Decontamination of Mutually Contaminated Models

Gilles Blanchard
Universität Potsdam

Clayton Scott
University of Michigan

Abstract

A variety of machine learning problems are characterized by data sets that are drawn from multiple different convex combinations of a fixed set of base distributions. We call this a *mutual contamination* model. In such problems, it is often of interest to recover these base distributions, or otherwise discern their properties. This work focuses on the problem of classification with multiclass label noise, in a general setting where the noise proportions are unknown and the true class distributions are nonseparable and potentially quite complex. We develop a procedure for decontamination of the contaminated models from data, which then facilitates the design of a consistent discrimination rule. Our approach relies on a novel method for estimating the error when projecting one distribution onto a convex combination of others, where the projection is with respect to a statistical distance known as the separation distance. Under sufficient conditions on the amount of noise and purity of the base distributions, this projection procedure successfully recovers the underlying class distributions. Connections to novelty detection, topic modeling, and other learning problems are also discussed.

1 Introduction

This paper considers a general framework for multiclass classification with label noise. As we develop later, this framework encompasses or relates to several other machine learning problems, including novelty detection, crowdsourcing, topic modeling, and learning from partial labels. Each of these applications involves *mutual contamination models*, meaning that observed data are drawn from mixtures of the underlying probability distributions of interest, and therefore may be viewed as being contaminated.

We begin by stating the problem of multiclass classification with label noise in a general setting. There are L classes, each governed by a class-conditional distribution P_i. The learner observes training random samples drawn from the contaminated distributions

$$\tilde{P}_i = \sum_{j=1}^{L} \pi_{ij} P_j, \tag{1}$$

$i = 1, \ldots, L$, where $\pi_{ij} \geq 0$ and $\sum_j \pi_{ij} = 1$ for each i. Here π_{ij} is the probability that an instance with observed label i is actually a realization of P_j. We consider a quite general setting where the π_{ij} are *unknown*, and the distributions P_i are not amenable to parametric modeling. We want to stress that the issue of contamination of training data comes *in addition to*, and is *different from*, the usual source of uncertainty about the labels in "noisy classification", which traditionally simply means that the supports of the true P_i can overlap (which we also allow here). In the contamination setting, training a conventional classifier directly on the contaminated training data will in general lead to a biased classifier at test time (Scott et al., 2013), and in particular be asymptotically inconsistent.

Our contribution is to establish general sufficient conditions on the true class-conditional distributions and label noise mixing proportions for recovery of the true distributions, which then enables the design of a consistent discrimination rule. The sufficient conditions essentially state that the examples drawn from $\tilde{P}_i$ come "mostly" from P_i, and that the distributions $\{P_i\}$ are "pure" with respect to each other. Both of these notions are made precise below.

At the heart of our approach is a novel technique for "decontamination" of the contaminated models. In particular, we show that under the sufficient conditions, if we project the contaminated distribution $\tilde{P}_i$

Appearing in Proceedings of the 17^{th} International Conference on Artificial Intelligence and Statistics (AISTATS) 2014, Reykjavik, Iceland. JMLR: W&CP volume 33.

onto the convex hull of the other contaminated distributions, the "residual" distribution is the true class-conditional distribution P_i. The projection is with respect to a statistical distance known as the separation distance, and a key part of our contribution is the development of a universally consistent estimator of the error of this projection. This estimator is valid regardless of whether the sufficient conditions for label noise recovery hold, and is therefore of independent interest.

1.1 Motivation

The label noise model (1) is relevant in a number of applications. For example, Scott et al. (2013) describe a problem in nuclear particle classification. Each class corresponds to a different type of nuclear particle, and label noise comes from the fact that it is impossible to eliminate different particle types from the background, so that every training sample for a given particle type is always contaminated with particles of other types.

Another potential application of the label noise model is crowdsourcing, where the labels of training data sets are provided by a mixture of expert and non-expert sources. One approach to crowdsourcing is to assume each annotator makes mistakes according to model (1), where the contamination probabilities are specific to each annotator. Then there is a label nois problem for each annotator. The "two-coin" model of Raykar et al. (2010) studies this setup in the two-class setting.

As a second formulation of crowdsourcing, consider unlabeled data representing some mixture of the true class-conditional distributions P_i, and suppose the data is labeled by crowdsourcing. One simple labelling model is as follows: with probability $1-\alpha$, an example is labeled by an expert, in which case the correct label is always assigned. With probability α, a random guesser assigns the label according to some predetermined distribution over class that is independent of the actual example. This fits the contamination model (1), and furthermore, as we argue in Sec. 5, satisfies our sufficient conditions on the noise.

A third relevant application is novelty detection. Suppose that for $i < L$ and $j \neq i$, we have $\pi_{ij} = 0$. In other words, the first $L-1$ data samples are uncontaminated. The last data sample, on the other hand, is drawn from a mixture of all classes. We may think of this as a semi-supervised learning problem, where the first $L-1$ random samples represent training data, while the last sample is an unlabeled testing sample. Then P_L corresponds to a novel class that is not represented among the $L-1$ training classes. We argue in Sec. 5 that this setting also satisfies our sufficient conditions on the noise, thus yielding a consistent discrimination rule for multiclass novelty detection.

Several previous works on classification with label noise also adopt (1) or an equivalent model, and we refer the reader to these works for additional applications of the label noise setting (Lawrence and Schölkopf, 2001; Bouveyron and Girard, 2009; Long and Servido, 2010; Manwani and Sastry, 2011; Stempfel and Ralaivola, 2009; Natarajan et al., 2013). Significantly, our works provides estimates of the proportions π_{ij}, which are assumed known in several of these earlier works.

Topic modeling is another problem that is closely related to multiclass label noise. In topic modeling, the base distributions P_i correspond to topics, and each $\tilde{P}_i$ corresponds to a document, viewed as a mixture of topics. An important difference between topic modeling and classification with label noise is that in topic modeling, typically far more documents are observed than topics. Nonetheless, as we discuss in Sec. 5, our sufficient conditions on the purity of the base distributions, when specialized to the discrete setting of topic modeling, coincide with the "separability" condition that has been widely adopted in that area.

Finally, we note that multiclass label noise has strong similarities to the problem of learning from partial labels (Cour et al., 2011). In this problem, every training data point is labeled by a subset $S \subset \{1, \ldots, L\}$ of possible labels, as opposed to a single label as in standard classification. The true label is one of the elements of the subset, but it is not known which one. Therefore, collecting all training examples sharing a common S gives a contamination model $\tilde{P}_S = \sum_{j \in S} \pi_{S,j} P_j$, and there is a contamination model (with associated data) for every observed S. We conjecture that this problem can be converted to a multiclass label noise problem, satisfying our sufficient conditions on the noise, through appropriate resampling of the data. Some insight is given in Sec. 5, although a full development is deferred to future work.

1.2 Related Work

Our work extends the recent work of Scott et al. (2013), reviewed in the next section, which studies label noise for binary classification ($L = 2$). We find there to be some significant differences between the multiclass and binary cases. Indeed the decontamination procedure is considerably more complex, as are the sufficient conditions for recovery. Multiclass label noise has received little attention in the literature. Existing theoretical work on label noise (of which we are aware) focuses on the binary case; we refer the reader to Scott et al. (2013) for a recent review.

The aforementioned projection in distribution space is accomplished by means of a problem we call multi-

sample mixture proportion estimation, which is apparently new. It generalizes a two-sample version developed by Blanchard et al. (2010). It is also similar to the problem of semi-supervised class proportion estimation; the latter problem essentially assumes that the projection error is zero, and is only concerned with estimating the mixing weights giving the projection (Hall, 1981; Titterington, 1983; Latinne et al., 2001; Du Plessis and Sugiyama, 2012). When the projection error is nonzero, as in our work, such methods are inconsistent.

1.3 Outline

In the next section we review the work of Scott et al. (2013) on label noise in the binary setting. Section 3 presents the aforementioned mixture proportion estimation problem and our universally consistent estimator. Section 4 introduces sufficient conditions under which this estimator successfully decontaminates the noisy distributions in the multiclass label noise setting. The final section connects our results more concretely to the other learning problems mentioned above. Proofs are contained in the supplemental file.

2 Label Noise in the Binary Case

Scott et al. (2013) study label noise in the binary case, $L = 2$. Their work hinges on the following result, which appears originally in Blanchard et al. (2010).

Proposition 1. *Given probability distributions F_0, F_1 on a measurable space $(\mathcal{X}, \mathcal{C})$, define*

$$\kappa^*(F_0|F_1) = \max\left\{\kappa \in [0,1] \middle| \exists \text{ a distribution } G \text{ s.t. } F_0 = (1-\kappa)G + \kappa F_1\right\}; \quad (2)$$

If $F_0 \neq F_1$, then $\kappa^*(F_0|F_1) < 1$ and the above supremum is attained for a unique distribution G (which we refer to as the residue *of F_0 w.r.t. F_1). Furthermore, the following equivalent characterization holds:*

$$\kappa^*(F_0|F_1) = \inf_{C \in \mathcal{C}, F_1(C) > 0} \frac{F_0(C)}{F_1(C)}. \quad (3)$$

The number $\kappa^*(F_0|F_1)$ can be understood at the maximum possible proportion of F_1 present in F_0. The result implies that the function $1 - \kappa^*(F_0|F_1)$ is a statistical distance, i.e., a functional that is nonnegative and equal to zero iff $F_0 = F_1$. This quantity has been called the *separation distance*, and has arisen previously in studies of Markov chain convergence (Aldous and Diaconis, 1987).

For the label noise/contamination model (1), the following two conditions are shown by Scott et al. (2013) to be sufficient for decontamination:

- $\pi_{12} + \pi_{21} < 1$,
- $\kappa^*(P_1|P_2) = \kappa^*(P_2|P_1) = 0$.

The former condition bounds the total amount of label noise, while the latter condition is referred to as *mutual irreducibility*, and says that it is not possible to write P_1 as a nontrivial mixture of P_2 and some other distribution, and *vice versa*.

In particular, by a simple algebraic manipulation of (1), it is shown that under the first condition

$$\tilde{P}_1 = (1-\kappa_1)P_1 + \kappa_1 \tilde{P}_2 \quad (4)$$

$$\tilde{P}_2 = (1-\kappa_2)P_2 + \kappa_2 \tilde{P}_1, \quad (5)$$

for unique $\kappa_1, \kappa_2 < 1$. It is then shown that if P_1 and P_2 are mutually irreducible, then $\kappa_1 = \kappa^*(\tilde{P}_1|\tilde{P}_2)$ and $\kappa_2 = \kappa^*(\tilde{P}_2|\tilde{P}_1)$ and therefore P_1 and P_2 are the respective residues of $\tilde{P}_1$ w.r.t. $\tilde{P}_2$ and $\tilde{P}_2$ w.r.t. $\tilde{P}_1$. This establishes decontamination in the population case.

In the sample case, the problem of estimating κ^* is referred to as *mixture proportion estimation*. The universally consistent estimator $\widehat{\kappa}$ of Blanchard et al. (2010) for κ^* is applied to estimate κ_1 and κ_2. Now all terms in (4) and (5) can be estimated except for P_1 in (4) and P_2 in (5). Therefore (4) and (5) can be solved for (estimates of) P_1 and P_2, which are then available for the design of a consistent discrimination rule.

Multiclass label noise can be treated analogously, although the generalization is far from trivial.

3 Multi-Sample Mixture Proportion Estimation

Let $(\mathcal{X}, \mathcal{C})$ be a measurable space, and let F_i be probability distributions on this space, $i = 0, 1, \ldots, M$. Let S_M denote the $(M-1)$-dimensional simplex $\left\{\boldsymbol{\mu} = (\mu_1, \ldots, \mu_M) \in \mathbb{R}^M \middle| \forall i\ \mu_i \geq 0 \text{ and } \sum_i \mu_i = 1\right\}$. For $\boldsymbol{\mu} \in S_M$, denote the probability distribution $F_{\boldsymbol{\mu}} := \sum_{i=1}^M \mu_i F_i$. We first define the maximum collective contaminating proportion of distributions $(F_i)_{1 \leq i \leq M}$ in F_0 by the following generalization of the binary case:

Definition 1. *Given probability distributions $F_i, i = 0, \ldots, M$, define*

$$\kappa^*(F_0|F_1, \ldots, F_M) = \max_{\boldsymbol{\mu} \in S_M} \kappa^*(F_0|F_{\boldsymbol{\mu}}). \quad (6)$$

If there is a unique $F_{\boldsymbol{\mu}}$ achieving the maximum in this definition, it may be thought of as the projection of F_0 onto the convex hull of $F_1, \ldots, F_M$ with respect to the separation distance. However, there exist simple examples for which $F_{\boldsymbol{\mu}}$ attaining κ^* is not unique; for example, suppose $M = 2$. Let F_0 be uniform on $\{0, 1, 2\}$,

F_1 uniform on $\{0,1\}$, and F_2 uniform on $\{1,2\}$. Then $\kappa^* = \frac{2}{3}$ and any $\boldsymbol{\mu}$ is optimal. Later, when we return to the label noise problem, we give sufficient conditions, analogous to those in the binary case, for the unicity of $\boldsymbol{\mu}$, so that the projection is well defined.

Observe that the following equivalent characterizations for κ^* hold:

$$\begin{aligned}
&\kappa^*(F_0|F_1,\ldots,F_M) \\
&\quad = \max_{\boldsymbol{\mu}\in S_M} \inf_{C\in\mathcal{C}:F_{\boldsymbol{\mu}}(C)>0} \frac{F_0(C)}{F_{\boldsymbol{\mu}}(C)} \qquad (7)\\
&\quad = \max\Big\{\kappa\in[0,1]\Big|\exists \boldsymbol{\mu}\in S_M \text{ and a distribution } G \\
&\qquad \text{s.t. } F_0 = (1-\kappa)G + \kappa F_{\boldsymbol{\mu}}\Big\} \qquad (8)\\
&\quad = \max\Big\{\sum_{i=1}^M \nu_i \,\Big|\, \nu_i\geq 0, \sum_{i=1}^M \nu_i \leq 1, \text{ and} \\
&\qquad \exists \text{ a distribution } G \text{ s.t.} \\
&\qquad F_0 = \left(1-\sum_{i=1}^M \nu_i\right) G + \sum_{i=1}^M \nu_i F_i\Big\}. \qquad (9)
\end{aligned}$$

The equivalence of (6), (7) and (8) is a straightforward consequence of the equivalent definitions (2) and (3) in the 2-class case. The equivalence of (8) and (9) is also clear from the representation $\kappa = \sum_i \nu_i$ and $\boldsymbol{\mu} = (\nu_1,\ldots,\nu_M)/\kappa$. The fact that the maxima are well-defined is justified formally in the supplemental material. Finally, if the maximum is attained for a unique G in (8)-(9), we refer to G as the *residue* of F_0 w.r.t. $F_1,\ldots,F_M$.

Suppose that for $m = 0,1,\ldots,M$, we observe

$$X_1^m,\ldots,X_{n_m}^m \overset{iid}{\sim} F_m$$

The random samples need not be independent of each other. Let $\widehat{F}_i$ denote the corresponding empirical distributions. By way of notation, set $\boldsymbol{n} := (n_0,n_1,\ldots,n_M)$ and write $\boldsymbol{n}\to\infty$ to indicate $\min_i\{n_i\}\to\infty$.

To define the estimator, let $\mathcal{C}_k \subset \mathcal{C}$, $k\geq 1$ be VC classes with VC dimensions $V_k < \infty$. For each $0\leq m\leq M, k\geq 1$ and $\delta>0$, define

$$\epsilon_m^k(\delta) := 3\sqrt{\frac{V_k\log(n+1) - \log\delta/2}{n_m}}.$$

By the VC inequality (Devroye et al., 1996), the following property holds:

P1 For each value of k, m, and $\delta>0$, the following holds with probability at least $1-\delta$:

$$\sup_{C\in\mathcal{C}_k}\left|F_m(C) - \widehat{F}_m(C)\right| \leq \epsilon_m^k(\delta),$$

where the probability is with respect to the draw of $X_1^m,\ldots,X_{n_m}^m \overset{iid}{\sim} F_m$.

We also require that $(\mathcal{C}_k)_{k\geq 1}$ possess the following universal approximation property:

P2 For any probability distribution Q, any $C^*\in\mathcal{C}$:

$$\liminf_{k\to\infty}\inf_{C\in\mathcal{C}_k} Q(C\Delta C^*) = 0,$$

where $C_1\Delta C_2 := (C_1\backslash C_2)\cup(C_2\backslash C_1)$ is the symmetric difference.

Examples of such classes include histograms, decision trees, neural networks, and generalized linear classifiers (Devroye et al., 1996).

Given $\boldsymbol{\delta} := (\delta_0,\ldots,\delta_M)$, a vector with positive components, we define the estimator

$$\begin{aligned}
&\widehat{\kappa}(\widehat{F}_0|\widehat{F}_1,\ldots,\widehat{F}_M;\boldsymbol{\delta}) \\
&\quad = \max_{\boldsymbol{\mu}\in S_M}\inf_k\inf_{C\in\mathcal{C}_k} \frac{\widehat{F}_0(C) + \epsilon_0^k(c\delta_0 k^{-2})}{\left(\widehat{F}_{\boldsymbol{\mu}}(C) - \sum_i \mu_i\epsilon_i^k(c\delta_i k^{-2})\right)_+},
\end{aligned} \qquad (10)$$

where $c := 6/\pi^2$, and the ratio is defined as $+\infty$ if the denominator is zero. Also, denote $\widehat{\boldsymbol{\mu}}$ an arbitrary point where the above maximum is attained; and finally $\widehat{\boldsymbol{\nu}} = \widehat{\kappa}\widehat{\boldsymbol{\mu}}$ (where the explicit dependence on $\widehat{F}_0$, etc., has been omitted to simplify notation). $\widehat{\boldsymbol{\nu}}$ is an estimate of the contaminating proportions $(\nu_1,\ldots,\nu_M)$ achieving the maximum in (9).

Proposition 2. *Let $(\mathcal{C}_k)_{k\geq 1}$ be a sequence of classes of sets with finite VC dimension and having the universal approximation property* **P2**.

- *(a) It holds that $\widehat{\kappa}\geq\kappa^*$ with probability at least $1-\sum_{i=0}^M\delta_i$.*
- *(b) If for all i, $\delta_i(\boldsymbol{n}) = 1/n_i$, then $\widehat{\kappa}$ converges in probability to κ^* as $\boldsymbol{n}\to\infty$, for any family of generating distributions $(F_i, i=0,\ldots,M)$.*
- *(c) Let $\mathcal{B}_* := \text{Arg Max}_{\boldsymbol{\mu}\in S_M}\kappa^*(F_0|F_{\boldsymbol{\mu}})$, i.e., $\mathcal{B}_*$ is the set of all mixture weights $\boldsymbol{\mu}$ attaining the max. in (6). Then under the assumptions of point (b), $d(\widehat{\boldsymbol{\mu}},\mathcal{B}_*)$ converges to zero in probability as $\boldsymbol{n}\to\infty$ (where d is any continuous distance function on the simplex). A similar result holds for the convergence in probability of $\widehat{\boldsymbol{\nu}}$ to the set of weight vectors attaining κ^* in (9).*

The definition of the estimator (10) is an intuitive generalization of the consistent estimator $\widehat{\kappa}(\widehat{F}_0|\widehat{F}_1)$ of

Blanchard et al. (2010), with an added maximum operation over $\boldsymbol{\mu} \in S_M$. The proof of the above result, however, does not follow trivially from the pointwise consistency of $\widehat{\kappa}(\widehat{F}_0|\widehat{F}_{\boldsymbol{\mu}})$ to $\kappa^*(F_0|F_{\boldsymbol{\mu}})$ for all fixed $\boldsymbol{\mu} \in S_M$, and requires a careful compactness argument. Furthermore, point (c) entails that the weights defining the residue are estimated consistently if they are unique (i.e., when $\mathcal{B}_*$ contains a single element), which they are in the label noise setting under the sufficient conditions given below.

Practical feasibility. The focus of this work is on identifiability and existence of a consistent estimator. While a suitable practical implementation is left for future work, we observe that there exist a priori reasonable strategies to compute $\widehat{\kappa}$. Because the $\epsilon^k(\ldots)$ terms become eventually larger than 1 for large k, the infimum over k can be limited to $k \leq k_{max}(n)$. The inner infimum loop can then be computed exactly if the classes $\mathcal{C}_k$ are finite (for instance, this is the case for the pieces of dyadic regular partitions of order k). The outer maximum loop can in turn be solved by gradient ascent (in alternating steps with the infimum solving step), which will converge to the global maximum since the inverse of the objective function is a linear function of $\boldsymbol{\mu}$.

Henceforth $\boldsymbol{\delta}$ will be omitted from the notation for $\widehat{\kappa}$, and δ_i taken to equal $1/n_i$.

4 Multiclass Label Noise

The model in (1) may be concisely expressed as $\tilde{\boldsymbol{P}} = \Pi \boldsymbol{P}$, where

$$\boldsymbol{P} = \begin{bmatrix} P_1 \\ \vdots \\ P_L \end{bmatrix}, \quad \tilde{\boldsymbol{P}} = \begin{bmatrix} \tilde{P}_1 \\ \vdots \\ \tilde{P}_L \end{bmatrix},$$

and $\Pi = [\pi_{ij}]$ is an $L \times L$ matrix with nonnegative entries and rows summing to one. We begin with identifiability assumptions that enable decontamination.

4.1 Assumptions

We start with a generalization of mutual irreducibility:

Lemma 1. *The following conditions on the family of distributions $P_1, \ldots, P_L$ are equivalent:*

- *It is not possible to write*

$$\sum_{i \in I} \epsilon_i P_i = \alpha \left(\sum_{i \notin I} \epsilon_i P_i \right) + (1 - \alpha) H,$$

  where $I \subset \{1, \ldots, L\}$ such that $1 \leq |I| < L$, ϵ_i are such that $\epsilon_i \geq 0$ and $\sum_{i \in I} \epsilon_i = \sum_{i \notin I} \epsilon_i = 1$, $\alpha \in (0, 1]$, and H is a distribution.

- *If $\sum_{i=1}^{L} \gamma_i P_i$ is a distribution, then $\gamma_i \geq 0 \ \forall i$.*

Definition 2. *We say the distributions $\{P_i\}_{1 \leq i \leq L}$ are* jointly irreducible *iff the conditions in Lemma 1 hold.*

Joint irreducibility says that every convex combination of some portion of the P_is is irreducible w.r.t. every convex combination of the remaining P_is. It generalizes the notion of mutual irreducibility from the two-class case (see Sec. 2). One case where it holds is when the support of each P_i contains some region with positive probability that does not intersect the supports of the other $P_j, j \neq i$. If each P_i is a class-conditional distribution, this means that every class has some exemplars that could not possibly arise from another class. This assumption is not unreasonable in many, and perhaps most, applications of interest. Joint irreducibility can still hold even when all P_i have the same support, as in the case of Gaussian densities with a common variance (Scott et al., 2013).

We will also make assumptions on the contamination weight matrix Π. Let $\boldsymbol{\pi}_i$ be the transpose of the i-th row of Π, which is a discrete probability distribution on $\{1, \ldots, L\}$. Let $\boldsymbol{e}_i$ denote the length L vector with 1 in the ith position and zeros elsewhere.

Lemma 2. *The following conditions on $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_L$ are equivalent:*

(a) For each ℓ, the residue of $\boldsymbol{\pi}_\ell$ with respect to $\{\boldsymbol{\pi}_j, j \neq \ell\}$ is $\boldsymbol{e}_\ell$.

(b) For every ℓ there exists a decomposition $\boldsymbol{\pi}_\ell = \kappa_\ell \boldsymbol{e}_\ell + (1 - \kappa_\ell)\boldsymbol{\pi}'_\ell$ where $\kappa_\ell > 0$ and $\boldsymbol{\pi}'_\ell$ is a convex combination of $\boldsymbol{\pi}_j$ for $j \neq \ell$.

(c) Π is invertible and Π^{-1} is a matrix with strictly positive diagonal entries and nonpositive off-diagonal entries.

Definition 3. *We say that Π is* recoverable *iff the conditions in Lemma 2 hold.*

This assumption ensures that the amount of contamination is not too high. Some intuition is given by condition (b). Fig. 1 depicts the case $L = 3$. In panel (i), condition (b) is satisfied, while in panel (ii), condition (b) is not satisfied. Clearly, in panel (i), the diagonal entries $\pi_{\ell\ell}$ are much larger than they are in panel (ii), and consequently the off-diagonal entries of Π will be smaller. Note that for $L = 2$ classes, from condition (c), recoverability is equivalent to $\pi_{12} + \pi_{21} < 1$, the condition developed by Scott et al. (2013).

We exhibit a setting where Π is guaranteed to be recoverable. Assume $\boldsymbol{c} := \sum_{i=1}^{M} c_i \boldsymbol{e}_i$ is a contaminating "background noise" which is common to all observed classes, albeit in possibly different proportions, i.e.,

$\boldsymbol{\pi}_i = \kappa_i \boldsymbol{c} + (1-\kappa_i)\boldsymbol{e}_i$, with $\kappa_i \in [0,1)$. Geometrically, this means we shift by various amounts the vertices $\boldsymbol{e}_i$ of the simplex towards a common point $\boldsymbol{c}$, or equivalently that each $\boldsymbol{\pi}_\ell$ belongs to the segment $(\boldsymbol{c}\boldsymbol{e}_\ell]$. Panel (iii) illustrates this situation.

Then we have

$$\sum_i \frac{c_i}{1-\kappa_i}\boldsymbol{\pi}_i = \left(1 + \sum_i \frac{c_i\kappa_i}{1-\kappa_i}\right)\boldsymbol{c},$$

so that for instance for $\ell = 1$ and $\kappa_1 > 0$,

$$\frac{c_1}{1-\kappa_1}\boldsymbol{\pi}_1 + \sum_{i\geq 2}\frac{c_i}{1-\kappa_i}\boldsymbol{\pi}_i = \left(1 + \sum_i \frac{c_i\kappa_i}{1-\kappa_i}\right)\left(\frac{1}{\kappa_1}\boldsymbol{\pi}_1 - \frac{1-\kappa_1}{\kappa_1}\boldsymbol{e}_1\right),$$

and finally

$$\left(1 + \sum_i \frac{c_i\kappa_i}{1-\kappa_i}\right)(1-\kappa_1)\boldsymbol{e}_1 + \kappa_1 \sum_{i\geq 2}\frac{c_i}{1-\kappa_i}\boldsymbol{\pi}_i = \boldsymbol{\pi}_1\left(1 + \sum_{i\geq 2}\frac{c_i\kappa_i}{1-\kappa_i}\right),$$

which is also valid when $\kappa_1 = 0$. Since all of the above coefficients are positive, this implies condition (c) of Lemma 2 after normalization.

In Sec. 5 we consider various applications where the recoverability assumption is satisfied.

4.2 Decontamination

The following result shows that the recoverability and joint irreducibility conditions ensure decontamination in the population (infinite sample) case. This result is applied in the next subsection, in conjunction with the estimator of Section 3, to establish a consistent discrimination rule.

Proposition 3. *If Π is recoverable and $P_1, \ldots, P_L$ are jointly irreducible, then for each ℓ, P_ℓ is the residue of $\tilde{P}_\ell$ w.r.t. $\{\tilde{P}_j, j \neq \ell\}$. Furthermore, in the representation*

$$\tilde{P}_\ell = (1-\kappa_\ell)P_\ell + \sum_{j\neq\ell}\nu_{\ell j}\tilde{P}_j, \tag{11}$$

where $\kappa_\ell = \kappa^*(\tilde{P}_\ell \,|\, \{\tilde{P}_j, j \neq \ell\})$, $\kappa_\ell < 1$ and the $\nu_{\ell j}$ are unique.

The proof of this result shows that under the joint irreducibility assumption, decontamination of the $\tilde{P}_\ell$ is equivalent to decontamination of the discrete distributions $\boldsymbol{\pi}_\ell$, which may be viewed as contaminated versions of the $\boldsymbol{e}_i$. In other words, the same weights κ_ℓ and $\nu_{\ell j}$ uniquely give the solutions of both approximation problems. The desired solution of the discrete problem is guaranteed by the recoverability assumption on Π (so that $\boldsymbol{e}_\ell$ is the residue of $\boldsymbol{\pi}_\ell$ w.r.t. $\{\boldsymbol{\pi}_j\}_{j\neq\ell}$), and this ensures, by the equivalence of the decontamination problems, that P_ℓ is the residue of $\tilde{P}_\ell$ w.r.t. $\{\tilde{P}_j, j \neq \ell\}$.

This equivalence leads to a second insight: By a construction in the proof of Lemma 2, Π^{-1} can be expressed explicitly in terms of the optimal weights κ_ℓ and $\nu_{\ell j}$. Therefore, under the assumptions of Proposition 3, Π^{-1} can be consistently estimated (i.e., recovered) via the estimator in Section 3.

4.3 A Consistent Discrimination Rule

The decontamination result gives a way to accurately estimate the true class-conditional error probabilities, which can then be converted into a consistent discrimination rule. These details now follow. For any classifier, that is, any measurable function $f : \mathcal{X} \to \{1, \ldots, L\}$, denote $R_i(f) := P_i(f(X) \neq i)$, where X follows P_i. As a performance measure we adopt the *minmax* criterion, which seeks to minimize

$$R(f) := \max_{1\leq i\leq L} R_i(f).$$

The optimal performance is the *minmax error*, $R^* := \inf_f R(f)$, where the inf is over all classifiers. Note that other performance measures could also be analyzed; we focus on the minmax criterion for concreteness.

The crux of classifier design is accurate error estimation. Denote $\tilde{R}_{j\ell}(f) := \tilde{P}_j(f(X) \neq \ell)$. To motivate an estimator for $R_\ell(f)$, the expression in (11) implies

$$R_\ell(f) = \frac{\tilde{R}_{\ell\ell}(f) - \sum_{j\neq\ell}\nu_{\ell j}\tilde{R}_{j\ell}(f)}{1-\kappa_\ell}.$$

If $X_1^j, \ldots, X_{n_j}^j \overset{iid}{\sim} \tilde{P}_j$ are the observed data from class j, then

$$\widehat{\tilde{R}}_{j\ell}(f) := \frac{1}{n_j}\sum_{i=1}^{n_j}\mathbf{1}_{\{f(X_i^j)\neq\ell\}}$$

estimates $\tilde{R}_{j\ell}(f)$. Now let $(\widehat{\nu}_{\ell j})_{j\neq\ell}$ be any vector achieving the maximum in the definition of $\widehat{\kappa}_\ell := \widehat{\kappa}(\widehat{\tilde{P}}_\ell | \{\widehat{\tilde{P}}_j, j \neq \ell\})$, so that $\widehat{\kappa}_\ell = \sum_{j\neq\ell}\widehat{\nu}_{\ell j}$. By Proposition 2, this vector converges to the unique weights $\nu_{\ell j}$ in Proposition 3, motivating the following estimator:

$$\widehat{R}_\ell(f) := \frac{\widehat{\tilde{R}}_{\ell\ell}(f) - \sum_{j\neq\ell}\widehat{\nu}_{\ell j}\widehat{\tilde{R}}_{j\ell}(f)}{1-\widehat{\kappa}_\ell}.$$

It is well known in statistical learning theory that consistency of a learning algorithm follows from uniform

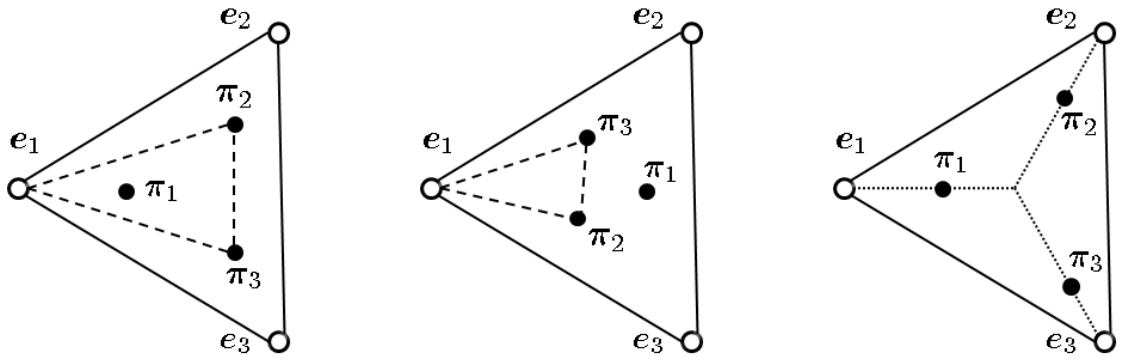

Figure 1: Illustration of the recoverability condition when $L = 3$. Panel (i): Low noise, Π recoverable. Each $\boldsymbol{\pi}_\ell$ can be written as a convex combination of $\boldsymbol{e}_\ell$ and the other two $\boldsymbol{\pi}_j$ (with a positive weight on $\boldsymbol{e}_\ell$), depicted here for $\ell = 1$. Panel (ii): High noise, Π not recoverable. Panel (iii): The setting of "common background noise."

control of error estimates over a class of classifiers $\mathcal{F}$, whose complexity may grow with sample size. Thus we introduce a sequence $(\mathcal{F}_k)_{k\geq 1}$ of sets of classifiers with VC dimension $V_k < \infty$. Since we are in a multiclass setting, we adopt the following generalization of VC dimension to multiclass, namely, define the VC dimension of class $\mathcal{F}$ to be the maximum (conventional) VC dimension of the family of sets $\{x : f(x) \neq \ell\}_{f\in\mathcal{F}}$, over $\ell = 1, \ldots, L$. This particular multiclass generalization of VC dimension allows our analysis to proceed using only the standard VC inequality for binary classifiers (or equivalently, for sets, as in Section 3). Indeed, the VC inequality (Devroye et al., 1996) gives uniform convergence of $\widehat{\tilde{R}}_{j\ell}(f)$ to $\tilde{R}_{j\ell}(f)$ over f in $\mathcal{F}_k$, and together with consistency of $\widehat{\kappa}_\ell$ and $(\widehat{\nu}_{\ell j})$ implies the following result. Similar to Section 3, set $\boldsymbol{n} := (n_1, \ldots, n_L)$ and write $\boldsymbol{n} \to \infty$ to indicate $\min\{n_\ell\} \to \infty$.

Proposition 4. *Let $k(\boldsymbol{n})$ take positive integer values and be such that as $\boldsymbol{n} \to \infty$,*

$$\frac{V_{k(\boldsymbol{n})} \log n_\ell}{n_\ell} \to 0,$$

$1 \leq \ell \leq L$. Then under the assumptions of Proposition 3, $\sup_{f\in\mathcal{F}_{k(\boldsymbol{n})}} |R_\ell(f) - \widehat{R}_\ell(f)| \overset{i.p.}{\to} 0$ as $\boldsymbol{n} \to \infty$.

This result allows us to analyze the estimation error of a learning algorithm based on $\widehat{R}_\ell$. To control the approximation error, we choose $(\mathcal{F}_k)$ such that

P3 For any $P_1, \ldots, P_L$, $\lim_{k\to\infty} \inf_{f\in\mathcal{F}_k} R(f) = R^*$.

This condition is analogous to the condition for our family of sets $(\mathcal{C}_k)$ in Section 3.

Let us now define a discrimination rule based on the above error estimates. Define $\widehat{R}(f) := \max_\ell \widehat{R}_\ell(f)$. Let τ_k be any sequence of positive numbers tending to zero. Let $\widehat{f}_k$ denote any classifier

$$\widehat{f}_k \in \Big\{ f \in \mathcal{F}_k \,:\, \widehat{R}(f) \leq \inf_{f\in\mathcal{F}_k} \widehat{R}(f) + \tau_k \Big\}.$$

The introduction of τ_k lets us avoid assuming the existence of an empirical risk minimizer. Finally, define the discrimination rule $\widehat{f} := \widehat{f}_{k(\boldsymbol{n})}$.

Theorem 1. *Let $k(\boldsymbol{n})$ take positive integer values and be such that as $\boldsymbol{n} \to \infty$, $k(\boldsymbol{n}) \to \infty$ and*

$$\frac{V_{k(\boldsymbol{n})} \log n_\ell}{n_\ell} \to 0, \tag{12}$$

$1 \leq \ell \leq L$. Then under **P3** *and the assumptions of Proposition 3, $R(\widehat{f}) \overset{i.p.}{\to} R^*$ as $\boldsymbol{n} \to \infty$.*

Thus, if the noise is recoverable and the distributions jointly irreducible, consistent classification is possible in the multiclass label noise setting.

5 Discussion

We now develop connections between the above theoretical framework for multiclass label noise and various applications mentioned in the introduction.

Consider the second crowdsourcing formulation, and suppose that unlabeled data are drawn according to $\sum_{j=1}^L \theta_j P_j$. Further suppose that with probability $1-\alpha$, the label is assigned correctly, while with probability α, the label is assigned according to some fixed distribution, say the uniform distribution $[\frac{1}{L}, \ldots, \frac{1}{L}]^T$. Let Y denote the true label of an example, and $\tilde{Y}$ the crowdsourced label. By Bayes' rule,

$$\begin{aligned}\pi_{ij} &= Pr(Y = j \,|\, \tilde{Y} = i) \\ &\propto Pr(\tilde{Y} = i \,|\, Y = j) Pr(Y = j) \\ &= ((1-\alpha)\mathbf{1}_{\{i=j\}} + \frac{\alpha}{L})\theta_j.\end{aligned}$$

In vector form, this means $\boldsymbol{\pi}_i \propto (1-\alpha)\theta_i \boldsymbol{e}_i + \frac{\alpha}{L}\boldsymbol{\theta}$, where $\boldsymbol{\theta} = [\theta_1, \ldots, \theta_L]^T$. This clearly satisfies the common background noise model described previously, and therefore our noise condition is satisfied.

Next, consider the semi-supervised novelty detection problem described in the introduction. In this case

$$\Pi = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & \ddots & \\ \pi_{L1} & \pi_{L2} & \cdots & \pi_{LL} \end{bmatrix}$$

because the first $L-1$ data samples are uncontaminated, while the last sample is an unlabeled testing sample where P_L governs the novelty class. It is easy to see that Π is recoverable if and only if $\pi_{LL} > 0$, using condition (c) of Lemma 2. Therefore, as long as there are at least some novel exemplars in the unlabeled data, the noise condition is satisfied, and under joint irreducibility we have a consistent discrimination rule for multiclass, semi-supervised novelty detection. We also note that the above is not the only approach to novelty detection using mixture proportion estimation (Sanderson and Scott, 2014).

Our work has an interesting connection to the problem of topic modeling, itself linked to nonnegative matrix factorization. In that problem, one also observes random samples from several contaminated distributions ("documents") $\tilde{P}_i$, but in the equation $\tilde{\boldsymbol{P}} = \Pi \boldsymbol{P}$, the matrix Π is no longer square but now has far more rows than columns. The distributions P_i are "topics," and the proportion π_{ij} reflects the prevalence of topic j in document i. Existing work on topic modeling typically represents the topics and documents as discrete distributions on a finite vocabulary. One interesting connection to our work is that in this discrete setting, our joint irreducibility condition is equivalent to a condition that has been previously shown to be sufficient for identifiability of the topics. In particular, in the discrete setting, this assumption states that for each topic, there exists at least one word occuring with a positive probability in the given topic, and with a probability of zero in the other topics (Donoho and Stodden, 2004; Arora et al., 2012).

Finally, we return to the problem of learning from partial labels (Cour et al., 2011). For the sake of argument we consider a simple form of the problem. Suppose there are $L = 3$ classes, and that each observed data point is labeled $A = \{1, 2\}, B = \{1, 3\}$, or $C = \{2, 3\}$. The true label of each example is one of the labels in the associated subset. Grouping together observations according to the three "partial labels" (A, B, or C), these three data samples are described by the contamination models $\tilde{P}_S = \sum_{j \in S} \pi_{S,j} P_j$, where

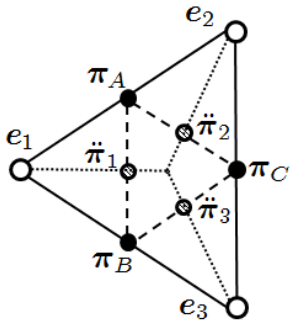

Figure 2: Learning from partial labels reduces to multiclass label noise in some cases.

$S \in \{A, B, C\}$. Each data sample arises from a convex combination of two of the three base distributions, as depicted in Fig. 2. Since these contaminated models lie on the boundary of the probability simplex, they do not satisfy the recoverability assumption on the noise. However, by resampling the data, we may obtain random samples from distributions that do satisfy the noise assumption. For simplicity, let's assume $\pi_{S,j} = 1/2$ for each $j \in S$, i.e., each contaminated model is an equal mixture of the two associated base distributions. For every pair of contaminated models, corresponding to subsets S and T, form a new random sample by resampling from the two observed samples, such that the new random samples are realizations of $\frac{1}{2}\tilde{P}_S + \frac{1}{2}\tilde{P}_T$. If Π and $\ddot{\Pi}$ denote the contamination proportions before and after resampling, then

$$\Pi = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}, \quad \ddot{\Pi} = \begin{bmatrix} \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{2} \end{bmatrix}.$$

The situation is depicted in Fig. 2. The rows $(\ddot{\boldsymbol{\pi}}_i)^T$ of $\ddot{\Pi}$ satisfy $\ddot{\boldsymbol{\pi}}_i = \frac{1}{4}\boldsymbol{e}_i + \frac{3}{4}\boldsymbol{c}$, where $\boldsymbol{c} = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]^T$, which is in the form of the common background noise model, implying our noise condition. Our framework can now be applied to the resampled data, assuming joint irreducibility, to give a consistent discrimination rule for learning from partial labels. If Π is perturbed slightly, our noise assumption still holds, and we further conjecture that this kind of resampling strategy may be applied when $L > 3$ and when the pattern of observed subsets is more general.

Acknowledgements

G. Blanchard was supported in part by the European Community's 7th Framework Programme under the E.U. grant agreement 247022 (MASH Project). C. Scott was supported in part by NSF Grants 0953135, 1047871, and 1217880.

References

D. Aldous and P. Diaconis. Strong uniform times and finite random walks. *Adv. Appl. Math.*, 8(1):69–97, 1987.

S. Arora, R. Ge, and A. Moitra. Learning topic models – going beyond SVD. In *FOCS*. 2012.

G. Blanchard, G. Lee, and C. Scott. Semi-supervised novelty detection. *Journal of Machine Learning Research*, 11:2973–3009, 2010.

C. Bouveyron and S. Girard. Robust supervised classification with mixture models: Learning from data with uncertain labels. *Journal of Pattern Recognition*, 42:2649–2658, 2009.

T. Cour, B. Sapp, and B. Taskar. Learning from partial labels. *J. Machine Learning Research*, 12:1501–1536, 2011.

L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, 1996.

D. Donoho and V. Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.

M. C. Du Plessis and M. Sugiyama. Semi-supervised learning of class balance under class-prior change by distribution matching. In J. Langford and J. Pineau, editors, *Proc. 29th Int. Conf. on Machine Learning*, pages 823–830, 2012.

P. Hall. On the non-parametric estimation of mixture proportions. *Journal of the Royal Statistical Society*, 43(2):147–156, 1981.

P. Latinne, M. Saerens, and C. Decaestecker. Adjusting the outputs of a classier to new a priori probabilities may signicantly improve classication accuracy: Evidence from a multi-class problem in remote sensing. In C. Sammut and A. H. Hoffmann, editors, *Proc. 18th Int. Conf. on Machine Learning*, pages 298–305, 2001.

N. Lawrence and B. Schölkopf. Estimating a kernel Fisher discriminant in the presence of label noise. *Proceedings of the International Conference in Machine Learning*, 2001.

P. Long and R. Servido. Random classification noise defeats all convex potential boosters. *Machine Learning*, 78:287–304, 2010.

N. Manwani and P. S. Sastry. Noise tolerance under risk minimization. Technical Report arXiv:1109.5231, 2011.

N. Natarajan, I. S. Dhillon, P. Ravikumar, and A. Tewari. Learning with noisy labels. In *Advances in Neural Information Processing Systems 26*, 2013.

V. C. Raykar, S. Yu, L. H. Zhao, G. H. Valadez, C. Florin, L. Bogoni, and L. Moy. Learning from crowds. *The Journal of Machine Learning Research*, 99:1297–1322, 2010.

T. Sanderson and C. Scott. Class proportion estimation with application to multiclass anomaly rejection. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014.

C. Scott, G. Blanchard, and G. Handy. Classification with asymmetric label noise: Consistency and maximal denoising. In *Proc. Conf. on Learning Theory,* JMLR W&CP, volume 30, pages 489–511. 2013.

G. Stempfel and L. Ralaivola. Learning SVMs from sloppily labeled data. In *Proc. 19th Int. Conf. on Artificial Neural Networks: Part I*, pages 884–893, 2009.

D. M. Titterington. Minimum distance nonparametric estimation of mixture proportions. *Journal of the Royal Statistical Society*, 45(1):37–46, 1983.

Supplemental Material for the AISTATS 2014 paper "Decontamination of Mutually Contaminated Models"

Gilles Blanchard and Clayton Scott

A Proofs for Section 3

A.1 κ^* and $\widehat{\kappa}$ are well-defined

Lemma A.1. *The maximum operation in the definition of κ^* and $\widehat{\kappa}$ is well-defined, that is, the outside supremum is attained at at least one point.*

We prove the statement for

$$\kappa^* = \max_{\boldsymbol{\mu}} \inf_{C\in\mathcal{C}:F_{\boldsymbol{\mu}}(C)>0} \frac{F_0(C)}{F_{\boldsymbol{\mu}}(C)}.$$

The argument for $\widehat{\kappa}$ is similar. Denote $G(\boldsymbol{\mu}) = \kappa^*(F_0|F_{\boldsymbol{\mu}}) = \inf_{C\in\mathcal{C}:F_{\boldsymbol{\mu}}(C)>0} \frac{F_0(C)}{F_{\boldsymbol{\mu}}(C)}$ the maximum proportion of the mixture $F_{\boldsymbol{\mu}}$ in the distribution F_0.

We argue that G is an upper semicontinuous function. To see this, define for each $C \in \mathcal{C}$ the function $g_C : S_M \to [0,\infty]$ as

$$g_C(\boldsymbol{\mu}) := \begin{cases} \frac{F_0(C)}{F_{\boldsymbol{\mu}}(C)} & \text{if } F_{\boldsymbol{\mu}}(C) > 0\,; \\ +\infty & \text{if } F_{\boldsymbol{\mu}}(C) = 0. \end{cases}$$

Then f_C is an upper semicontinuous function: if $\boldsymbol{\mu} \in S_M$ is such that $F_{\boldsymbol{\mu}}(C) > 0$, then f_C is continuous at point $\boldsymbol{\mu}$. Otherwise, $f_C(\boldsymbol{\mu}) = \infty$ and f_C is trivially upper semicontinuous at point $\boldsymbol{\mu}$. Clearly, one has $G(\boldsymbol{\mu}) = \inf_{C\in\mathcal{C}} f_C(\boldsymbol{\mu})$; as an infimum of upper semicontinuous functions, it is itself upper semicontinuous, and therefore attains its maximum on the compact set S_M.

A.2 Proof of Proposition 2

Point (a): We apply condition **P1** for all k, i with $\delta_{k,i} = c\delta_i/k^2$. By the union bound, with probability at least $1 - \sum_{i=0}^{M} \delta_i$, it holds simultaneously for all $k \geq 1$ and $i = 0, \ldots, M$ that

$$\forall k \geq 1\,, \qquad \forall i \in \{0,\ldots,M\} : \qquad \sup_{C\in\mathcal{C}_k} \left| F_i(C) - \widehat{F}_i(C) \right| \leq \epsilon_i^k(c\delta_i k^{-2}) \quad \text{(S.1)}$$

Recall the notation (from the proof of Lemma A.1) $G(\boldsymbol{\mu}) = \inf_{C\in\mathcal{C}:F_{\boldsymbol{\mu}}(C)>0} \frac{F_0(C)}{F_{\boldsymbol{\mu}}(C)}$ and introduce

$$\widehat{G}(\boldsymbol{\mu}) := \inf_k \inf_{C\in\mathcal{C}_k} \frac{\widehat{F}_0(C) + \epsilon_0^k(c\delta_0 k^{-2})}{\left(\widehat{F}_{\boldsymbol{\mu}}(C) - \sum_i \nu_i \epsilon_i^k(c\delta_i k^{-2})\right)_+}.$$

Observe that when (S.1) is satisfied, this implies that for all $\boldsymbol{\mu} \in S_M$, one has $G(\boldsymbol{\mu}) \leq \widehat{G}(\boldsymbol{\mu})$. Taking the maximum over $\boldsymbol{\mu}$ yields the first point.

Point (b): let $\epsilon > 0$ be an arbitrary positive constant. For any $\boldsymbol{\mu} \in S_M$, let $C_{\boldsymbol{\mu}} \in \mathcal{C}$ with $F_{\boldsymbol{\mu}}(C_{\boldsymbol{\mu}}) > 0$ be such that $\frac{F_0(C_{\boldsymbol{\mu}})}{F_{\boldsymbol{\mu}}(C_{\boldsymbol{\mu}})} \leq \kappa^* + \epsilon/4$.

By continuity of the function $\boldsymbol{\mu} \mapsto F_{\boldsymbol{\mu}}(C)$ for any fixed C, there exists for each $\boldsymbol{\mu} \in S_M$ an open neighborhood $N_{\boldsymbol{\mu}}$ of $\boldsymbol{\mu}$ for which both of the following conditions are realized for all $\boldsymbol{\mu'} \in N_{\boldsymbol{\mu}}$:

$$\frac{F_0(C_{\boldsymbol{\mu}})}{F_{\boldsymbol{\mu'}}(C_{\boldsymbol{\mu}})} \leq \kappa^* + \frac{\epsilon}{2}, \tag{S.2}$$

$$\text{and } F_{\boldsymbol{\mu'}}(C_{\boldsymbol{\mu}}) \geq \frac{1}{2} F_{\boldsymbol{\mu}}(C_{\boldsymbol{\mu}}). \tag{S.3}$$

(For the second condition, we have used the fact that $F_{\boldsymbol{\mu}}(C_{\boldsymbol{\mu}}) > 0$). By compactness of S_M, there exists a finite subset S_M^ϵ of S_M such that $(N_{\boldsymbol{\mu}})_{\boldsymbol{\mu}\in S_M^\epsilon}$ covers S_M.

Denote $F_{\min}^\epsilon := \frac{1}{2}\min_{\boldsymbol{\mu}\in S_M^\epsilon} F_{\boldsymbol{\mu}}(C_{\boldsymbol{\mu}})$; it is a positive quantity since $F_{\boldsymbol{\mu}}(C_{\boldsymbol{\mu}}) > 0$ for any $\boldsymbol{\mu}$, and S_M^ϵ is finite. For each $\boldsymbol{\mu} \in S_M$, denote $\zeta(\boldsymbol{\mu})$ an arbitrary element of the finite net S_M^ϵ such that $\boldsymbol{\mu} \in N_{\zeta(\boldsymbol{\mu})}$. By property (S.2), we have

$$\sup_{\boldsymbol{\mu}\in S_M} \frac{F_0(C_{\zeta(\boldsymbol{\mu})})}{F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})})} \leq \max_{\boldsymbol{\mu}\in S_M^\epsilon} \sup_{\boldsymbol{\mu'}\in N_{\boldsymbol{\mu}}} \frac{F_0(C_{\boldsymbol{\mu}})}{F_{\boldsymbol{\mu'}}(C_{\boldsymbol{\mu}})} \leq \kappa^* + \frac{\epsilon}{2}, \tag{S.4}$$

and by property (S.3):

$$\inf_{\boldsymbol{\mu}\in S_M} F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})}) \geq \min_{\boldsymbol{\mu}\in S_M^\epsilon} \inf_{\boldsymbol{\mu'}\in N_{\boldsymbol{\mu}}} F_{\boldsymbol{\mu'}}(C_{\boldsymbol{\mu}}) \geq F_{\min}^\epsilon. \tag{S.5}$$

Denote $\mathcal{C}_\epsilon := \{C_{\boldsymbol{\mu}}\,,\ \boldsymbol{\mu} \in S_M^\epsilon\}$. Let $\eta \in (0, F_{\min}^\epsilon/2)$ be another arbitrary positive constant. Consider the distribution $Q = \frac{1}{M+1}\sum_{i=0}^M F_i$, to which we apply condition **P2**. This entails that for each individual $C \in \mathcal{C}$ there exists a k_C and $\widetilde{C} \in \mathcal{C}_{k_C}$ with

$$Q(C\Delta\widetilde{C}) \leq \frac{\eta}{M+1},$$

implying for all $i \in \{0, \ldots, M\}$:

$$\left|F_i(C) - F_i(\widetilde{C})\right| \leq F_i(C\Delta\widetilde{C}) \leq (M+1)Q(C\Delta\widetilde{C}) \leq \eta,$$

and then also for all $\boldsymbol{\mu} \in S_M$:

$$\left|F_{\boldsymbol{\mu}}(\widetilde{C}) - F_{\boldsymbol{\mu}}(C)\right| \leq \sum_{i=1}^M \mu_i \left|F_i(C) - F_i(\widehat{C})\right| \leq \eta.$$

In what follows we use the shortened notation $\varepsilon_i^k \equiv \epsilon_i^k(c\delta_i k^{-2})$, and further define $\underline{\varepsilon}(\epsilon,\eta) := \max_i \max_{C\in\mathcal{C}_\epsilon} \varepsilon_i^{k_C}$. For fixed (ϵ,η), the quantity $\underline{\varepsilon}(\epsilon,\eta)$ is defined as a maximum of a finite number of functions decreasing to 0 as $\boldsymbol{n}\to\infty$, and therefore $\underline{\varepsilon}$ also decreases to zero. Below, we assume that all components of $\boldsymbol{n}$ are chosen big enough so that $F^\epsilon_{\min} - \eta - 2\underline{\varepsilon}(\epsilon,\eta) > 0$. It holds with probability $1-\sum_{i=0}^M \delta_i$ that

$$
\begin{aligned}
\widehat{\kappa} &\le \sup_{\boldsymbol{\mu}\in S_M} \inf_k \inf_{C\in\mathcal{C}_k} \frac{F_0(C)+2\varepsilon_0^k}{\left(F_{\boldsymbol{\mu}}(C) - 2\sum_i \mu_i\varepsilon_i^k\right)_+} \\
&\le \sup_{\boldsymbol{\mu}\in S_M} \inf_{C\in\mathcal{C}} \frac{F_0(\widetilde{C})+2\varepsilon_0^{k_C}}{\left(F_{\boldsymbol{\mu}}(\widetilde{C}) - 2\sum_i \mu_i\varepsilon_i^{k_C}\right)_+} \\
&\le \sup_{\boldsymbol{\mu}\in S_M} \inf_{C\in\mathcal{C}} \frac{F_0(C)+\eta+2\varepsilon_0^{k_C}}{\left(F_{\boldsymbol{\mu}}(C) - \eta - 2\sum_i \mu_i\varepsilon_i^{k_C}\right)_+} \\
&\le \sup_{\boldsymbol{\mu}\in S_M} \frac{F_0(C_{\zeta(\boldsymbol{\mu})})+\eta+2\varepsilon_0^{k_{C_{\zeta(\boldsymbol{\mu})}}}}{\left(F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})}) - \eta - 2\sum_i \mu_i\varepsilon_i^{k_{C_{\zeta(\boldsymbol{\mu})}}}\right)_+} \\
&\le \sup_{\boldsymbol{\mu}\in S_M} \frac{F_0(C_{\zeta(\boldsymbol{\mu})})+\eta+2\underline{\varepsilon}(\epsilon,\eta)}{\left(F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})}) - \eta - 2\underline{\varepsilon}(\epsilon,\eta)\right)_+} \\
&\le \left(\sup_{\boldsymbol{\mu}\in S_M} \frac{F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})})}{\left(F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})}) - \eta - 2\underline{\varepsilon}(\epsilon,\eta)\right)_+}\right) \sup_{\boldsymbol{\mu}\in S_M} \frac{F_0(C_{\zeta(\boldsymbol{\mu})})+\eta+2\underline{\varepsilon}(\epsilon,\eta)}{F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})})} \\
&\le \left(\frac{F^\epsilon_{\min}}{\left(F^\epsilon_{\min} - \eta - 2\underline{\varepsilon}(\epsilon,\eta)\right)_+}\right)\left(\sup_{\boldsymbol{\mu}\in S_M} \frac{F_0(C_{\zeta(\boldsymbol{\mu})})}{F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})})} + \sup_{\boldsymbol{\mu}\in S_M} \frac{\eta+2\underline{\varepsilon}(\epsilon,\eta)}{F_{\boldsymbol{\mu}}(C_{\zeta(\boldsymbol{\mu})})}\right) \\
&\le \left(\frac{F^\epsilon_{\min}}{\left(F^\epsilon_{\min} - \eta - 2\underline{\varepsilon}(\epsilon,\eta)\right)_+}\right)\left(\kappa^* + \frac{\epsilon}{2}\right) + \frac{\eta+2\underline{\varepsilon}(\epsilon,\eta)}{\left(F^\epsilon_{\min} - \eta - 2\underline{\varepsilon}(\epsilon,\eta)\right)_+},
\end{aligned}
$$

where we have used (S.4) and (S.5) for the last inequality. By choosing first η small enough, then all components of $\boldsymbol{n}_0$ big enough, the r.h.s. of the above inequality can be made smaller than $\kappa^*+\epsilon$, for all $\boldsymbol{n}\succ\boldsymbol{n}_0$ ($\succ$ indicates the inequality holds for all components). Since $\sum_{i=0}^M \delta_i \to 0$ as $\boldsymbol{\mu}\to 0$, this implies the second part of the proposition.

For the last point of the proposition, consider an arbitrary open set Ω containing the set $\mathcal{B}^*$. Then $\Omega^c := S_M \setminus \Omega$ is a compact set; therefore, the function $G(\boldsymbol{\mu}) := \inf_{C\in\mathcal{C}, F_{\boldsymbol{\mu}}(C)>0} \frac{F_0(C)}{F_{\boldsymbol{\mu}}(C)}$, being upper semicontinuous (see proof of Lemma A.1), attains its supremum $\widetilde{\kappa}$ on Ω^c. Observe that $\widetilde{\kappa} > \kappa^*$ must hold, otherwise we would have a contradiction with the definition of $\mathcal{B}^*$. Finally, we have:

$$
\begin{aligned}
\mathbb{P}\left[\widehat{\boldsymbol{\mu}}\notin\Omega\right] &\le \mathbb{P}\left[\widehat{\boldsymbol{\mu}}\notin\Omega; G(\widehat{\boldsymbol{\mu}}) \le \widehat{G}(\widehat{\boldsymbol{\mu}})\right] + \mathbb{P}\left[G(\widehat{\boldsymbol{\mu}}) > \widehat{G}(\widehat{\boldsymbol{\mu}})\right] \\
&\le \mathbb{P}\left[\widehat{\kappa}\ge\widetilde{\kappa}\right] + \sum_{i=1}^M \delta_i,
\end{aligned}
$$

where we have used that $\widehat{\kappa} = \widehat{G}(\widehat{\boldsymbol{\mu}})$ by definition, and the argument used in the proof of point (a). By point (b), the first probability converges to 0 as $\boldsymbol{\mu} \to \infty$. Thus, the probability that $\widehat{\boldsymbol{\mu}} \in \Omega$ must converge to 1 as $\boldsymbol{n} \to \infty$. This applies in particular to any open set of the form $\Omega_\epsilon := \{\boldsymbol{\mu} : d(\boldsymbol{\mu}, \mathcal{B}^*) < \epsilon\}$, hence the conclusion.

B Proofs for Section 4

B.1 Proof of Lemma 1

Suppose the first condition does not hold, so that

$$\sum_{i \in I} \epsilon_i P_i = \alpha \left(\sum_{i \notin I} \epsilon_i P_i \right) + (1 - \alpha) H.$$

Then $\sum_i \gamma_i P_i = H$, where $\gamma_i = \frac{\epsilon_i}{1-\alpha}$ for $i \in I$, and $\gamma_i = -\frac{\alpha \epsilon_i}{1-\alpha}$ for $i \notin I$. Since $\sum_{i \notin I} \epsilon_i = 1$, at least one $\gamma_i < 0$, so the second condition is violated.

Now suppose the second condition is violated, say $\sum_i \gamma_i P_i = H$. Let $I = \{i \,|\, \gamma_i \geq 0\}$, which has fewer than K elements by assumption. Since $\sum_i \gamma_i = 1$, we also know $1 \leq |I|$ and further that $\Gamma := \sum_{i \in I} \gamma_i > 1$. A violation of the first condition is obtained by $\epsilon_i = \gamma_i / \Gamma$ for $i \in I$, $\epsilon_i = -\gamma_i / (\Gamma - 1)$ for $i \notin I$ (noting that $\sum_{i \notin I} (-\gamma_i) = \Gamma - 1$), and $\alpha = (\Gamma - 1)/\Gamma$.

B.2 Proof of Lemma 2

(a) $\Rightarrow$ (b): Follows immediately from the definition of the residue.

(b) $\Rightarrow$ (c): By assumption, there exists $\kappa > 0$ such that $\boldsymbol{\pi}_1 = \kappa \boldsymbol{e}_1 + (1 - \kappa) \boldsymbol{\eta}_1$, where $\boldsymbol{\eta}_1 = \sum_{i=2}^{L} \mu_i \boldsymbol{\pi}_i$, with $\mu_i \geq 0$, for all $2 \leq i \leq L$. Thus,

$$\boldsymbol{e}_1 = \kappa^{-1} \boldsymbol{\pi}_1 - \sum_{i=2}^{L} \frac{(1 - \kappa)}{\kappa} \mu_i \boldsymbol{\pi}_i \,;$$

a similar relation holds for all rows. This implies that Π is invertible and allows to identify (for instance) the first row of Π^{-1} as $(\kappa^{-1}, -\frac{(1-\kappa)}{\kappa}\mu_2, \ldots, -\frac{(1-\kappa)}{\kappa}\mu_L)$. This implies (c).

(c) $\Rightarrow$ (a): Without loss of generality, consider $\ell = 1$ and the problem of identifying $\kappa^*(\pi_1 | (\pi_i)_{2 \leq i \leq L})$, and the associated residue (if it exists). According to characterization (9), this corresponds to the optimization problem

$$\max_{\boldsymbol{\nu}, \boldsymbol{\gamma}} \sum_{i=2}^{L} \nu_i \;\; s.t. \;\; \boldsymbol{\pi}_1 = (1 - \sum_{i \geq 2} \nu_i) \boldsymbol{\gamma} + \sum_{i \geq 2} \nu_i \boldsymbol{\pi}_i,$$

over $\boldsymbol{\gamma} \in S_L$ and $\boldsymbol{\nu} = (\nu_2, \ldots, \nu_L) \in \Delta_{L-1} = \Big\{ (\nu_2, \ldots, \nu_L) | \nu_i \geq 0; \sum_{i=2}^{L} \nu_i \leq 1 \Big\}$.

We now reformulate this problem. First, note that the constraint implies that admissible $\boldsymbol{\nu}$ are such that $\sum_{i\geq 2}\nu_i < 1$, otherwise we would have a linear relation between the $\boldsymbol{\pi}_i$, contradicting invertibility of Π.

Then for an admissible $\boldsymbol{\nu}$, denote $\boldsymbol{\eta}(\boldsymbol{\nu}) := (1-\sum_{i\geq 2}\nu_i)^{-1}(1,-\nu_2,\ldots,-\nu_L)$. Observe that the constraint of the optimization problem is equivalent to $\Pi^T\boldsymbol{\eta} = \boldsymbol{\gamma}$, or $\boldsymbol{\eta} = (\Pi^T)^{-1}\boldsymbol{\gamma}$. The inverse mapping of $\boldsymbol{\eta}$ to $\boldsymbol{\nu}$ is $\boldsymbol{\nu}(\boldsymbol{\eta}) = \eta_1^{-1}(-\eta_2,\ldots,-\eta_L)$, so that the objective of the optimization can be rewritten as

$$-\frac{\sum_{i=2}^L \eta_i}{\boldsymbol{e}_1^T\boldsymbol{\eta}} = -\frac{\mathbf{1}^T\boldsymbol{\eta}}{\boldsymbol{e}_1^T\boldsymbol{\eta}} + 1 = 1 - \frac{1}{\boldsymbol{e}_1^T\boldsymbol{\eta}} = 1 - \frac{1}{\boldsymbol{e}_1^T(\Pi^T)^{-1}\boldsymbol{\gamma}},$$

where $\mathbf{1}$ denotes a L-dimensional vector with all coordinates equal to 1. So finding the point of maximum of the above problem is equivalent to the program

$$\max_{\boldsymbol{\gamma}\in S_L} \boldsymbol{e}_1^T(\Pi^T)^{-1}\boldsymbol{\gamma} \;\; s.t. \;\; \boldsymbol{\nu}((\Pi^T)^{-1}\boldsymbol{\gamma}) \in \Delta_{L-1}$$

The above objective function has the form $\boldsymbol{a}^T\boldsymbol{\gamma}$, where $\boldsymbol{a}$ is the first column of Π^{-1} which, by assumption, has its first coordinate positive and the others nonpositive. Therefore, the unconstrainted maximum over $\boldsymbol{\gamma} \in S_M$ is attained uniquely for $\boldsymbol{\gamma} = \boldsymbol{e}_1$. We now check that this value also satisfies the required constraint. Observe that $(\Pi^T)^{-1}\boldsymbol{e}_1$ is the (transpose of) the first row of Π^{-1}, denote this vector as $\boldsymbol{b} = (b_1,\ldots,b_L)$. We want to ensure that $\boldsymbol{\nu}(\boldsymbol{b}) = b_1^{-1}(-b_2,\ldots,-b_L) \in \Delta_{L-1}$. By assumption, $\boldsymbol{b}$ has its first coordinate positive and the others nonpositive, ensuring all components of $\boldsymbol{\nu}(\boldsymbol{b})$ are nonnegative. Furthermore, the sum of the components of $\boldsymbol{\nu}(\boldsymbol{b})$ is

$$\sum_{i=2}^L -\frac{b_i}{b_1} = 1 - \frac{\sum_{i=1}^L b_i}{b_1} = 1 - \frac{1}{b_1} \leq 1;$$

the last equality is because the rows of Π^{-1} sum to 1 (since Π is a stochastic matrix, so is its inverse). It follows that $\boldsymbol{\nu}((\Pi^T)^{-1}\boldsymbol{e}_1) \in \Delta_{L-1}$. Thus, the unique maximum of the optimization problem is attained for $\boldsymbol{\gamma} = \boldsymbol{e}_1$, establishing (a).

B.3 Proof of Proposition 3

We start with the following Lemma:

Lemma B.1. *If Π is recoverable, then $\boldsymbol{\pi}_1,\ldots,\boldsymbol{\pi}_L$ are linearly independent. If $P_1,\ldots,P_L$ are jointly irreducible, then they are linearly independent. If $\boldsymbol{\pi}_1,\ldots,\boldsymbol{\pi}_L$ are linearly independent and $P_1,\ldots,P_L$ are linearly independent, then $\tilde{P}_1,\ldots,\tilde{P}_L$ are linearly independent.*

Proof of the lemma: The first statement follows from characterization (c) of Lemma 2: if Π is recoverable, it is invertible and thus has full rank.

For the second statement, suppose $\sum_i \beta_i P_i = 0$ is a nontrivial linear relation. Let j be any index such that $\beta_j \geq 0$. Then $\sum_i \gamma_i P_i = P_j$, where $\gamma_i = \beta_i$ if $i \neq j$, and $\gamma_j = \beta_j + 1$. Since at least one $\beta_i < 0$, $i \neq j$, joint irreducibility is violated.

For the third part, suppose $\sum_i \alpha_i \tilde{P}_i = 0$. Since $\tilde{P}_i = \boldsymbol{\pi}_i^T \boldsymbol{P}$, this implies $\sum_i \alpha_i \boldsymbol{\pi}_i^T \boldsymbol{P} = 0$, which implies $\sum_i \alpha_i \boldsymbol{\pi}_i = \mathbf{0}$, which implies $\alpha_i = 0$.

Proof of Proposition 3: Consider $\ell = 1$, the other cases being similar. Suppose G is such that

$$\tilde{P}_1 = (1 - \sum_{j \geq 2} \nu_j) G + \sum_{j \geq 2} \nu_j \tilde{P}_j. \tag{S.6}$$

Note that $\tilde{P}_1, \ldots, \tilde{P}_L$ are linearly independent by Lemma B.1. This implies $\sum_{j \geq 2} \nu_j < 1$, because otherwise $\tilde{P}_1 = \sum_{j \geq 2} \nu_j \tilde{P}_j$.

Therefore, any G satisfying (S.6) has the form $\sum_{i=1}^L \gamma_i P_i$. The weights γ_i clearly sum to one, and by joint irreducibility, they are nonnegative. That is, $\boldsymbol{\gamma} := [\gamma_1, \ldots, \gamma_L]^T$ is a discrete distribution. Thus, Eqn. (S.6) is equivalent to

$$\boldsymbol{\pi}_1^T \boldsymbol{P} = (1 - \sum_{j \geq 2} \nu_j) \boldsymbol{\gamma}^T \boldsymbol{P} + \sum_{j \geq 2} \nu_j \boldsymbol{\pi}_j^T \boldsymbol{P}.$$

By linear independence of $P_1, \ldots, P_L$ (see Lemma B.1) and taking the transpose, this gives

$$\boldsymbol{\pi}_1 = (1 - \sum_{j \geq 2} \nu_j) \boldsymbol{\gamma} + \sum_{j \geq 2} \nu_j \boldsymbol{\pi}_j.$$

Therefore $\kappa^*(\tilde{P}_1 | \{\tilde{P}_j, j \neq 1\}) = \kappa^*(\boldsymbol{\pi}_1 | \{\boldsymbol{\pi}_j, j \neq 1\}) < 1$, and there is a one-to-one correspondence between feasible G in the definition of $\kappa^*(\tilde{P}_1 | \{\tilde{P}_j, j \neq 1\})$ and feasible $\boldsymbol{\gamma}$ in the definition of $\kappa^*(\boldsymbol{\pi}_1 | \{\boldsymbol{\pi}_j, j \neq 1\})$. Since Π is recoverable, the residue of $\boldsymbol{\pi}_1$ w.r.t. $\{\boldsymbol{\pi}_j, j \neq 1\}$ is $\boldsymbol{\gamma} = \boldsymbol{e}_1$, and so the residue of $\tilde{P}_1$ w.r.t. $\{\tilde{P}_j, j \neq 1\}$ is $G = \boldsymbol{e}_1^T \boldsymbol{P} = P_1$.

To see uniqueness of the maximizing ν_j, suppose

$$\tilde{P}_1 = (1 - \kappa^*) G + \sum_{j \geq 2} \nu_j \tilde{P}_j = (1 - \kappa^*) G + \sum_{j \geq 2} \nu'_j \tilde{P}_j.$$

Lemma B.1 implies $\nu_j = \nu'_j$.

B.4 Proof of Proposition 4

For brevity we at times omit the dependence of the errors and their estimates on f. For any f,

$$\begin{aligned}
|R_\ell(f) - \widehat{R}_\ell(f)| &= \left| \frac{\tilde{R}_{\ell\ell} - \sum_{j \neq \ell} \nu_{\ell j} \tilde{R}_{j\ell}}{1 - \kappa_\ell} - \frac{\widehat{\tilde{R}}_{\ell\ell} - \sum_{j \neq \ell} \widehat{\nu}_{\ell j} \widehat{\tilde{R}}_{j\ell}}{1 - \widehat{\kappa}_\ell} \right| \\
&\leq \left| \frac{\tilde{R}_{\ell\ell} - \sum_{j \neq \ell} \nu_{\ell j} \tilde{R}_{j\ell}}{1 - \kappa_\ell} - \frac{\widehat{\tilde{R}}_{\ell\ell} - \sum_{j \neq \ell} \widehat{\nu}_{\ell j} \widehat{\tilde{R}}_{j\ell}}{1 - \kappa_\ell} \right| \\
&\qquad + \left| \frac{\widehat{\tilde{R}}_{\ell\ell} - \sum_{j \neq \ell} \widehat{\nu}_{\ell j} \widehat{\tilde{R}}_{j\ell}}{1 - \kappa_\ell} - \frac{\widehat{\tilde{R}}_{\ell\ell} - \sum_{j \neq \ell} \widehat{\nu}_{\ell j} \widehat{\tilde{R}}_{j\ell}}{1 - \widehat{\kappa}_\ell} \right|
\end{aligned}$$

$$
\begin{aligned}
&\leq \frac{|\tilde{R}_{\ell\ell} - \widehat{\tilde{R}}_{\ell\ell}| + \sum_{j\neq\ell} |\nu_{\ell j}\tilde{R}_{j\ell} - \widehat{\nu}_{\ell j}\widehat{\tilde{R}}_{j\ell}|}{1-\kappa_\ell} + \left|\frac{1}{1-\kappa_\ell} - \frac{1}{1-\widehat{\kappa}_\ell}\right| \\
&= \frac{|\tilde{R}_{\ell\ell} - \widehat{\tilde{R}}_{\ell\ell}| + \sum_{j\neq\ell} \left(|\nu_{\ell j}\tilde{R}_{j\ell} - \widehat{\nu}_{\ell j}\tilde{R}_{j\ell} + \widehat{\nu}_{\ell j}\tilde{R}_{j\ell} - \widehat{\nu}_{\ell j}\widehat{\tilde{R}}_{j\ell}|\right)}{1-\kappa_\ell} \\
&\qquad + \left|\frac{1}{1-\kappa_\ell} - \frac{1}{1-\widehat{\kappa}_\ell}\right| \\
&\leq \frac{|\tilde{R}_{\ell\ell} - \widehat{\tilde{R}}_{\ell\ell}| + \sum_{j\neq\ell} \left(|\nu_{\ell j} - \widehat{\nu}_{\ell j}| + |\tilde{R}_{j\ell} - \widehat{\tilde{R}}_{j\ell}|\right)}{1-\kappa_\ell} + \left|\frac{1}{1-\kappa_\ell} - \frac{1}{1-\widehat{\kappa}_\ell}\right|.
\end{aligned}
$$

The VC inequality [1] implies that for any $\epsilon > 0$, $\sup_{f\in\mathcal{F}_{k(\boldsymbol{n})}} |R_{i\ell}(f) - \widehat{R}_{i\ell}(f)| \leq \epsilon$ with probability tending to 1, since (12) holds, and by our convention for multiclass VC dimension. Noting that $\kappa_\ell < 1$ by Proposition 3, the other terms tend to zero in probability by consistency of $\widehat{\kappa}_\ell$ and the $\widehat{\nu}_{\ell j}$. This completes the proof.

B.5 Proof of Theorem 1

Consider the decomposition into estimation and approximation errors,

$$
R(\widehat{f}) - R^* = R(\widehat{f}) - \inf_{f\in\mathcal{F}_{k(\boldsymbol{n})}} R(f) + \inf_{f\in\mathcal{F}_{k(\boldsymbol{n})}} R(f) - R^*.
$$

The approximation error converges to zero by **P3** and since $k(\boldsymbol{n}) \to \infty$. To analyze the estimation error, let $\epsilon > 0$. For each positive integer k, let $f_k^* \in \mathcal{F}_k$ such that $R(f_k^*) \leq \inf_{f\in\mathcal{F}_k} R(f) + \frac{\epsilon}{4}$. Then

$$
\begin{aligned}
R(\widehat{f}) - \inf_{f\in\mathcal{F}_{k(\boldsymbol{n})}} R(f) &= R(\widehat{f}_{k(\boldsymbol{n})}) - \inf_{f\in\mathcal{F}_{k(\boldsymbol{n})}} R(f) \\
&\leq R(\widehat{f}_{k(\boldsymbol{n})}) - R(f^*_{k(\boldsymbol{n})}) + \frac{\epsilon}{4} \\
&\leq \widehat{R}(\widehat{f}_{k(\boldsymbol{n})}) - \widehat{R}(f^*_{k(\boldsymbol{n})}) + \frac{\epsilon}{2} \\
&\quad \text{(with probability tending to 0, by Proposition 4)} \\
&\leq \tau_{k(\boldsymbol{n})} + \frac{\epsilon}{2} \\
&\leq \epsilon
\end{aligned}
$$

where the last step holds for $\boldsymbol{n}$ sufficiently large. The result now follows.

References

[1] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, 1996.