
Statistical Hypothesis Testing for Class-Conditional Label Noise

Rafael Poyiadzi¹

Weisong Yang¹

Niall Twomey²

Raul Santos-Rodriguez¹

¹University of Bristol, England

²Cookpad Ltd, and University of Bristol, England (honorary)

Abstract

In this work we aim to provide machine learning practitioners with tools to answer the question: *is there class-conditional flipping noise in my labels?* In particular, we present hypothesis tests to reliably check whether a given dataset of instance-label pairs has been corrupted with *class-conditional label noise*. While previous works explore the direct estimation of the noise rates, this is known to be hard in practice and does not offer a real understanding of how trustworthy the estimates are. These methods typically require *anchor points* – examples whose true posterior is either 0 or 1. Differently, in this paper we assume we have access to a set of anchor points whose true posterior is approximately 1/2. The proposed hypothesis tests are built upon the asymptotic properties of Maximum Likelihood Estimators for Logistic Regression models and accurately distinguish the presence of class-conditional noise from uniform noise. We establish the main properties of the tests, including a theoretical and empirical analysis of the dependence of the power on the test on the training sample size, the number of anchor points, the difference of the noise rates and the use of realistic relaxed anchors.

1 INTRODUCTION

When a machine learning practitioner is presented with a new dataset, a first fundamental question is that of data quality (Lawrence [2017]) as this will determine any subsequent learning. However, in supervised settings, an equally important but often overlooked question is the verification of the quality of labels. For instance, in standard data collections, data curators usually rely on annotators from online platforms, where individual annotators cannot be uncondi-

tionally trusted as they have been shown to perform inconsistently (Jindal et al. [2016]). Labels are also expected to not be ideal in situations where the data is harvested from the web (Fergus et al. [2005], Schroff et al. [2011]). In general this is a product of annotations not being carried out by (perfect) domain experts. While the existing literature focuses on estimating the distortion(s) present in the labels (see Section 4), in this paper, we take a step back and our main contribution is the design and analysis of hypothesis testing procedures that would allow us (under certain assumptions we introduce later) to measure the likelihood that the labels of a provided dataset have been corrupted. With this information at hand, the practitioner can then make a more informed choice of the learning strategy better suited to the given dataset or even to revisit the data annotation process altogether.

Focusing on binary classification, the goal is to train a classifier $g : \mathcal{X} \rightarrow \{-1, +1\}$, from a labelled dataset $\mathcal{D}_n^{\text{train}} = \{(x_i, y_i)\}_{i=1}^n \in (\mathbb{R}^d \times \{-1, 1\})$, with the objective of achieving a low miss-classification error: $\mathbb{P}_{X,Y}(g(X) \neq Y)$. While it is generally assumed that the training dataset is drawn from the distribution for which we wish to minimise the error for $\mathcal{D}_n^{\text{train}} \sim p(X, Y)$, this is often not the case. Instead, the task requires us to train a classifier on a corrupted version of the dataset $\tilde{\mathcal{D}}_n^{\text{train}} \sim p(X, \tilde{Y})$ whilst still hoping to achieve a low error rate on the clean distribution. In order to narrow down the task, in this work we focus on one particular type of corruption that is pervasive in the examples described above, *instance-independent label noise*, where labels are flipped with a certain probability. This rate can either be uniform across the entire data-generating distribution or conditioned on the true class of the data point.

A key ingredient in our procedure is the input from the user in the form of a set of *anchor points*. Differently from previous works, we assume anchor points for which the true posterior distribution $\mathbb{P}(Y = 1 | X = x)$ is (approximately) 1/2. For a specific instance x , this requirement means that, given the current set of features, an expert would not be able to provide *any* help to identify the correct class label. While

this will be shown to be convenient for theoretical purposes, finding such anchor points might be rather difficult to accomplish in practice, so we show how to relax this notion to a more realistic $\eta(x) \approx 1/2$, which could be understood as *probably* carrying little information about the label. In a setting with multiple annotators, this could be understood as the team disagreeing.

Our approach is based on the asymptotic properties of the *Maximum Likelihood Estimate* (MLE) solution for Logistic Regression models, and the relationship between the true and noisy posteriors. On the theoretical side, we show that when the asymptotic properties of MLE hold and the user provides a single anchor point, we can devise hypothesis tests to assess the presence of class-conditional label corruption in the dataset. We then further extend this ideas to allow for richer sets of anchor points and illustrate how these lead to gains in the *power* of the tests.

In Section 2 we cover the necessary background on MLE asymptotics, noisy labels and define the necessary tools. In Section 3 we illustrate how to carry a z-test using anchor points on the presence of class-conditional noise. In Section 4 we discuss related work and in Section 5 we present experimental findings.

2 BACKGROUND

We assume we are provided with a dataset $(X, y) = \{(x_i, y_i)\}_{i=1}^N \in (\mathbb{R}^d \times \{-1, 1\})$, and our task is to assess whether the labels have been corrupted with class-conditional *flipping noise*. We use y to denote the true label, and \tilde{y} to denote the noisy label. In general, the *tilda* ($\tilde{\cdot}$) will be used to denote the noisy version of an object and *hat* ($\hat{\cdot}$) to denote an estimated version, unless specified otherwise. We assume the feature vectors (x) have been augmented with *ones* such that we have $x \rightarrow (1, x)$.

We assume the data is generated from the following process:

$$y_i \sim \text{Bernoulli}(\eta_i), \quad \eta_i = \sigma(\theta_0^\top x_i) = \frac{1}{1 + \exp(-\theta_0^\top x_i)}$$

And, following the Maximum Likelihood Estimation procedure we have:

$$\hat{\theta}_n := \operatorname{argmax}_{\theta \in \Theta} \ell_n(\theta | D_n) = \operatorname{argmax}_{\theta \in \Theta} \prod_{i=1}^n \ell_i(\theta | x_i, y_i)$$

where: $l(\theta | x_i, y_i) = \frac{y_i+1}{2} \cdot \log \eta_i + \frac{1-y_i}{2} \cdot \log(1 - \eta_i)$.

In this setting, the following can be shown (See for example Chapter 4 of Van der Vaart [2000]):

$$\sqrt{n}(\hat{\theta}_n - \theta_0) \xrightarrow{D} \mathcal{N}(0, I_n(\theta_0)^{-1}) \quad (1)$$

where I_{θ_0} denotes the Fisher-Information Matrix defined as:

$$I_n(\theta_0) = \mathbb{E}_\theta \left(-\frac{\partial^2 \ell_n(\theta; Y | x)}{\partial \theta \partial \theta^\top} \right) = \mathbb{E}_\theta (-H_n(\theta; Y | x))$$

where the expectation is with respect to the conditional distribution, and H_n denotes the Hessian matrix. An important remark is that we can express $I_n(\theta_0) = n \cdot I_i(\theta_0)$, which shows that as $n \rightarrow \infty$, the variance of the MLE estimator goes to 0.

We will consider two types of flipping noise, which we formally define below. In both cases the noise rates are independent of the instance: $\mathbb{P}(\tilde{Y} = -i | Y = i, X = x) = \mathbb{P}(\tilde{Y} = -i | Y = i)$.

Definition 2.1. Bounded Uniform Noise (UN)

In this setting the per-class noise rates are identical: $\mathbb{P}(\tilde{Y} = 1 | Y = -1) = \mathbb{P}(\tilde{Y} = -1 | Y = 1) = \tau$ and bounded: $\tau < 0.50$. We will denote this setting with $UN(\tau)$, and a dataset $\mathcal{D} = (X, y)$ inflicted by $UN(\tau)$ by: \mathcal{D}_τ .

Definition 2.2. Bounded Class-Conditional Noise (CCN)

In this setting the per-class noise rates are different, $\alpha \neq \beta$ and bounded $\alpha + \beta < 1$ with: $\mathbb{P}(\tilde{Y} = -1 | Y = 1) = \alpha$ and $\mathbb{P}(\tilde{Y} = 1 | Y = -1) = \beta$. We will denote this setting with $CCN(\alpha, \beta)$, and a dataset $\mathcal{D} = (X, y)$ inflicted by $CCN(\alpha, \beta)$ by: $\mathcal{D}_{\alpha, \beta}$.

An object of central interest in classification settings is the posterior predictive distribution: $\eta(x) = \mathbb{P}(Y = 1 | X = x)$. Its noisy counterpart, $\tilde{\eta}(x) = \mathbb{P}(\tilde{Y} = 1 | X = x)$, under the two settings, $UN(\tau)$ and $CCN(\alpha, \beta)$, can be expressed as,

$$\tilde{\eta}(x) = \begin{cases} (1 - \alpha - \beta) \cdot \eta(x) + \beta & \text{if (CCN)} \\ (1 - 2\tau) \cdot \eta(x) + \tau & \text{if (UN)} \end{cases} \quad (2)$$

See Appendix 7.1 for full derivation.

We consider loss functions that have the margin property: $\ell(y, f(x)) = \psi(yf(x))$, where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a scorer, and $g(x) = \text{sign}(f(x))$ is the predictor. Let $f^* = \operatorname{argmin}_{f \in \mathcal{F}} \mathbb{E}_{X, Y} \psi(Yf(X))$ and $\tilde{f}^* = \operatorname{argmin}_{f \in \mathcal{F}} \mathbb{E}_{X, \tilde{Y}} \psi(\tilde{Y}f(X))$ denote the minimisers under the clean and noisy distributions, under model-class \mathcal{F} .

Definition 2.3. Uniform Noise robustness (Ghosh et al. [2015])

Empirical risk minimization under loss function ℓ is said to be noise-tolerant if $\mathbb{P}_{X, Y}(g^(X) = Y) = \mathbb{P}_{X, Y}(\tilde{g}^*(X) = Y)$*

Theorem 2.1. Sufficient conditions for robustness to uniform noise

Under uniform noise $\tau < 0.50$, and a margin loss function, $\ell(y, f(x)) = \psi(yf(x))$ satisfying: $\psi(f(X)) + \psi(-f(X)) = K$ for a positive constant K , we have that $\tilde{g}^(x) = \text{sign}(\tilde{f}^*(x))$ obtained from: $\tilde{f}^* = \operatorname{argmin}_{f \in \mathcal{F}} \mathbb{E}_{X, \tilde{Y}} \psi(\tilde{Y}f(X))$ is robust to uniform noise.*

For the proof of the theorem, see Appendix 7.2.

Several loss functions satisfy this, such as: the *square*, *unhinged* (linear), *logistic*, and more. We now introduce our definition of anchor points¹.

Definition 2.4. (Anchor Points) *An instance x is called an anchor point if we are provided with its true posterior $\eta(x)$. Let \mathcal{A}_s^k denote a collection of k anchor points, with $\eta(x) = s \forall x \in \mathcal{A}_s^k$. Furthermore, let us also define $\mathcal{A}_{s,\delta}^k$, to imply that $\eta(x_i) = s + \varepsilon_i$, for $\varepsilon_i \sim \mathbb{U}([- \delta, \delta])$, with $0 \leq \delta \ll 1$ (respecting $0 \leq \eta(x) \leq 1$). Also let $\mathcal{A}_{s,\delta} = \mathcal{A}_{s,\delta}^1$.*

The three cases we will be referring to are as follows:

$$\begin{aligned} \mathcal{A}_1^k &\rightarrow \eta(x) = 1 &\rightarrow \tilde{\eta}(x) = 1 - \alpha \\ \mathcal{A}_{1/2}^k &\rightarrow \eta(x) = 1/2 &\rightarrow \tilde{\eta}(x) = \frac{1 - \alpha + \beta}{2} \\ \mathcal{A}_0^k &\rightarrow \eta(x) = 0 &\rightarrow \tilde{\eta}(x) = \beta \end{aligned}$$

The first and last examples, \mathcal{A}_1^k and \mathcal{A}_0^k , have been used in the past in different scenarios. In this work we will make use of the second case, $\mathcal{A}_{1/2}^k$. These instances should be understood as *providing no information* with regards to the label.

3 HYPOTHESIS TESTS BASED ON ANCHOR POINTS

In this section we propose a framework to devise hypothesis tests to examine the presence of class-conditional label noise in a given dataset, assuming we are provided with an anchor point(s).

Our procedure is based on a two-sided *z-test* (see for example Chapter 8 of Casella and Berger [2001]) with a simple null hypothesis, and a composite alternative hypothesis (Eq.5). We first define the distribution under the null hypothesis (Eq.6), and under the alternative hypothesis (Eq.7), when provided with one strict anchor point. In this setting, for a fixed *level of significance* (Type I error) (Eq.8), we first derive a region for retaining the null hypothesis (Eq.9), and then we analyse the *power* (Prop.3.1) of the test (where we have that Type II Error = 1 - *power*). Table 1 summarises these. We then extend the approach to examine scenarios that include: 1) having multiple strict anchors, 2) having multiple relaxed anchors, and 3) not having anchors.

With the application of the *delta method* (See for example Chapter 3 of Van der Vaart [2000]) on Eq.1, we can get an asymptotic distribution for the predictive posterior:

$$\sqrt{n}(\hat{\eta}(x) - \eta(x)) \xrightarrow{D} \mathcal{N}\left(0, [\eta(x)\{1 - \eta(x)\}]^2 \cdot x^\top I_{\theta_0}^{-1} x\right) \quad (3)$$

¹Different notions related to our definition of anchor points have been used before in the literature – under different names – we review their uses and assumptions in Section 4

	Retain \mathcal{H}_0	Reject \mathcal{H}_0
\mathcal{H}_0 True	Correct	Type I Error (Eq.8)
\mathcal{H}_0 False	Type II Error	Correct (Eq.10)

Table 1: Identifying Type I and Type II errors.

This fails in the case of $\eta(x) \in \{0, 1\}$, so instead we work with $1/2$. Which, together with the approximation of the Fisher-Information matrix with the empirical Hessian, we get:

$$\hat{\eta}(x) \xrightarrow{D} \mathcal{N}\left(\frac{1}{2}, \frac{1}{16} \cdot x^\top \hat{H}_n x\right) \quad (4)$$

where $\hat{H}_n = (X^\top DX)^{-1}$, where D is a diagonal matrix, with $D_{ii} = \hat{\eta}_i(1 - \hat{\eta}_i)$, where $\hat{\eta}_i = \sigma(x_i^\top \hat{\theta})$.

For the settings: $(\mathcal{D}, \mathcal{A}_{1/2}^k)$ and $(\mathcal{D}_\tau, \mathcal{A}_{1/2}^k)$, for an $x \in \mathcal{A}_{1/2}^k$ we get: $\tilde{\eta}(x) = \frac{1}{2}$. While for $(\mathcal{D}_{\alpha,\beta}, \mathcal{A}_{1/2}^k)$ we get: $\tilde{\eta}(x) = \frac{1 - \alpha + \beta}{2}$. Note that under $(\mathcal{D}_\tau, \mathcal{A}_{1/2}^k)$, we also have $[\tilde{\eta}(x)\{1 - \tilde{\eta}(x)\}]^2 = \frac{1}{16}$ similarly to $(\mathcal{D}, \mathcal{A}_{1/2}^k)$.

We now have all the necessary ingredients to define our two-sided hypothesis test.

3.1 A HYPOTHESIS TEST FOR CLASS-CONDITIONAL LABEL NOISE

We now define our null hypothesis (\mathcal{H}_0) and (implicit) alternative hypothesis (\mathcal{H}_1) as follows:

$$\begin{aligned} \mathcal{H}_0 &: \alpha = \beta \\ \mathcal{H}_1 &: \alpha \neq \beta \end{aligned} \quad (5)$$

Under the null and the alternative hypotheses, we have the following distributions for the estimated posterior of the anchor:

$$\mathcal{H}_0 : \hat{\eta}(x) \sim \mathcal{N}\left(\frac{1}{2}, \frac{1}{16} \cdot x^\top \hat{H}_n x\right) = \mathcal{N}\left(\frac{1}{2}, v(x)\right) \quad (6)$$

$$\begin{aligned} \mathcal{H}_1 : \hat{\eta}(x) &\sim \mathcal{N}\left(\frac{1 - \alpha + \beta}{2}, \frac{[(1 - \alpha + \beta)(\beta - \alpha)]^2}{16} \cdot x^\top \hat{H}_n x\right) \\ &= \mathcal{N}\left(\frac{1 + \alpha - \beta}{2}, \tilde{v}(x)\right) \end{aligned} \quad (7)$$

Level of Significance and Power of the test The *level of significance* (also known as Type I Error) is defined as follows:

$$a = \mathbb{P}(\text{reject } \mathcal{H}_0 \mid \mathcal{H}_0 \text{ is True}) \quad (8)$$

Rearranging Eq.6 we get the following:

$$\frac{\hat{\eta}(x) - 1/2}{\sqrt{v(x)}} \sim \mathcal{N}(0, 1)$$

Which for a chosen level of *significance* (a) allows us to define a region of retaining the null \mathcal{H}_0 . We let $z_{a/2}$ and $z_{1-a/2}$ denote the lower and upper critical values for retaining the null at a level of significance of a .

Retain \mathcal{H}_0 if:

$$\begin{aligned} z_{a/2} &\leq \frac{\hat{\eta}(x) - 1/2}{\sqrt{v(x)}} \leq z_{1-a/2} \\ z_{a/2} \cdot \sqrt{v(x)} + 1/2 &\leq \hat{\eta}(x) \leq z_{1-a/2} \cdot \sqrt{v(x)} + 1/2 \end{aligned} \quad (9)$$

Using the region of retaining the null hypothesis, we can now derive the *power* of the test.

Proposition 3.1. Power of the test

Under the distributions for the estimated posterior under the null and alternative hypotheses in Eqs.6&7, based on the definition of the hypotheses in Eq.5, the test has power:

$$1 - b_1 = \mathbb{P}(\text{reject } \mathcal{H}_0 \mid \mathcal{H}_0 \text{ is False}) \quad (10)$$

$$b_1 = \Phi\left(\frac{z \cdot \sqrt{v(x)} + \frac{\beta - \alpha}{2}}{\sqrt{\tilde{v}(x)}}\right) - \Phi\left(\frac{-z \cdot \sqrt{v(x)} + \frac{\beta - \alpha}{2}}{\sqrt{\tilde{v}(x)}}\right) \quad (11)$$

See Appendix 7.3 for the full derivation.

3.2 MULTIPLE ANCHOR POINTS

In this section we discuss how the properties of the test change in the setting where multiple anchors points are provided. The following expectation (\mathbb{E}), and variance (\mathbb{V}) operators, unless specified otherwise, are with respect to randomness in $\hat{\theta}_{mle}$ (see Eq.1).

Let $\hat{\eta}_i$ correspond to the i th instance in $\mathcal{A}_{1/2}^k$. Then for $\bar{\eta} = \frac{1}{k} \sum_{i=1}^k \hat{\eta}_i$ we have:

$$\bar{\eta} \sim \mathcal{N}\left(\frac{1}{2}, \frac{1}{16} \cdot \bar{x}^\top H \bar{x}\right)$$

where $\bar{x} = \frac{1}{k} \sum_{i=1}^k x_i$ with $x_i \in \mathcal{A}_{1/2}^k \forall i$.

For the full derivation see Appendix 7.4.

Anchors chosen at random We have that $x \in \mathcal{A}_{1/2}^k \rightarrow x^\top \beta_0 = 0$, so that for an orthonormal basis U , $x = Ur$. Without loss of generalisation we let $U_{:,0} = \frac{\beta_0}{\|\beta_0\|_2}$, and therefore $\eta(x) = 1/2 \rightarrow r_0 = 0$. In words: $\forall x \in \mathcal{A}_{1/2}^k$ we have that

x 's component in the direction of β_0 is 0. Pictorially, this is shown in Fig.2 (left). Anchor points are picked from the dotted line, which is described by the relationship $\eta(x) = 1/2$.

Now we make the assumption that x 's are random with $r_j \sim \mathbb{U}([-c, c])$. Therefore, $\mathbb{E}r_j = 0$, and $\mathbb{V}r_j = \frac{c^2}{3}$. In the following we use the subscript S in the operator \mathbb{E}_S to denote the randomness in choosing the set \mathcal{A} . In words: we assume that the set $\mathcal{A}_{1/2}^k$ is chosen uniformly at random from the set of all anchor points.

Combining these we get:

$$\begin{aligned} \mathbb{E}_S v(x) &= \mathbb{E}_S x^\top H x = \mathbb{E}_S r^\top U H U^\top r \\ &= \frac{dc^2}{3} \cdot \text{tr}(U H U^\top) = \frac{dc^2 q}{3} \end{aligned}$$

where $q = \text{tr}(H)$.

While for k anchor points chosen independently at random, we get:

$$\begin{aligned} \mathbb{E}_S v(\bar{x}) &= \mathbb{E}_S \left[\frac{1}{k^2} \sum_{i,j} x_i^\top H x_j \right] = \mathbb{E}_S \left[\frac{1}{k^2} \sum_{i,j} r_i^\top U H U^\top r_j \right] \\ &= \frac{dc^2}{3k} \cdot \text{tr}(U H U^\top) = \frac{dc^2 q}{3k} \end{aligned}$$

Following the same derivation as above we get:

$$b_k = \Phi\left(\frac{z \cdot \sqrt{v(\bar{x})} + \frac{\beta - \alpha}{2}}{\sqrt{\tilde{v}(\bar{x})}}\right) - \Phi\left(\frac{-z \cdot \sqrt{v(\bar{x})} + \frac{\beta - \alpha}{2}}{\sqrt{\tilde{v}(\bar{x})}}\right)$$

If we let $v = \mathbb{E}_S v(x)$ (similarly $\tilde{v} = \mathbb{E}_S \tilde{v}(x)$), then we have seen that $\mathbb{E}_S v(\bar{x}) = \frac{v}{k}$ (Reminder: expectations are with respect to the randomness in picking the anchor points). Then we have:

$$\frac{b_k}{b_1} = \frac{\Phi\left(\frac{z\sqrt{v} + h\sqrt{k}}{\sqrt{\tilde{v}}}\right) - \Phi\left(\frac{-z\sqrt{v} + h\sqrt{k}}{\sqrt{\tilde{v}}}\right)}{\Phi\left(\frac{z\sqrt{v} + h}{\sqrt{\tilde{v}}}\right) - \Phi\left(\frac{-z\sqrt{v} + h}{\sqrt{\tilde{v}}}\right)} \leq 1 \quad (12)$$

with $h = \frac{\beta - \alpha}{2}$.

In Fig.1 we compare the power ($1 - b$) of the test, as a function of the difference between the noise rates ($\beta - \alpha$), and number of anchor points used (k). We observe that a larger number of anchor points leads to a higher value for power.

3.3 MULTIPLE RELAXED ANCHORS-POINTS

In this section we see how the properties of the test change in the setting where the anchors do not have a perfect $\eta(x) = 1/2$. We now consider the case of $\mathcal{A}_{1/2,\delta}^k$.




Figure 1: Power of the hypothesis test as a function of $\beta - \alpha$, for a range of different k s. We set $v = (1/16) \cdot (X^\top DX)^{-1} = 0.1$.

Let x be such that $\eta(x) = \frac{1}{2} + \varepsilon$, where $\varepsilon \sim \mathcal{U}[-\delta, \delta]$ with $0 < \delta \ll 1$. (Note: by definition $\delta \leq 1/2$.)

For one instance we have the following:

$$\mathbb{E}_{\hat{\theta}} \hat{\eta} = 1/2 + \varepsilon, \quad \text{and} \quad \mathbb{E}_S \mathbb{E}_{\hat{\theta}} \hat{\eta} = 1/2$$

For the variance component we have:

$$[\hat{\eta}(1 - \hat{\eta})]^2 = \left[\left(\frac{1}{2} + \varepsilon \right) \left(\frac{1}{2} - \varepsilon \right) \right]^2 \approx \frac{1}{16} - \frac{\varepsilon^2}{2}$$

ignoring terms of order higher than ε^2 , under the assumption that $\delta \ll 1$.

When combined, for $\eta(x)$ for an $x \in \mathcal{A}_{1/2, \delta}^k$, we obtain:

$$\eta \sim \mathcal{N} \left(\frac{1}{2} + \varepsilon, \left(\frac{1}{16} - \frac{\varepsilon^2}{2} \right) \cdot x^\top H x \right)$$

which under expectation with respect to the randomness in picking the anchor point we get:

$$\eta \sim \mathcal{N} \left(\frac{1}{2}, \left(\frac{1}{16} - \frac{\delta^2}{6} \right) \cdot x^\top H x \right) \quad (13)$$

Moving on to consider multiple relaxed anchor points now, with $\bar{\eta} = \frac{1}{k} \sum_{i=1}^k \hat{\eta}_i$, for $x_i \in \mathcal{A}_{1/2, \delta}^k$ we have:

$$\mathbb{V}_S \mathbb{E}_{\hat{\theta}} \bar{\eta} = \mathbb{V}_S \left[\frac{1}{2} + \frac{1}{k} \sum_{i=1}^k \varepsilon_i \right] = \frac{\delta^2}{3k} \quad (14)$$

which shows that as we increase the number of anchor points (k) the average of the posteriors concentrates around $1/2$.

For the variance we have:

$$\mathbb{V}_{\hat{\theta}} \left[\frac{1}{k} \sum_{i=1}^k \hat{\eta}_i \right] = \left(\frac{1}{16} - \frac{\delta^2}{6} \right) \cdot \bar{x}^\top H \bar{x}$$




Figure 2: On the left we have the line from which the (strict) anchors are sampled from uniformly at random. On the right we show the rectangle from which the relaxed anchors are sampled from.

For the full derivation see Appendix 7.6.

Finally, bringing everything together we get:

$$\begin{aligned} \bar{\eta} &\sim \mathcal{N} \left(\frac{1}{2}, \left(\frac{1}{16} - \frac{\delta^2}{6} \right) \cdot \bar{x}^\top H \bar{x} \right) \\ &\approx \mathcal{N} \left(\frac{1}{2}, \frac{1}{16} \cdot \bar{x}^\top H \bar{x} \right) \end{aligned}$$

3.4 WHAT IF WE HAVE NO ANCHOR POINTS?

In the previous section we have shown that we can relax the hard constraint on the anchor points to be exactly $\eta = 1/2$, to $\eta \approx 1/2$. It might be natural then to ask if we need anchor points at all. If instead we were to sample points at random, then we would have the following:

$$\mathbb{E}_{p(x)} \eta(X) = \pi \quad (15)$$

The importance of needing for set of anchor points, either $\mathcal{A}_{1/2}^k$ or $\mathcal{A}_{1/2, \delta}^k$, is that, the anchor points would be centered around a known value $1/2$, as opposed to having no anchor points and sampling at random, where the anchor points would end up being centered around π .

Knowledge of the class priors could allow for a different type of hypothesis tests to asses the presence of label noise. We do not continue this discussion in the main document as it relies on very different type of information, but provide pointers in the Appendix 7.8.

3.5 PRACTICAL CONSIDERATIONS

Beyond Logistic Regression Our approach relies on the asymptotic properties of MLE estimators, and specifically of Logistic Regression. More complex models can be constructed in a similar fashion through polynomial feature expansion. However the extension of these tests to richer model-classes, such as Gaussian Processes, remains open.

Multi-class classification Multi-class classification setting can be reduced to *one-vs-all*, *all-vs-all*, or more general error-correcting output codes setups as described in Dietterich and Bakiri [1995], which rely on multiple runs of binary classification. In these settings then we could apply the proposed framework. The challenge would then be how to interpret $\eta = 1/2$.

Finding anchor points While it might not be straightforward for the user to provide us with instances whose true posterior is $\eta(x) = 1/2$, we do show how this can be relaxed, by allowing $\eta(x) \approx 1/2$. We then show how multiple anchor points can be stacked, improving the properties of the test.

4 RELATED WORK

There exist multiple works in the field of weak supervision where instead of being provided with the ground truth labels, the dataset is annotated with a weak version of them, usually derived from the true label and potentially influenced by exogenous variables. It goes beyond the scope of this work to discuss this field but these works and the references therein offer an overview of the field: Patrini [2016], Menon et al. [2015], Cid-Sueiro et al. [2014], Perello-Nieto et al. [2017], Frénay and Verleysen [2013].

As an instance of weak supervision, we briefly discuss approaches in the literature that relate to tackling the problem of learning with the presence of (flipping) noise in the labels. As already discussed in Theorem 2.1, in the case of uniform noise, under mild assumptions, we have robust risk minimisation. However, in the case of class-conditional noise, we do not have similar guarantees.

One common approach is to proceed by correcting the loss to be minimised. To better illustrate this, we consider the notion of *mixing matrix* $M \in [0, 1]^{c \times c}$, where $M_{i,j} = \mathbb{P}(\bar{y} = e^j \mid y = e^i)$ and follow the exposition from Patrini [2016]. Introducing the notation $e^i \in \{0, 1\}^c$ to denote the i th standard canonical vector with only the i th position being non-zero, and the multi-class class posterior $p(y \mid x) \in [0, 1]^c$, let us define the loss $l : \Delta^{c-1} \rightarrow \mathbb{R}^c$ in vector form, where $\Delta^{c-1} \subset [0, 1]^c$ is the c -dimensional simplex, as follows:

$$\ell(\hat{p}(y \mid x)) = (\ell(e^1, \hat{p}(y \mid x)), \dots, \ell(e^c, \hat{p}(y \mid x)))^\top \in \mathbb{R}^c$$

where in the case of the *cross-entropy* loss we have,

$$\ell(e^i, \hat{p}(y \mid x)) = -(e^i)^\top \log \hat{p}(y \mid x) = -\log \hat{p}(y = e^i \mid x)$$

Theorem 1 from Patrini [2016] then states that,

$$\operatorname{argmin}_{\hat{p}(y|x)} \mathbb{E}_{x, \bar{y}} M^{-1} \ell(\hat{p}(y \mid x)) = \operatorname{argmin}_{\hat{p}(y|x)} \mathbb{E}_{x, y} \ell(y, \hat{p}(y \mid x))$$

Anchor points and perfect samples Using these formulations, we are in a position where, if we have access to M , we can correct the training procedure to obtain an unbiased estimator. However, M is rarely known and difficult to estimate. Works on estimating M rely on having access to ‘perfect samples’ and can be traced back to Scott et al. [2013], and the idea was later adapted and generalised in Patrini [2016], Menon et al. [2015], Liu and Tao [2015], Perello-Nieto et al. [2020] to the multi-class setting. Interestingly, in Patrini et al. [2017] authors do not explicitly define these perfect samples, but rather assume they do exist in a large enough (validation) dataset X' - obtaining good experimental results. Similarly, Xia et al. [2019] also work by not explicitly requiring anchor points, but rather assuming their existence.




Figure 3: An illustration of how CCN alters the class distributions of two Gaussian class-conditional distributions.

Noisy examples An alternative line is followed by Northcutt et al. [2019, 2017], where the aim is to identify the *specific* examples that have been inflicted with noise. This is a non-trivial task unless certain assumptions can be made about the per-class distributions, and their shape. For example, if we can assume that the supports of the two classes do not overlap (i.e. $\eta(x)(1 - \eta(x)) \in \{0, 1\} \forall x$), then we can identify mislabelled instances using per-class densities. If this is not the case, then it would be difficult to differentiate between a mislabelled instance and an instance for which $\eta(x)(1 - \eta(x)) \in (0, 1)$. A different assumption could be uni-modality, which would again provide a prescription for identifying mislabelled instances through density estimation tools.

Distilled examples The authors in Cheng et al. [2020] go the opposite direction by trying to identify instances that *have not been corrupted* \rightarrow the *distilled examples*. At a first step the authors assume knowledge of an upper-bound² (Theorem 2 of Cheng et al. [2020]) which allows them to define sufficient conditions for identifying whether an instance is *clean*. At a second step they aim at estimating the (local) noise rate based on the neighbourhood of an instance (see Theorem 3 of Cheng et al. [2020]).

Informative priors The field of Bayesian Statistics is often concerned with constructing informative prior distributions that reflect the knowledge of experts. While it might

²The paper aims at tackling instance-dependent noise.

be challenging eliciting information from experts and modeling it quantitatively; it is often a necessary, and useful, step in low-data settings. Methods that rely on experts quantifying their beliefs include the *conditional means* approach of Bedrick et al. [1996] where a prior distribution is derived from the potential outcomes of given input, and Greenland [2001] where priors beliefs are first expressed in weakly informative ranges (see Gelman et al. [2008] for a discussion.) Other works rely on constructing priors based on the specification of quantiles (Garthwaite et al. [2005]).

Similarly to the first set of works we introduce and exploit *anchor points*, but not for directly estimating the mixing matrix, but rather to devise hypothesis tests to assess the likelihood of the *whole dataset* having been inflicted with noise.

5 EXPERIMENTS

In order to illustrate the properties of the tests, for the experiments we consider a synthetic dataset where the per-class distributions are Gaussians, with means $[1, 1]^\top$ and $[-1, -1]^\top$, with identity as scale. For this setup we know that anchor points should lie on the line $y = -x$, and draw them uniformly at random $x \in [-4, 4]$.

We analyse the following parameters of interest:

1. $N \in [500, 1000, 2000, 5000]$: the training sample size.
2. $(\alpha - \beta) \in [-0.05, 0.10, 0.20]$: the difference between the per-class noise rates.
3. $k \in [1, 2, 4, 8, 16, 32]$: the number of anchor points.
4. $\delta \in [0, 0.05, 0.10]$: how relaxed the anchor points are: $\eta(x) \in [0.50 - \delta, 0.50 + \delta]$.

For all combinations of N and $(\alpha - \beta)$ we perform 500 runs. In each run, we generate a clean version of the data \mathcal{D} , and then proceed by corrupting it to obtain a separate version: $\mathcal{D}_{\alpha,\beta}$. For both datasets, we fit a Logistic Regression model. We sample both the anchor points and relaxed anchor points. Finally, we then compute the z-scores, and subsequently the corresponding p-values.

The box-plots should be read as follows: $Q1$, $Q2$ & $Q3$ separate the data into 4 equal parts. The inner box starts (at the bottom) at $Q1$ and ends (at the top) at $Q3$, with the horizontal line inside denoting the median ($Q2$). The whiskers extend to show $Q1 - 1.5 \cdot IQR$, and $Q3 + 1.5 \cdot IQR$. IQR denotes the *Interquartile Range* and $IQR = Q3 - Q1$.

In Figures 4, 5 and 6 we have the following: moving to the right we increase the relaxation of anchor points, and moving downwards we increase the training sample-size. On the subplot level, on the x-axis we vary the number of anchor points, and on the y-axis we have the p-values. In all subplots we indicate with a red dashed line the mark of

0.10, and with a blue one the mark of 0.05, which would serve as rejection thresholds for the null hypothesis.

The experiments are illustrative of the claims made earlier in the paper. Below we discuss the findings in the experiments and what they mean with regards to Type I and Type II errors. We discuss these points in two parts; we first discuss the effect on sample size (N), difference in noise rates ($|\alpha - \beta|$) and number of anchor points (k).

Size of training set (N) As the size of training set (N) increases, the power increases. This can be seen Figures 4, 5 & 6. By moving down the first column, and fixing a value for k , where N increases, we see the range of the purple box-plots decreasing, and essentially a larger volume of tests falling under the cut-off levels of significance (red and blue dashed lines). This is expected given that the variance of the MLE $\hat{\theta}_{MLE}$ vanishes as N increases, as is seen in Eq.1 and discussion underneath it. Subsequently, variance in Eq.3 also vanishes, which increases power (Eq.10).

Difference in noise rates ($|\alpha - \beta|$) As $|\alpha - \beta|$ increases, the power increases. This can be seen in Figures 4, 5 & 6, by fixing a particular subplot in the first column (for example, top-left one), and a value for k , we see again that the volume moves down. As presented in Eq.11, as $\beta - \alpha$ increases, the power also increases.

Number of anchor points (k) The same applies to the number of anchor points – as the number of anchor points (k) increases, the power of the test increases. This can be seen in all three figures by focusing in any subplot in the first column, and considering the purple box-plots moving to the right. In Eq.12 we see effect of k on the power.

In all three discussions above we focused on the first column of each of the figures – which shows results from experiments on strict anchors. What we also observe in this case (the first column of all figures) is that the p-values follow the uniform distribution under the null (as expected, given the null hypothesis is true) – shown by the green box-plots. Therefore the portion of Type I Errors = α (the level of significance Eq.8). When we relax the requirements for strict anchors to allow for values close to $1/2$, we introduce a bias in the lower and upper bounds in Eq.9 of $+\epsilon$. While $\mathbb{E}\epsilon = 0$ this shift on the boundaries of the retention region will increase Type I Error. On the other hand, in Eq.14 we see how this bias decreases as you increase the number of anchor points. Both of these phenomena are also shown experimentally by looking at the latter two columns of the figures.

Anchor point relaxation (δ) Lastly, we examine the effect of relaxing the strictness of the anchors (δ), $\eta(x) \in [0.50 - \delta, 0.50 + \delta]$ on the properties of the test. As just discussed we see that as we increase the number of anchor



Figure 5: Box-plots with fixed $|\beta - \alpha| = 0.10$. Red dotted line indicates the mark of 0.10, and the blue one 0.05.

points Type I Error decreases (volume of green box-plots under each of the cut-off points). We also observe that, as compared to only allowing strict anchors, the power is not affected significantly – with the effect decreasing as the number of anchor points increases. Furthermore, in the latter two columns we also observe the phenomena mentioned in the discussion concerning the first column only.

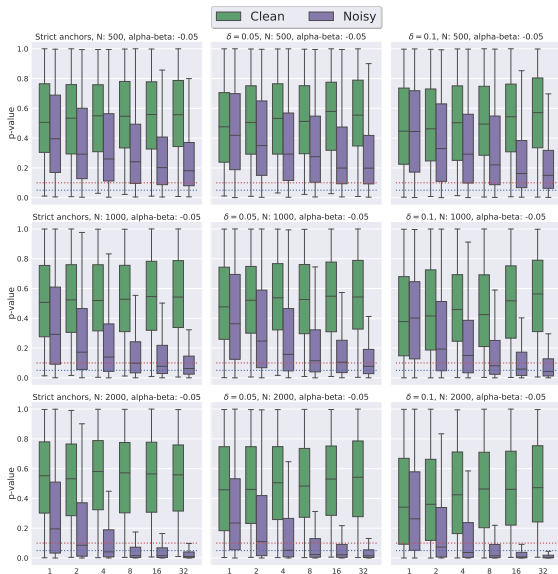


Figure 4: Box-plots with fixed $|\beta - \alpha| = 0.05$. Red dotted line indicates the mark of 0.10, and the blue one 0.05.




Figure 6: Box-plots with fixed $|\beta - \alpha| = 0.20$. Red dotted line indicates the mark of 0.10, and the blue one 0.05.

6 CONCLUSION & FUTURE WORK

In this work we introduce the first statistical hypothesis test for class-conditional label noise. Our approach requires the specification of anchor points, i.e. instances whose labels are highly uncertain under the true posterior probability distribution, and we show that the test’s significance and power is preserved over several relaxations on the requirements for these anchor points.

Our experimental analysis, which confirms the soundness of our test, explores many configurations of practical interest for practitioners using this test. Of particular importance for practitioners, since anchor specification is under their control, is the high correspondence shown theoretically and experimentally between the number of anchors and test significance.

Future work will cover both theoretical and experimental components. On the theoretical front, we are interested in extending this work to the multi-class setting (and understanding the implications of our requirements in this case), understanding the test’s value under a richer set of classification models, and further relaxing requirements on true posterior uncertainty for anchor points. Experimentally, we are particularly interested in applying the tests to diagnostically challenging healthcare problems and utilising clinical experts for anchor specification.

References

- Edward J Bedrick, Ronald Christensen, and Wesley Johnson. A new perspective on priors for generalized linear models. *Journal of the American Statistical Association*, 91(436): 1450–1460, 1996.
- George Casella and Roger Berger. *Statistical Inference*. Duxbury Resource Center, June 2001. ISBN 0534243126.
- Jiacheng Cheng, Tongliang Liu, Kotagiri Ramamohanarao, and Dacheng Tao. Learning with bounded instance and label-dependent label noise. In *International Conference on Machine Learning*, pages 1789–1799. PMLR, 2020.
- Jesus Cid-Sueiro, Dario Garcia-Garcia, and Raul Santos-Rodriguez. Consistency of losses for learning from weak labels. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, pages 197–210, 2014.
- Thomas G. Dietterich and Ghulum Bakiri. Solving multi-class learning problems via error-correcting output codes. *J. Artif. Int. Res.*, 2(1):263–286, January 1995.
- R. Fergus, L. Fei-Fei, P. Perona, and A. Zisserman. Learning object categories from google’s image search. In *IEEE International Conference on Computer Vision*, volume 2, pages 1816–1823 Vol. 2, 2005.
- Benoît Fréney and Michel Verleysen. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5):845–869, 2013.
- Paul H Garthwaite, Joseph B Kadane, and Anthony O’Hagan. Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association*, 100(470):680–701, 2005.
- Andrew Gelman, Aleks Jakulin, Maria Grazia Pittau, Yu-Sung Su, et al. A weakly informative default prior distribution for logistic and other regression models. *Annals of applied Statistics*, 2(4):1360–1383, 2008.
- Aritra Ghosh, Naresh Manwani, and PS Sastry. Making risk minimization tolerant to label noise. *Neurocomputing*, 160:93–107, 2015.
- Sander Greenland. Putting background information about relative risks into conjugate prior distributions. *Biometrics*, 57(3):663–670, 2001.
- Ishan Jindal, Matthew Nokleby, and Xuewen Chen. Learning deep networks from noisy labels with dropout regularization. In *IEEE International Conference on Data Mining*, 2016.
- Neil D Lawrence. Data readiness levels. *arXiv preprint arXiv:1705.02245*, 2017.
- Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38(3):447–461, 2015.
- 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*, pages 125–134. PMLR, 2015.
- Curtis G. Northcutt, Tailin Wu, and Isaac L. Chuang. Learning with confident examples: Rank pruning for robust classification with noisy labels. In *Uncertainty in Artificial Intelligence*, 2017.
- Curtis G Northcutt, Lu Jiang, and Isaac L Chuang. Confident learning: Estimating uncertainty in dataset labels. *arXiv preprint arXiv:1911.00068*, 2019.
- Giorgio Patrini. *Weakly supervised learning via statistical sufficiency*. PhD thesis, ANU College of Engineering & Computer Science, The Australian National University, 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*, pages 1944–1952, 2017.
- Miquel Perello-Nieto, Raul Santos-Rodriguez, and Jesus Cid-Sueiro. Adapting supervised classification algorithms to arbitrary weak label scenarios. In *Advances in Intelligent Data Analysis*, pages 247–259, 2017.
- Miquel Perello-Nieto, Raul Santos-Rodriguez, Dario Garcia-Garcia, and Jesus Cid-Sueiro. Recycling weak labels for multiclass classification. *Neurocomputing*, 400:206–215, 2020.
- Florian Schroff, Antonio Criminisi, and Andrew Zisserman. Harvesting image databases from the web. *IEEE transactions on pattern analysis and machine intelligence*, 33: 754–66, 04 2011.
- Clayton Scott, Gilles Blanchard, and Gregory Handy. Classification with asymmetric label noise: Consistency and maximal denoising. In *Conference on learning theory*, pages 489–511. PMLR, 2013.
- Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? *arXiv preprint arXiv:1906.00189*, 2019.

7 APPENDIX

7.1 DERIVATION OF THE NOISY POSTERIOR

We can compute the noisy posterior: $\tilde{\eta}(x) = \mathbb{P}(\tilde{y} = 1 \mid X = x)$ as follows:

$$\begin{aligned} \tilde{\eta}(x) &= \mathbb{P}(\tilde{y} = 1, y = 1 \mid X = x) \\ &\quad + \mathbb{P}(\tilde{y} = 1, y = 0 \mid X = x) \\ &= \mathbb{P}(\tilde{y} = 1 \mid y = 1)\mathbb{P}(y = 1 \mid X = x) \\ &\quad + \mathbb{P}(\tilde{y} = 1 \mid y = 0)\mathbb{P}(y = 0 \mid X = x) \\ &= (1 - \alpha) \cdot \eta(x) + \beta \cdot (1 - \eta(x)) \\ &= \begin{cases} (1 - \alpha - \beta) \cdot \eta(x) + \beta & \text{if (CCN)} \\ (1 - 2\tau) \cdot \eta(x) + \tau & \text{if (UN)} \end{cases} \end{aligned}$$

7.2 SUFFICIENT CONDITIONS FOR ROBUSTNESS TO UNIFORM NOISE

In the case of uniform-noise we have that:

$$\begin{aligned} \text{sign}(\tilde{\eta}(x) - 0.50) &= \text{sign}((1 - 2\tau) \cdot (\eta(x) - 0.50)) \\ &= \text{sign}(\eta(x) - 0.50) \end{aligned}$$

where the last equality holds under the assumptions in Def. 2.1.

$$\begin{aligned} \mathbb{E}_{X,Y} l(Y, f(X)) &= \mathbb{E}_X \mathbb{E}_{Y|X} l(Y, f(X)) = \mathbb{E}_X \mathbb{E}_{Y|X} \psi(Yf(X)) \\ \text{but we have access to noisy versions of labels} \\ \mathbb{E}_X \mathbb{E}_{\tilde{Y}|X} \psi(\tilde{Y}f(X)) \\ &= \mathbb{E}_X \left[\tilde{\eta} \cdot \psi(f(X)) + (1 - \tilde{\eta}) \cdot \psi(-f(X)) \right] \\ &= \mathbb{E}_X \left[[(1 - \tau)\eta + \tau(1 - \eta)] \cdot \psi(f(X)) \right. \\ &\quad \left. + [(1 - \tau)(1 - \eta) + \tau\eta] \cdot \psi(-f(X)) \right] \\ &= \mathbb{E}_X \left[(1 - 2\tau) \cdot \mathbb{E}_{Y|X} \psi(Yf(X)) \right. \\ &\quad \left. + \tau \cdot [\psi(f(X)) + \psi(-f(X))] \right] \\ &= (1 - 2\tau) \cdot \mathbb{E}_{X,Y} l(Y, f(X)) + \tau \cdot \text{BIAS}_\tau(\psi) \end{aligned}$$

This implies that if we train under uniform-noise with rate: $\tau < 0.50$, with a loss with the property that $\psi(f(X)) + \psi(-f(X)) = K$ for a constant K then risk minimisation is tolerant to noise Ghosh et al. [2015].

7.3 DERIVATION FOR PROPOSITION 3.1

We let L and U denote the lower and upper bounds in Eq.9 respectively, and let $\varepsilon \sim \mathcal{N}(0, 1)$.

$$\begin{aligned} b_1 &= \mathbb{P}(\text{retain } \mathcal{H}_0 \mid \mathcal{H}_0 \text{ is False}) \\ &= \mathbb{P}\left(L \leq \hat{\eta}(x) \leq U \mid \hat{\eta}(x) \sim \mathcal{N}\left(\frac{1 + \alpha - \beta}{2}, \tilde{v}(x)\right)\right) \\ &= \mathbb{P}\left(L \leq \hat{\eta}(x) \leq U \mid \hat{\eta}(x) = \frac{1 + \alpha - \beta}{2} + \sqrt{\tilde{v}(x)} \cdot \varepsilon\right) \\ &= \mathbb{P}\left(\frac{L - \frac{1 + \alpha - \beta}{2}}{\sqrt{\tilde{v}(x)}} \leq \varepsilon \leq \frac{U - \frac{1 + \alpha - \beta}{2}}{\sqrt{\tilde{v}(x)}}\right) \\ &= \mathbb{P}\left(\frac{-z \cdot \sqrt{v(x)} + h}{\sqrt{\tilde{v}(x)}} \leq \varepsilon \leq \frac{z \cdot \sqrt{v(x)} + h}{\sqrt{\tilde{v}(x)}}\right) \\ &= \Phi\left(\frac{z \cdot \sqrt{v(x)} + h}{\sqrt{\tilde{v}(x)}}\right) - \Phi\left(\frac{-z \cdot \sqrt{v(x)} + h}{\sqrt{\tilde{v}(x)}}\right) \end{aligned}$$

where we have used: $h = \frac{\beta - \alpha}{2}$, for ease of notation.

7.4 MEAN & VARIANCE FOR MULTIPLE ANCHORS-POINTS

For the expectation we have:

$$\mathbb{E}\left[\frac{1}{k} \sum_{i=1}^k \hat{\eta}_i\right] = \frac{1}{k} \sum_{i=1}^k \mathbb{E} \hat{\eta}_i = \frac{1}{2}$$

And, for the variance we have:

$$\mathbb{V}\left[\frac{1}{k} \sum_{i=1}^k \hat{\eta}_i\right] = \frac{1}{k^2} \left[\sum_{i=1}^k \mathbb{V} \hat{\eta}_i + 2 \cdot \sum_{i,j>i} \text{Cov}(\hat{\eta}_i, \hat{\eta}_j) \right]$$

with

$$\begin{aligned} \mathbb{V} \hat{\eta}_i &= \frac{1}{16} \cdot x_i^\top H x_i \\ \text{and } \text{Cov}(\hat{\eta}_i, \hat{\eta}_j) &= \frac{1}{16} \cdot x_i^\top H x_j \end{aligned}$$

For the derivation of $\text{Cov}(\hat{\eta}_i, \hat{\eta}_j)$ see Appendix 7.5.

7.5 COVARIANCE OF ESTIMATED POSTERiors FOR THE CASE OF MULTIPLE ANCHOR-POINTS

In this section we estimate $\text{Cov}(\hat{\eta}(x_i), \hat{\eta}(x_j))$ for the multiple anchors.

We will make use of the following: Let $\hat{\theta}$ be such that $\mathbb{E}\hat{\theta} = \theta_0$, then we have:

$$\begin{aligned} f(\hat{\theta}) &\approx f(\theta_0) + (\hat{\theta} - \theta_0)^T \nabla f|_{\theta_0} \\ &\quad + \frac{1}{2} (\hat{\theta} - \theta_0)^T \nabla^2 f|_{\theta_0} (\hat{\theta} - \theta_0) \\ &= f(\theta_0) + \frac{1}{2} (\hat{\theta} - \theta_0)^T \nabla^2 f|_{\theta_0} (\hat{\theta} - \theta_0) \end{aligned}$$

Let $\hat{\eta}(x_i) = \hat{\eta}_i$, then we have:

$$\begin{aligned} \text{Cov}(\hat{\eta}_i, \hat{\eta}_j) &= \mathbb{E}[(\hat{\eta}_i - \mathbb{E}\hat{\eta}_i)(\hat{\eta}_j - \mathbb{E}\hat{\eta}_j)] \\ &= \mathbb{E}[\hat{\eta}_i \hat{\eta}_j] - \eta_i \eta_j \end{aligned}$$

A few useful derivations (the hat ($\hat{\eta}(x)$) is implied), we also let $\nabla = \nabla_{\theta}$:

1. $\nabla \eta_i = x_i \eta_i (1 - \eta_i) = x_i \gamma_i$
2. $\nabla \gamma_i = \nabla[\eta_i - \eta_i^2] = (1 - 2\eta_i) x_i \gamma_i$
3. $\nabla[\eta_i \eta_j] = \eta_j \cdot x_i \gamma_i + \eta_i \cdot x_j \gamma_j$
4. $\nabla[\eta_j \cdot x_i \gamma_i] = x_i [\eta_j \nabla \gamma_i + \gamma_i \nabla \eta_j]$
 $= x_i [\eta_j \cdot (1 - 2\eta_i) x_i \gamma_i + \gamma_i \cdot x_j \gamma_j]$
 $= x_i x_i^\top \cdot \eta_j (1 - 2\eta_i) \gamma_i + x_i x_j^\top \cdot \gamma_i \gamma_j$
5. $\nabla^2[\eta_i \eta_j] = x_i x_i^\top \cdot \eta_j (1 - 2\eta_i) \gamma_i$
 $+ 2 \cdot x_i x_j^\top \cdot \gamma_i \gamma_j$
 $+ x_j x_j^\top \cdot \eta_i (1 - 2\eta_j) \gamma_j$

by plugging $\eta_i = \eta_j = 1/2$ in item 5 above, combined with Eq.1, and approximation of the Fisher-Information matrix with the hessian, we get:

$$\text{Cov}(\hat{\eta}(x_i), \hat{\eta}(x_j)) = \frac{1}{16} x_i^\top \hat{H} x_j$$

7.6 VARIANCE FOR MULTIPLE RELAXED ANCHOR-POINTS

For the variance we have:

$$\mathbb{V}_{\hat{\theta}} \left[\frac{1}{k} \sum_{i=1}^k \hat{\eta}_i \right] = \frac{1}{k^2} \left[\sum_{i=1}^k \mathbb{V}_{\hat{\theta}} \hat{\eta}_i + 2 \cdot \sum_{i,j>i} \text{Cov}_{\hat{\theta}}(\hat{\eta}_i, \hat{\eta}_j) \right]$$

with

$$\mathbb{E}_S \mathbb{V}_{\hat{\theta}} \hat{\eta}_i = \left(\frac{1}{16} - \frac{\delta^2}{6} \right) \cdot x_i^\top H x_i$$

$$\text{and } \mathbb{E}_S \text{Cov}_{\hat{\theta}}(\hat{\eta}_i, \hat{\eta}_j) = \left(\frac{1}{16} - \frac{\delta^2}{6} \right) \cdot x_i^\top H x_j$$

For the derivation of $\text{Cov}(\hat{\eta}_i, \hat{\eta}_j)$ see Appendix 7.7.

7.7 COVARIANCE OF ESTIMATED POSTERiors FOR THE CASE OF MULTIPLE RELAXED ANCHOR-POINTS

In the section we estimate: $\text{Cov}(\hat{\eta}(x_i), \hat{\eta}(x_j))$ for the case of having multiple relaxed anchors.

We continue from Item 5 of Appendix 7.5:

$$\begin{aligned} \nabla^2[\eta_i \eta_j] &= x_i x_i^\top \cdot \eta_j (1 - 2\eta_i) \gamma_i \\ &\quad + 2 \cdot x_i x_j^\top \cdot \gamma_i \gamma_j \\ &\quad + x_j x_j^\top \cdot \eta_i (1 - 2\eta_j) \gamma_j \end{aligned}$$

We let $\eta_i = 1/2 + \varepsilon_i$ and $\eta_j = 1/2 + \varepsilon_j$, then we have: (for the first and third terms above)

$$\begin{aligned} \eta_j (1 - 2\eta_i) \gamma_i &= \left(\frac{1}{2} + \varepsilon_j \right) (-2\varepsilon_i) \left(\frac{1}{2} + \varepsilon_i \right) \left(\frac{1}{2} - \varepsilon_i \right) \\ \mathbb{E}_S[\eta_j (1 - 2\eta_i) \gamma_i] &= \mathbb{E}_S \left[\left(\frac{1}{2} + \varepsilon_j \right) (-2\varepsilon_i) \left(\frac{1}{2} + \varepsilon_i \right) \left(\frac{1}{2} - \varepsilon_i \right) \right] \\ &= \mathbb{E}_S \left[(-\varepsilon_i) \left(\frac{1}{2} + \varepsilon_i \right) \left(\frac{1}{2} - \varepsilon_i \right) \right] \\ &= \mathbb{E}_S \left[-\varepsilon_i \left(\frac{1}{4} - \varepsilon_i^2 \right) \right] = 0 \end{aligned}$$

For the second term we have:

$$\begin{aligned} \gamma_i &= \left(\frac{1}{2} - \varepsilon_i \right) \left(\frac{1}{2} + \varepsilon_i \right) = \frac{1}{4} - \varepsilon_i^2 \\ \mathbb{E}_S \gamma_i \gamma_j &= \mathbb{E}_S \left(\frac{1}{4} - \varepsilon_i^2 \right) \left(\frac{1}{4} - \varepsilon_j^2 \right) \\ &= \frac{1}{16} - \frac{\delta^2}{6} + \delta^4 \\ &\approx \frac{1}{16} - \frac{\delta^2}{6} \end{aligned}$$

7.8 TEST BASED ON PRIORS

Another important relationship is that between the clean and noisy class priors: $\tilde{\pi} = \mathbb{P}(\tilde{y} = 1)$:

$$\begin{aligned} \mathbb{P}(\tilde{y} = 1) &= \mathbb{P}(\tilde{y} = 1, y = 1) + \mathbb{P}(\tilde{y} = 1, y = 0) \\ &= \mathbb{P}(\tilde{y} = 1 | y = 1) \mathbb{P}(y = 1) \\ &\quad + \mathbb{P}(\tilde{y} = 1 | y = 0) \mathbb{P}(y = 0) \\ &= (1 - \alpha) \cdot \pi + \beta \cdot (1 - \pi) \end{aligned}$$

which under the two settings, *UN* and *CCN*, gives:

$$\tilde{\pi} = \begin{cases} (1 - \alpha - \beta) \cdot \pi + \beta & \text{if (CCN)} \\ (1 - 2\tau) \cdot \pi + \tau & \text{if (UN)} \end{cases} \quad (16)$$

The relationships in Eq.16, combined with the knowledge of the true class priors, would allow someone to carry out Binomial Hypothesis Tests for presence of label noise. These tests would not need to rely on MLE asymptotics.