}) \cdot t(\text{FIT}(h(X) = \tilde{y}, \tilde{Y} = \tilde{y})) \\
&\quad + [e_+ \cdot \mathbb{P}(\tilde{Y} = -1) + e_- \cdot \mathbb{P}(\tilde{Y} = +1)] \cdot \sum_{\tilde{y}} \mathbb{P}(h(X) = \tilde{y}) \cdot t(\text{FIT}(h(X) = \tilde{y}, \tilde{Y} = -\tilde{y})) \\
&\leq \max_{h \in \mathcal{H}^*} \frac{e_+ + e_-}{2} \cdot \left[t(\max_y \text{FIT}(h(X) = y, \tilde{Y} = y)) + t(\min_y \text{FIT}(h(X) = y, \tilde{Y} = -y)) \right]
\end{aligned}$$

Thus, $\text{Bias}_f(h_f^*, g^*) \geq \max_{h \in \mathcal{H}^*} \text{Bias}_f(h, \tilde{g}^*)$. According to Theorem 6, the corresponding f -divergence measure is \mathcal{H}^* -robust.

□

A.12 PROOF OF THEOREM 8: ROBUSTNESS OF D_f

Proof. Earlier we proved Theorem 11, next we show presented conditions in Eqn. (8) and (9) and $t(x)$ can be satisfied by the listed divergences:

The proof for $\Delta_f^y(h_f^*, g^*)$ can be viewed as a special case of $\Delta_f^y(h, \tilde{g}^*)$. The following derivations will therefore focus on $\Delta_f^y(h, \tilde{g}^*)$ and will not repeat for $\Delta_f^y(h_f^*, g^*)$.

Jenson-Shannon For Jenson-Shannon, we have $f^*(u) = -\log(2 - e^u)$, and

$$\begin{aligned}\tilde{g}^*(y, y') &= \log \frac{2 \cdot \mathbb{P}(h(X) = y, \tilde{Y} = y')}{\mathbb{P}(h(X) = y, \tilde{Y} = y') + \mathbb{Q}(h(X) = y, \tilde{Y} = y')} \\ &= \log \frac{2 \cdot \mathbb{P}(h(X) = y | \tilde{Y} = y')}{\mathbb{P}(h(X) = y | \tilde{Y} = y') + \mathbb{P}(h(X) = y)}\end{aligned}$$

Therefore,

$$\begin{aligned}\Delta_f^y(h, g) &= \mathbb{P}(h(X) = -y) \cdot \log \frac{4 \cdot \text{FIT}(h(X) = -y, \tilde{Y} = y)}{(1 + \text{FIT}(h(X) = -y, \tilde{Y} = y))^2} \\ &\quad + \mathbb{P}(h(X) = y) \cdot \log \frac{4 \cdot \text{FIT}(h(X) = y, \tilde{Y} = y)}{(1 + \text{FIT}(h(X) = y, \tilde{Y} = y))^2}\end{aligned}$$

$t(x) = \log \frac{4x}{(1+x)^2}$ satisfies the requirement specified in Theorem 11.

Squared-Hellinger For Squared-Hellinger, we have $f^*(u) = \frac{u}{1-u}$, and

$$\tilde{g}^*(y, y') = 1 - \sqrt{\frac{\mathbb{P}(h(X) = y)}{\mathbb{P}(h(X) = y | \tilde{Y} = y)}}$$

Therefore

$$\begin{aligned}\Delta_f^y(h, g) &= \mathbb{P}(h(X) = -y) \cdot \left[2 - \sqrt{\text{FIT}(h(X) = -y, \tilde{Y} = y)} - \frac{1}{\sqrt{\text{FIT}(h(X) = -y, \tilde{Y} = y)}} \right] \\ &\quad + \mathbb{P}(h(X) = y) \cdot \left[2 - \sqrt{\text{FIT}(h(X) = y, \tilde{Y} = y)} - \frac{1}{\sqrt{\text{FIT}(h(X) = y, \tilde{Y} = y)}} \right]\end{aligned}$$

Clearly $t(x) = 2 - \sqrt{x} - \frac{1}{\sqrt{x}}$ satisfies the requirement specified in Theorem 11.

Pearson χ^2

$$\begin{aligned}\Delta_f^y(h, \tilde{g}^*) &= \mathbb{E}[\tilde{g}^*(h(X), y) - f^*(\tilde{g}^*(h(X), y))] \\ &= \mathbb{E}[\tilde{g}^*(h(X), y) - \tilde{g}^*(h(X), y) - \frac{1}{4}(\tilde{g}^*(h(X), y))^2] \\ &= -\frac{1}{4}\mathbb{E}[(\tilde{g}^*(h(X), y))^2]\end{aligned}$$

$$\begin{aligned}
\Delta_f^y(h, \tilde{g}^*) &= -\mathbb{P}(h(X) = y) \cdot \left(\frac{\mathbb{P}(h(X) = y, \tilde{Y} = y)}{\mathbb{P}(h(X) = y) \cdot \mathbb{P}(\tilde{Y} = y)} - 1 \right)^2 \\
&\quad - \mathbb{P}(h(X) = -y) \cdot \left(\frac{\mathbb{P}(h(X) = -y, \tilde{Y} = y)}{\mathbb{P}(h(X) = -y) \cdot \mathbb{P}(\tilde{Y} = y)} - 1 \right)^2 \\
&= -\mathbb{P}(h(X) = y) \cdot (\text{FIT}(h(X) = y, \tilde{Y} = y) - 1)^2 \\
&\quad - \mathbb{P}(h(X) = -y) \cdot (\text{FIT}(h(X) = -y, \tilde{Y} = y) - 1)^2
\end{aligned}$$

Correspondingly $t(x) = -(x - 1)^2$, which satisfies the requirement specified in Theorem 11.

Neyman \mathcal{X}^2 For Neyman \mathcal{X}^2 we have $f^*(u) = 2 - 2\sqrt{1-u}$, and

$$\begin{aligned}
\Delta_f^y(h, \tilde{g}^*) &= \mathbb{E}[\tilde{g}^*(h(X), y) - f^*(\tilde{g}^*(h(X), y))] \\
&= \mathbb{E}[\tilde{g}^*(h(X), y) + 2\sqrt{1 - \tilde{g}^*(h(X), y)} - 2]
\end{aligned}$$

Since

$$\tilde{g}^*(y, y') = 1 - \left(\frac{Q(h(X) = y, \tilde{Y} = y')}{P(h(X) = y, \tilde{Y} = y')} \right)^2 = 1 - \left(\frac{\mathbb{P}(h(X) = y)}{\mathbb{P}(h(X) = y | \tilde{Y} = y')} \right)^2$$

Therefore

$$\tilde{g}^*(h(X) = y, y') + 2\sqrt{1 - \tilde{g}^*(h(X) = y, y')} - 2 = - \left(\frac{\mathbb{P}(h(X) = y)}{\mathbb{P}(h(X) = y | \tilde{Y} = y')} - 1 \right)^2$$

And further we have

$$\begin{aligned}
\Delta_f^y(h, \tilde{g}^*) &= -\mathbb{P}(h(X) = y) \cdot \left(\frac{\mathbb{P}(h(X) = y)}{\mathbb{P}(h(X) = y | \tilde{Y} = y)} - 1 \right)^2 \\
&\quad - \mathbb{P}(h(X) = -y) \cdot \left(\frac{\mathbb{P}(h(X) = -y)}{\mathbb{P}(h(X) = -y | \tilde{Y} = y)} - 1 \right)^2 \\
&= -\mathbb{P}(h(X) = y) \cdot (\text{FIT}(h(X) = y, \tilde{Y} = y)^{-1} - 1)^2 \\
&\quad - \mathbb{P}(h(X) = -y) \cdot (\text{FIT}(h(X) = -y, \tilde{Y} = y)^{-1} - 1)^2
\end{aligned}$$

$t(x) = -(x^{-1} - 1)^2$ satisfies the requirement specified in Theorem 11.

KL For KL, we have $f^*(u) = e^{u-1}$, and

$$\tilde{g}^*(y, y') = 1 + \log \frac{\mathbb{P}(h(X) = y | \tilde{Y} = y')}{\mathbb{P}(h(X) = y')}$$

Therefore

$$\begin{aligned}
\Delta_f^y(h, g) &= \mathbb{P}(h(X) = -y) \cdot [1 + \log \text{FIT}(h(X) = -y, \tilde{Y} = y) - \text{FIT}(h(X) = -y, \tilde{Y} = y)] \\
&\quad + \mathbb{P}(h(X) = y) \cdot [1 + \log \text{FIT}(h(X) = y, \tilde{Y} = y) - \text{FIT}(h(X) = y, \tilde{Y} = y)]
\end{aligned}$$

Clearly $t(x) = 1 + \log x - x$ satisfies the requirement specified in Theorem 11.

Reverse-KL For Reverse-KL, we have $f^*(u) = -1 - \log(-u)$, and

$$\tilde{g}^*(y, y') = -\log \frac{\mathbb{P}(h(X) = y)}{\mathbb{P}(h(X) = y | \tilde{Y} = y)}$$

Therefore

$$\begin{aligned}\Delta_f^y(h, g) &= \mathbb{P}(h(X) = -y) \cdot \left[1 - \log \text{FIT}(h(X) = -y, \tilde{Y} = y) - \frac{1}{\text{FIT}(h(X) = -y, \tilde{Y} = y)} \right] \\ &\quad + \mathbb{P}(h(X) = y) \cdot \left[1 - \log \text{FIT}(h(X) = y, \tilde{Y} = y) - \frac{1}{\text{FIT}(h(X) = y, \tilde{Y} = y)} \right]\end{aligned}$$

Clearly $t(x) = 1 - \log x - \frac{1}{x}$ satisfies the requirement specified in Theorem 11. \square

B SUPPLEMENTARY EXPERIMENT RESULTS

In our experiment settings, D_f measures fail to work well on almost all sparse high noise setting. This is largely due to the super unbalanced noisy labels, e.g., for each pair, the ratio of samples between the two classes is in the range of $[\frac{1}{3}, \frac{1}{2}]$.

B.1 SUPPLEMENTARY TABLE OF TABLE 3: METHODS COMPARISON WITHOUT BIAS CORRECTION

In Table 7, the performance of Pearson χ^2 , Jeffrey divergence on MNIST, Fashion MNIST, CIFAR-10 and CIFAR-100 are included in Table 7.

Dataset	Noise	CE	BLC	FLC	DMI	PL	Pearson	Jeffrey
MNIST	Sparse, Low	97.21	95.23	97.37	97.76	98.59	99.24(99.07±0.16)	99.24(99.11±0.08)
	Sparse, High	48.55	55.86	49.67	49.61	60.27	58.63(58.58±0.05)	49.21(49.17±0.04)
	Uniform, Low	97.14	94.27	95.51	97.72	99.06	99.13(99.03±0.09)	99.14(99.06±0.05)
	Uniform, High	93.25	85.92	87.75	95.50	97.77	97.89(97.76±0.10)	97.94(97.80±0.12)
	Random (0.2)	98.26	97.46	97.61	98.82	99.25	99.28(99.27±0.02)	99.29(99.22±0.11)
	Random (0.7)	97.00	93.52	87.74	95.47	98.52	98.70(98.54±0.10)	98.67(98.53±0.15)
Fashion MNIST	Sparse, Low	84.36	86.02	88.15	85.65	88.32	88.93(88.81±0.11)	88.96(88.68±0.21)
	Sparse, High	43.33	46.97	47.63	47.16	51.92	44.62(44.36±0.21)	45.57(45.26±0.28)
	Uniform, Low	82.98	84.48	86.58	83.69	89.31	87.27(87.15±0.09)	88.13(87.85±0.18)
	Uniform, High	79.52	78.10	82.41	77.94	84.69	85.30(85.26±0.06)	84.92(84.63±0.26)
	Random (0.2)	85.47	83.40	77.61	86.21	89.78	89.65(89.44±0.21)	89.74(89.33±0.29)
	Random (0.7)	82.05	78.41	73.42	80.89	87.22	86.72(86.29±0.32)	87.21(87.19±0.04)
CIFAR-10	Sparse, Low	87.20	72.96	76.17	92.32	91.35	91.40(91.24±0.27)	91.55(91.24±0.15)
	Sparse, High	61.81	56.30	66.12	27.94	69.70	46.36(46.27±0.07)	46.21(45.78±0.27)
	Uniform, Low	85.68	72.73	77.12	90.39	91.70	92.37(92.27±0.07)	92.17(92.02±0.08)
	Uniform, High	71.38	54.41	64.22	82.68	83.42	83.61(83.09±0.38)	83.80(83.73±0.05)
	Random (0.5)	78.40	59.31	68.97	85.06	86.47	86.03(85.56±0.32)	86.04(85.75±0.19)
	Random (0.7)	68.26	38.59	54.39	77.91	57.81	76.92 (76.82±0.07)	79.46(79.08±0.25)
CIFAR-100	Uniform	63.87	51.40	60.04	64.39	67.94	68.42(68.16±0.16)	68.86(68.63±0.18)
	Sparse	40.45	36.57	43.39	40.53	44.25	37.54(37.50±0.07)	37.43(37.06±0.25)
	Random (0.2)	65.84	61.21	61.52	66.23	62.92	69.90(69.72±0.15)	69.56(69.42±0.14)
	Random (0.5)	56.92	22.21	55.88	56.06	49.62	60.81(60.36±0.26)	60.95(60.73±0.15)

Table 7: Experiment results comparison (w/o bias correction): The best performance in each setting (row) is highlighted in blue. All f -divergences will be highlighted if they are better than the baselines we compare to. We report the maximum accuracy of each D_f measures along with (mean \pm standard deviation).

B.2 D_f MEASURES WITH BIAS CORRECTION ON MNIST

In Table 8, we test the impact of bias correction on MNIST with 4 noise settings. Except for the spare high noise setting which is a huge challenge for all implemented methods, experiment results of other 3 noise settings further demonstrate the negligible effect of bias term in the optimization of D_f measures.

Noise	J-S	Gap	PS	Gap	KL	Gap	Jeffrey	Gap
Sparse, Low	98.88(98.82±0.06)	-0.27	99.05(98.98±0.05)	-0.19	99.29(99.19±0.09)	+0.08	99.13(99.06±0.06)	-0.11
Sparse, High	21.39(21.36±0.03)	-37.54	49.22(49.17±0.05)	-9.41	49.07(49.05±0.02)	-0.07	49.14(49.06±0.09)	-0.07
Uniform, Low	99.18(99.10±0.05)	+0.05	99.13(99.01±0.10)	+0.00	99.30(99.24±0.08)	+0.20	99.20(99.12±0.09)	+0.06
Uniform, High	97.76(97.68±0.07)	-0.10	97.72(97.65±0.06)	-0.17	97.91(97.74±0.14)	-0.23	98.16(97.98±0.14)	+0.22

Table 8: D_f measures with bias correction on MNIST: Digits highlighted in **blue** means better than baseline methods, in **red** means better than without bias correction. PS: Pearson.

B.3 D_f MEASURES WITH BIAS CORRECTION ON FASHION MNIST

In Table 9, we test the impact of bias correction on Fashion MNIST with 4 noise settings. We can reach the same conclusion on bias correction as MNIST.

Noise	J-S	Gap	PS	Gap	KL	Gap	Jeffrey	Gap
Sparse, Low	89.37(88.83±0.34)	+0.57	87.99(87.90±0.13)	-0.94	82.29(82.03±0.22)	-7.48	82.04(81.65±0.26)	-6.92
Sparse, High	X	X	39.02(38.85±0.10)	-5.60	46.98(46.23±0.62)	+8.02	38.94(38.75±0.15)	-6.63
Uniform, Low	88.98(88.61±0.26)	+0.40	87.72(87.66±0.06)	+0.45	89.04(88.78±0.18)	+0.72	89.05(88.87±0.15)	+0.92
Uniform, High	85.56(85.33±0.19)	-0.06	85.57(85.03±0.37)	+0.27	85.15(84.94±0.15)	-0.54	84.76(84.48±0.31)	-0.16

Table 9: D_f measures with bias correction on Fashion MNIST: Digits highlighted in **blue** means better than baseline methods, in **red** means better than without bias correction. X: experiment failed to stabilize. PS: Pearson.

C EXPERIMENT DETAILS

C.1 NOISE TRANSITION MATRIX FOR SECTION 5.2: ROBUSTNESS OF D_f MEASURES

We use CIFAR-10 dataset together with the uniform noise transition matrix to flip the noisy labels. In the following noise transition matrix, e is in $[0.00, 0.01, 0.02, \dots, 0.09]$ and the noise rate of each set of noisy labels is $9 * e$.

$$\begin{bmatrix} 1 - 9 * e & e & e & e & e & e & e & e & e \\ e & 1 - 9 * e & e & e & e & e & e & e & e \\ e & e & 1 - 9 * e & e & e & e & e & e & e \\ e & e & e & 1 - 9 * e & e & e & e & e & e \\ e & e & e & e & 1 - 9 * e & e & e & e & e \\ e & e & e & e & e & 1 - 9 * e & e & e & e \\ e & e & e & e & e & e & 1 - 9 * e & e & e \\ e & e & e & e & e & e & e & 1 - 9 * e & e \\ e & e & e & e & e & e & e & e & 1 - 9 * e \\ e & e & e & e & e & e & e & e & e \end{bmatrix}$$

C.2 NOISE TRANSITION MATRIX FOR MNIST AND FASHION MNIST DATASET

Sparse-low noise matrix:

$$\begin{bmatrix} 0.7 & 0.3 & 0. & 0. & 0. & 0. & 0. & 0. & 0. \\ 0.2 & 0.8 & 0. & 0. & 0. & 0. & 0. & 0. & 0. \\ 0. & 0. & 0.7 & 0.3 & 0. & 0. & 0. & 0. & 0. \\ 0. & 0. & 0.2 & 0.8 & 0. & 0. & 0. & 0. & 0. \\ 0. & 0. & 0. & 0. & 0.7 & 0.3 & 0. & 0. & 0. \\ 0. & 0. & 0. & 0. & 0.2 & 0.8 & 0. & 0. & 0. \\ 0. & 0. & 0. & 0. & 0. & 0. & 0.7 & 0.3 & 0. \\ 0. & 0. & 0. & 0. & 0. & 0. & 0.7 & 0.3 & 0. \\ 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0.7 & 0.3 \\ 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0.8 \end{bmatrix}$$

Sparse-high noise matrix:

$$\begin{bmatrix} 0.3 & 0.7 & 0. & 0. & 0. & 0. & 0. & 0. & 0. \\ 0.2 & 0.8 & 0. & 0. & 0. & 0. & 0. & 0. & 0. \\ 0. & 0. & 0.3 & 0.7 & 0. & 0. & 0. & 0. & 0. \\ 0. & 0. & 0.2 & 0.8 & 0. & 0. & 0. & 0. & 0. \\ 0. & 0. & 0. & 0. & 0.3 & 0.7 & 0. & 0. & 0. \\ 0. & 0. & 0. & 0. & 0.2 & 0.8 & 0. & 0. & 0. \\ 0. & 0. & 0. & 0. & 0. & 0.3 & 0.7 & 0. & 0. \\ 0. & 0. & 0. & 0. & 0. & 0.2 & 0.8 & 0. & 0. \\ 0. & 0. & 0. & 0. & 0. & 0. & 0.7 & 0.3 & 0.7 \\ 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0.2 & 0.8 \end{bmatrix}$$

Uniform-low noise matrix:

0.258	0.075	0.09	0.085	0.07	0.082	0.077	0.091	0.092	0.08
0.08	0.253	0.09	0.085	0.07	0.082	0.077	0.091	0.092	0.08
0.08	0.075	0.268	0.085	0.07	0.082	0.077	0.091	0.092	0.08
0.08	0.075	0.09	0.263	0.07	0.082	0.077	0.091	0.092	0.08
0.08	0.075	0.09	0.085	0.248	0.082	0.77	0.91	0.92	0.08
0.08	0.075	0.09	0.085	0.07	0.26	0.077	0.091	0.092	0.08
0.08	0.075	0.09	0.085	0.07	0.082	0.255	0.091	0.092	0.08
0.08	0.075	0.09	0.085	0.07	0.082	0.077	0.269	0.092	0.08
0.08	0.075	0.09	0.085	0.07	0.082	0.077	0.091	0.27	0.08
0.08	0.075	0.09	0.085	0.07	0.082	0.077	0.091	0.092	0.258

Uniform-high noise matrix:

0.58	0.045	0.047	0.055	0.053	0.022	0.068	0.054	0.056	0.02
0.05	0.0575	0.047	0.055	0.053	0.022	0.068	0.054	0.056	0.02
0.05	0.045	0.577	0.055	0.053	0.022	0.068	0.054	0.056	0.02
0.05	0.045	0.047	0.585	0.053	0.022	0.068	0.054	0.056	0.02
0.05	0.045	0.047	0.055	0.583	0.022	0.068	0.054	0.056	0.02
0.05	0.045	0.047	0.055	0.053	0.552	0.068	0.054	0.056	0.02
0.05	0.045	0.047	0.055	0.053	0.022	0.598	0.054	0.056	0.02
0.05	0.045	0.047	0.055	0.053	0.022	0.068	0.584	0.056	0.02
0.05	0.045	0.047	0.055	0.053	0.022	0.068	0.054	0.586	0.02
0.05	0.045	0.047	0.055	0.053	0.022	0.068	0.054	0.056	0.55

Random 0.2 noise matrix:

Random 0.7 noise matrix:

0.36	0.07	0.08	0.07	0.08	0.07	0.07	0.07	0.07	0.07	0.07
0.06	0.39	0.07	0.06	0.07	0.07	0.07	0.07	0.07	0.07	0.07
0.07	0.07	0.38	0.08	0.07	0.07	0.07	0.07	0.07	0.07	0.07
0.08	0.07	0.07	0.36	0.07	0.08	0.07	0.07	0.07	0.07	0.07
0.07	0.07	0.07	0.07	0.37	0.07	0.07	0.07	0.07	0.07	0.07
0.07	0.07	0.07	0.08	0.06	0.37	0.07	0.07	0.07	0.07	0.07
0.07	0.07	0.07	0.07	0.07	0.07	0.38	0.07	0.07	0.07	0.07
0.07	0.06	0.07	0.07	0.07	0.07	0.07	0.38	0.07	0.07	0.07
0.06	0.07	0.07	0.07	0.08	0.07	0.07	0.07	0.37	0.07	0.07
0.07	0.07	0.06	0.07	0.07	0.07	0.08	0.07	0.07	0.07	0.37

C.3 NOISE TRANSITION MATRIX FOR CIFAR-10 DATASET

Sparse-low noise matrix:

Sparse-high noise matrix:

Uniform-low noise matrix:

$$\begin{bmatrix} 0.82 & 0.03 & 0.01 & 0.023 & 0.017 & 0.022 & 0.021 & 0.018 & 0.019 & 0.02 \\ 0.02 & 0.83 & 0.01 & 0.023 & 0.017 & 0.022 & 0.021 & 0.018 & 0.019 & 0.02 \\ 0.02 & 0.03 & 0.81 & 0.023 & 0.017 & 0.022 & 0.021 & 0.018 & 0.019 & 0.02 \\ 0.02 & 0.03 & 0.01 & 0.823 & 0.017 & 0.022 & 0.021 & 0.018 & 0.019 & 0.02 \\ 0.02 & 0.03 & 0.01 & 0.023 & 0.817 & 0.022 & 0.021 & 0.018 & 0.019 & 0.02 \\ 0.02 & 0.03 & 0.01 & 0.023 & 0.017 & 0.822 & 0.021 & 0.018 & 0.019 & 0.02 \\ 0.02 & 0.03 & 0.01 & 0.023 & 0.017 & 0.022 & 0.821 & 0.018 & 0.019 & 0.02 \\ 0.02 & 0.03 & 0.01 & 0.023 & 0.017 & 0.022 & 0.021 & 0.818 & 0.019 & 0.02 \\ 0.02 & 0.03 & 0.01 & 0.023 & 0.017 & 0.022 & 0.021 & 0.018 & 0.819 & 0.02 \\ 0.02 & 0.03 & 0.01 & 0.023 & 0.017 & 0.022 & 0.021 & 0.018 & 0.019 & 0.82 \end{bmatrix}$$

Uniform-high noise matrix:

$$\begin{bmatrix} 0.46 & 0.07 & 0.04 & 0.05 & 0.06 & 0.04 & 0.06 & 0.07 & 0.08 & 0.07 \\ 0.05 & 0.48 & 0.04 & 0.05 & 0.06 & 0.04 & 0.06 & 0.07 & 0.08 & 0.07 \\ 0.05 & 0.07 & 0.45 & 0.05 & 0.06 & 0.04 & 0.06 & 0.07 & 0.08 & 0.07 \\ 0.05 & 0.07 & 0.04 & 0.46 & 0.06 & 0.04 & 0.06 & 0.07 & 0.08 & 0.07 \\ 0.05 & 0.07 & 0.04 & 0.05 & 0.47 & 0.04 & 0.06 & 0.07 & 0.08 & 0.07 \\ 0.05 & 0.07 & 0.04 & 0.05 & 0.06 & 0.45 & 0.06 & 0.07 & 0.08 & 0.07 \\ 0.05 & 0.07 & 0.04 & 0.05 & 0.06 & 0.04 & 0.47 & 0.07 & 0.08 & 0.07 \\ 0.05 & 0.07 & 0.04 & 0.05 & 0.06 & 0.04 & 0.06 & 0.48 & 0.08 & 0.07 \\ 0.05 & 0.07 & 0.04 & 0.05 & 0.06 & 0.04 & 0.06 & 0.07 & 0.49 & 0.07 \\ 0.05 & 0.07 & 0.04 & 0.05 & 0.06 & 0.04 & 0.06 & 0.07 & 0.08 & 0.48 \end{bmatrix}$$

Random 0.5 noise matrix:

$$\begin{bmatrix} 0.55 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.05 & 0.56 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.05 & 0.05 & 0.55 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.05 & 0.05 & 0.06 & 0.54 & 0.05 & 0.05 & 0.05 & 0.04 & 0.06 & 0.06 \\ 0.05 & 0.05 & 0.05 & 0.05 & 0.56 & 0.05 & 0.05 & 0.05 & 0.04 & 0.05 \\ 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.54 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.04 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.55 & 0.05 & 0.05 & 0.05 \\ 0.04 & 0.04 & 0.05 & 0.05 & 0.06 & 0.05 & 0.04 & 0.56 & 0.05 & 0.05 \\ 0.06 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.55 & 0.05 \\ 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.05 & 0.55 \end{bmatrix}$$

Random 0.7 noise matrix:

$$\begin{bmatrix} 0.37 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 \\ 0.06 & 0.38 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.08 & 0.07 \\ 0.07 & 0.07 & 0.36 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.08 \\ 0.07 & 0.07 & 0.07 & 0.37 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 \\ 0.07 & 0.07 & 0.08 & 0.07 & 0.37 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 \\ 0.07 & 0.08 & 0.07 & 0.07 & 0.07 & 0.36 & 0.07 & 0.07 & 0.07 & 0.06 \\ 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.37 & 0.07 & 0.07 & 0.07 \\ 0.07 & 0.06 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.37 & 0.07 & 0.07 \\ 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.07 & 0.38 & 0.06 \\ 0.07 & 0.07 & 0.07 & 0.08 & 0.07 & 0.07 & 0.07 & 0.07 & 0.06 & 0.37 \end{bmatrix}$$

C.4 NOISE TRANSITION MATRIX FOR CIFAR-100 DATASET

For sparse noise matrix, we randomly divide 100 classes into 50 disjoint pairs, the flipping probability (T_{ji}, T_{ij}) in each pair is randomly chosen from $(0.05, 0.75), (0.1, 0.70), (0.15, 0.65), (0.2, 0.6)$.

C.5 PARAMETER SETTINGS ON NOISED DATASET

MNIST, Fashion MNIST, CIFAR-10 For experiments on MNIST and Fashion-MNIST datasets, we use the convolutional neural network used in DMI for DMI, PL and f -divergences. All the experiments are performed with batch size 128. PL and f -divergences adopt two kinds of learning rate setting and trained for 80 epochs, either with initial learning rate 5e-4 or 1e-3, then decay 0.2, 0.5, 0.2 every 20 epochs. We choose the default learning rate setting for DMI, BLC and FLC. For DMI’s convolutional neural network, Adam(Kingma & Ba (2014)) with default parameters is used as the optimizer, while for loss-correction’s fully-connected neural network case we use AdaGrad(Duchi et al. (2010)) in order to be consistent with their works.

CIFAR-10 and CIFAR-100 For all methods and both datasets, we unify the model to be an 18-layer PreAct Resnet (He et al. (2016)) and train it using SGD with a momentum of 0.9, a weight decay of 0.0005, and a batch size of 128. All methods firstly train with CE warm-up for 120 (CIFAR-10) or 240 (CIFAR-100) epochs on CIFAR-10 and CIFAR-100 respectively. For DMI, BLC and FLC, we use the default learning rate settings. For PL and f-divergences, we train 100 epochs after the warm-up with initial learning rate 0.01, and decays 0.1 every 30 epochs.

Clothing 1M For clothing 1M, we use pre-trained ResNet50, SGD optimizer with momentum 0.9 and weight decay 1e-3. The initial learning rate is 0.002. All mentioned f -divergences trained 40 epochs, after 10 epochs, the learning rate becomes 5e-5. Then it decays 0.2, 0.5 consequently for every 5 epochs. We compare with reported best result for all our baseline methods.

C.6 PARAMETER SETTINGS ON CLEAN DATASET

We adopt the same setting (except for the number of epochs and the learning rate setting) as used in the noised dataset for each dataset.

MNIST, Fashion MNIST For CE, we trained the model for 40 epochs. The initial learning rate is 5e-4, and it decays 0.2 after 20 epochs. For D_f measures, the learning rate setting is the same as that in the noised dataset.

CIFAR-10 For CE, we trained the model for 300 epochs. Learning rate is 0.1 for first 150 epochs. From 150-th epoch to 250-th epoch, the learning rate is 0.01. Then, 0.001 till the end. For D_f measures, we trained the model for 240 epochs. The initial learning rate is 0.1, and it decays 0.1 for every 60 epochs.

CIFAR-100 For CE, we trained the model for 200 epochs. Learning rate is 0.1 for first 60 epochs. From 61-th epoch to 120-th epoch, the learning rate is 0.02 (save the model at 120-th epoch as a warm-up model for D_f measures). From 121-th epoch to 160-th epoch, the learning rate is 0.004. Then, 0.0008 till the end. For D_f measures, we load pre-trained CE model and trained for another 100 epochs. The initial learning rate is 0.01, and it decays 0.1 for every 30 epochs.

C.7 COMPUTING INFRASTRUCTURE

In our experiments, we use a GPU cluster (8 TITAN V GPUs and 16 GeForce GTX 1080 GPUs) for training and evaluation.