



HAL
open science

SPOCC: Scalable POSSibilistic Classifier Combination - toward robust aggregation of classifiers

Mahmoud Albardan, John Klein, Olivier Colot

► **To cite this version:**

Mahmoud Albardan, John Klein, Olivier Colot. SPOCC: Scalable POSSibilistic Classifier Combination - toward robust aggregation of classifiers. *Expert Systems with Applications*, 2020, 150, pp.113332. 10.1016/j.eswa.2020.113332 . hal-02494735

HAL Id: hal-02494735

<https://hal.science/hal-02494735>

Submitted on 22 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

SPOCC: Scalable POSSibilistic Classifier Combination - toward robust aggregation of classifiers

Mahmoud Albardan, John Klein and Olivier Colot

*Univ. Lille, CNRS, Centrale Lille, UMR 9189 - CRISTAL - Centre de Recherche en
Informatique Signal et Automatique de Lille, F-59000 Lille, France*

Abstract

We investigate a problem in which each member of a group of learners is trained separately to solve the same classification task. Each learner has access to a training dataset (possibly with overlap across learners) but each trained classifier can be evaluated on a validation dataset.

We propose a new approach to aggregate the learner predictions in the possibility theory framework. For each classifier prediction, we build a possibility distribution assessing how likely the classifier prediction is correct using frequentist probabilities estimated on the validation set. The possibility distributions are aggregated using an adaptive t-norm that can accommodate dependency and poor accuracy of the classifier predictions. We prove that the proposed approach possesses a number of desirable classifier combination robustness properties. Moreover, the method is agnostic on the base learners, scales well in the number of aggregated classifiers and is incremental as a new classifier can be appended to the ensemble by building upon previously computed parameters and structures. A python implementation can be downloaded at this link.

Keywords: robust classifier combination, agnostic aggregation, information fusion, classification, possibility theory

1. Introduction

Classification is a supervised machine learning task consisting of assigning objects (inputs) to discrete categories (classes). When several predictors have been trained to solve the same classification task, a second level of algorithmic procedure is necessary to reconcile the classifier predictions and deliver a single one. Such a procedure is known as classifier combination, fusion or aggregation. When each individual classifier is trained using the same training algorithm (but under different circumstances) the aggregation procedure is referred to as an ensemble method. When each classifier may be generated by different training algorithms, the aggregation procedure is referred to as a multiple classifier system. In both cases, the set of individual classifiers is called a classifier ensemble.

Email address: {mahmoud.albardan,john.klein,olivier.colot}@univ-lille1.fr
(Mahmoud Albardan, John Klein and Olivier Colot)

Classifier combination comes either from a choice of the programmer or is imposed by context. In the first case, combination is meant to increase classification performances by either increasing the learning capacity or mitigating overfitting. For instance, boosting (Schapire, 1990) and bagging (Breiman, 1996) can be regarded as such approaches. In the second case, it is not possible to learn a single classifier. A typical situation of this kind occurs when the dataset is dispatched on several machines in a network and sequential learning (such as mini-batch gradient descent) is not possible either to preserve network load or for some privacy or intellectual property reasons. In this decentralized learning setting, a set of classifiers are trained locally and the ensemble is later aggregated by a meta-learner.

In this article, we address classifier aggregation in a perspective that is in line with the decentralized setting assuming that the meta-learner has access to evaluations of the individual classifiers on a fraction of each local dataset which is not used for training. We do not make any assumption on the base learner models and we do not assume base learners are trained on i.i.d. samples, however it is assumed that the union of the fractions of (local) datasets used by the meta-learner is i.i.d.. We introduce a number of desirable robustness properties for the aggregation procedure in this context. We investigate fault tolerance (ability to discard classifiers whose predictions are noise), robustness to adversarial classifiers (ability to thwart classifiers with abnormal error rates) and robustness to redundant information (when classifier predictions are highly dependent).

We introduce an aggregation procedure in the framework of possibility theory. We prove that these robustness properties are verified asymptotically (when the size of the validation set is large) for this new approach. The mechanism governing the aggregation essentially relies on estimates of probabilities of class labels, classifier predictions or classifier correct predictions. There are many related works (Huang & Suen, 1995; Kim & Ghahramani, 2012; Lacoste et al., 2014) dealing with classifier combination using similar information. We believe we are the first to do so in the framework of possibility theory but more importantly these above referenced work are not proved to possess theoretical robustness guarantees. An asymptotic optimality property is verified by an approach from Balakrishnan & Mojirsheibani (2015). This property is stronger than most of the properties that we state except for robustness to classifier dependency. Also, two technical conditions are necessary for the property to hold while our results have no such conditions to check. Similar remarks hold w.r.t. (Biau et al., 2016) which shares some ideas with (Balakrishnan & Mojirsheibani, 2015). Another piece of work with strong properties (oracle inequalities) is exponential weight aggregation (Rigollet & Tsybakov, 2012) but the properties are non-exact¹ and hold in expectation or with high probability while our properties rely on almost sure convergence. Also, exponential weight aggregation is a linear combination model while our method is non-linear.

In addition, the form of aggregation robustness achieved by our method does not jeopardize other important aspects of aggregation such as scalability and incrementalism, two aspects that Balakrishnan & Mojirsheibani (2015) fails

¹Error rate is proved to convergence to the vicinity of the optimal one not exactly to the optimal one.

to possess. Our approach relies on a parametric model involving a number of parameters that is linear in the number of classifiers. Incrementalism is also preserved in the sense that a new classifier can be appended to the ensemble without implying to re-estimate the previously obtained parameter values or structures.

In the next section, we recall the classifier aggregation problem and formally define the robustness properties that we seek. In section 3, we introduce a new aggregation technique in the framework of possibility theory and we show that the desired properties hold asymptotically for this technique. Section 4 contains numerical experiments illustrating our results.

2. Problem statement

2.1. Classification

Let Ω denote a set of ℓ class labels $\Omega = \{\omega_1, \dots, \omega_\ell\}$. Let \mathbf{x} denote an input example with d entries. Most of the time, \mathbf{x} is a vector and lives in \mathbb{R}^d but sometimes some of its entries are categorical data and \mathbf{x} lives in an abstract space which does not necessarily have a vector space structure. Without loss of generality, we suppose that \mathbf{x} is a vector in the rest of this article.

A classification task consists in determining a prediction function c that maps any input \mathbf{x} to its actual class $y \in \Omega$. This function is obtained from a training set $\mathcal{D}_{\text{train}}$ which contains pairs $(\mathbf{x}^{(i)}, y^{(i)})$ where $y^{(i)}$ is the class label of example $\mathbf{x}^{(i)}$. Given K classifiers (each of them trained by one base learner), the label y assigned by the k^{th} classifier to the input \mathbf{x} is denoted by $c_k(\mathbf{x})$.

From a statistical point of view, training examples are instances of a random vector X whose distribution is unknown. Likewise, class labels are instances of a random variable Y whose distribution is also unknown. The training set is often alleged to contain i.i.d. samples of the joint distribution of (X, Y) .

2.2. Classifier performance estimates

The ultimate goal of machine learning is to obtain predictors that generalize well (w.r.t. unseen data at training time). Mathematically speaking, this means achieving the lowest possible expected loss between predictions and true values. When misclassification errors do not have different costs, the 0-1 loss function L is the standard choice:

$$L(y, c_k(\mathbf{x})) = \begin{cases} 0 & \text{if } y = c_k(\mathbf{x}) \\ 1 & \text{otherwise} \end{cases}.$$

In this case, the expected loss is the misclassification error rate of c_k . It is well known, that the error rate minimizer is the Bayes classifier c_\star :

$$c_\star(\mathbf{x}) = \arg \max_{y \in \Omega} p(Y = y | X = \mathbf{x}).$$

Obviously, since the conditional distributions of Y given $X = \mathbf{x}$ are unknown, we must try to find proxys of the Bayes classifier. The error rate (or risk) of classifier c_k is denoted by $r[c_k]$.

Although our goal is to achieve the lowest possible error rate, the performances of a classifier are not, in general, constant across true class labels and

predicted ones. This finer grained information will be instrumental to elicit our possibilistic ensemble of classifiers. This information is contained in the confusion matrix $\mathbf{M}^{(k)}$. Each entry of this matrix reads

$$M_{i,j}^{(k)} = \sum_{(\mathbf{x},y) \in \mathcal{D}_{\text{val}}} \mathbb{I}\{y = \omega_i\} \mathbb{I}\{c_k(\mathbf{x}) = \omega_j\}, \quad (1)$$

where \mathbb{I} denotes the indicator function. It is important to compute the confusion matrices using a validation set \mathcal{D}_{val} disjoint from $\mathcal{D}_{\text{train}}$ otherwise the estimates drawn from the matrix are biased. Actually, if n_{val} is the size of the validation set, then $\frac{M_{ij}^{(k)}}{n_{\text{val}}}$ is the maximum likelihood estimate of the joint probability $p(Y = \omega_i, c_k(X) = \omega_j)$. Also, the sum of the non-diagonal entries of $\mathbf{M}^{(k)}$ over n_{val} is an unbiased estimate of the error rate of c_k . Many other performance criterion estimates can be derived from a confusion matrix.

The classifier combination that we introduce in section 3 essentially relies on the information contained in those matrices. Computing those matrices can thus be regarded as the training phase of the combination method.

2.3. Agnostic combination of classifiers and position of the problem

Let C denote the random vector spanned by plugging X into the ensemble of classifiers:

$$C = \begin{pmatrix} c_1(X) \\ \vdots \\ c_K(X) \end{pmatrix}.$$

A realization of this random vector is denoted by \mathbf{c} or $\mathbf{c}(\mathbf{x})$ whenever the dependence on inputs must be made explicit. We place ourselves in the context where vectors \mathbf{c} can be pictured as new (learned) representation of inputs and we must be agnostic, i.e. we have no control on the base classifiers. In this context, the best aggregate classifier (Balakrishnan & Mojrshuibani, 2015) based on observed data is thus

$$c_*(\mathbf{x}) = \arg \max_{y \in \Omega} p(Y = y | C = \mathbf{c}(\mathbf{x})). \quad (2)$$

Again, the distributions involved in the above definition are unknown. Since \mathbf{c} lives in the discrete space Ω^K , it is possible to try to learn these distributions but this leads to very hard inference problems (Kim & Ghahramani, 2012; Balakrishnan & Mojrshuibani, 2015; Li et al., 2019b) and such statistical learning approaches do not scale well w.r.t. either ℓ or K . The smallest memory complexity among these references is achieved by Li et al. (2019b) who introduce a mixture model relying on tensor decomposition. If K' denotes the number of components in the decomposition, the number of parameters to learn is linear in $K' \times K$. Linearity in K can thus be claimed if $K' \ll K$ which cannot always be assumed. In addition, the Bayesian solutions introduced in these references do not allow to obtain an incremental² aggregation algorithm, an attribute that we believe is much desirable.

²Incremental aggregation means that a new classifier can be appended to the ensemble later without having to recompute everything from scratch.

Generally speaking, classifier combination consists in finding a function f capturing the relation between vectors \mathbf{c} and class labels y that achieves the closest possible performances as compared to c_* . In the approach presented in the next section, we leverage the flexibility of possibility theory to find one such function. The proposed possibilistic approach visits several functions, i.e. aggregation strategies, and select the one maximizing accuracy obtained on the validation dataset. The strategies in question are generalizations of logical rules as opposed to probabilistic approaches which resort to calculus rules. In this regard, the proposed solution both relies on artificial learning and artificial reasoning principles.

Besides, we have chosen to place ourselves in a framework in which each classifier can only deliver one piece of information, i.e. a predicted class label. This allows us to be completely agnostic on the nature of the base learners. Indeed, depending on the training algorithm and model employed by a learner, this latter may be able to deliver a score vector (usually in the form of a probability distribution). In this case, the above analysis no longer applies and the optimal aggregation consists in inferring posterior predictive probabilities of class labels given the scores. Examples of approaches in this line of work are reviewed in Tulyakov et al. (2008). It is possible to remain relatively agnostic on the nature of base classifiers by resorting to a statistical calibration step that allows to obtain prediction probabilities from non-probabilistic classifiers as done in Bella et al. (2013). Calibration will consume a significant portion of the validation set leaving a smaller amount to train the aggregation technique. Consequently, score based aggregation is out of the scope of this paper. Actually, score based aggregation is a leading follow-up of the approach introduced in the next section as mentioned in the concluding remarks in section 5.

2.4. Desirable properties for classifier combination

In terms of purely error rate related performances, the most desirable property for some aggregation function f is

$$r[f(\mathbf{c})] \longrightarrow r[c_*] \text{ as } n_{\text{val}} \longrightarrow \infty. \quad (3)$$

The aggregation technique studied in (Balakrishnan & Mojirsheibani, 2015) achieves a result of this kind (under two technical assumptions). Indeed this technique, which elaborates on (Huang & Suen, 1995), amounts to compute maximum likelihood estimates of the probabilities involved in (2). But classifier aggregation can also bring other types of guarantees which we refer to as robustness. Robustness is understood here as a form of fault tolerance, i.e. the ability to maintain a good level of predictions in several circumstances involving malfunctioning individual classifiers. There may be different causes behind malfunctioning classifiers, e.g. hardware failure or malicious hacks.

Among other possibilities, we have identified the following desirable properties in this scope:

- (a) robustness to random guess: if the error rate of c_k is $\frac{\ell-1}{\ell}$ then $f(\mathbf{c}) = f(\mathbf{c}_{-k})$ where $\mathbf{c}_{-k} \in \Omega^{K-1}$ is the same vector as \mathbf{c} but with its k^{th} entry deleted.

Property (a) means that if the predictions of c_k are in average no better than random guess then c_k has no influence on the aggregated classifier.

- (b) robustness to adversarial classifiers: if c_k has an error rate larger than random guess, i.e. $r[f(\mathbf{c})] > \frac{\ell-1}{\ell}$, then there is a classifier $c_k^{(\text{rec})}$ with an error rate lower than random guess such that $f(\mathbf{c}) = f(\tilde{\mathbf{c}})$ where $\tilde{c}_s = c_s$ for any $s \neq k$ and $\tilde{c}_k = c_k^{(\text{rec})}$.

Property (b) means that we can somewhat rectify the incorrect predictions of classifier c_k so that the aggregated classifier is identical to the one obtained from a non-adversarial situation.

- (c) robustness to redundant information: if there are two individual classifiers such that $c_k(\mathbf{x}) = c_{k'}(\mathbf{x})$ for any \mathbf{x} , then $f(\mathbf{c}) = f(\mathbf{c}_{-k})$.

Property (c) means that copies of classifiers have no influence on the aggregated classifier.

In the above, we assume that the aggregation function f is produced by a given algorithmic procedure and that this procedure applies for any $K > 1$. So $f(\mathbf{c}_{-k})$ is not the restriction of $f(\mathbf{c})$ but another function learned from the same algorithm by omitting classifier c_k .

Obviously, one can think of other properties or reshape them in different ways. For instance, a soft version of property (c) would be better in the sense that an ensemble contains rarely identical copies of a predictor but it contains very often highly dependent ones. This is a first attempt to formalize desirable robustness properties for classifier combination and we hope that more advanced declinations of these will be proposed in the future.

For the time being, our goal in this paper is to introduce an aggregation procedure that is compliant with properties (a) to (c). We will prove that these properties hold for the possibilistic approach that we introduce in the next section at least asymptotically for some of them. Observe that (3) asymptotically implies properties (a) to (c) so the added value of our approach as compared to (Balakrishnan & Mojirsheibani, 2015) relies on its incremental aspect as shown in 3.8 and scalability w.r.t. K as numerical experiments will illustrate in section 4.

2.5. Other related works

So far, we have mentioned only closely related works which perfectly fall in the same setting as ours, i.e. performing the same type of aggregation based on the same information. We remind that this paper is focused on a classifier aggregation paradigm in which one must be agnostic on the base learners. As explained before this immediately rules out a large number of methods such as those relying on classifier scores and ensemble methods. Score based combination most often assume that base learners exploit a given training algorithm. For instance, Loustau (2008) introduced an aggregation method tailored to SVMs while Guo et al. (2019) applied another one to combine deep nets. By definition, so do ensemble methods such as Liu & Zhang (2019). Other score based algorithms will require that training algorithms belong to the same class of models, typically probabilistic learners as in Hoang et al. (2019).

It must also be made clear that the addressed paradigm in this paper is not federated learning. In federated learning, there is no base learners. A group of remotely connected clients have access to a local dataset. Clients are meant to compute a parameter update based on their local data and send

this update to a meta-learner (Yao et al. (2019); Li et al. (2019a); Ji et al. (2019)) to collaboratively train a model. This means that parameter updates are aggregated and this is thus not a classifier aggregation problem.

Finally, an utmost important and original aspect of the method introduced in this paper is its robustness properties w.r.t. noise, adversaries and information redundancy. To the best of our knowledge, there is no prior art in classifier aggregation that has touched jointly these aspects for agnostic aggregation of predicted class labels. Robustness to noise is investigated by Kim & Ghosh (2019) but again in a very different setting which is deep fusion, i.e. deep learning from multiple inputs. Other references in the literature are focused on other aspects than robustness such as security for instance (Ma et al. (2019)). However, robustness is a hot topic in the supervised learning paradigm with important consequences in deep learning (Madry et al., 2018). But obtaining robust base learners does not ensure that the aggregation itself is robust.

3. Robust combination in the possibilistic framework

In this section, we introduce a new classifier combination approach in the possibility theory framework. Possibility theory (Zadeh, 1978; Dubois & Prade, 1988) is an uncertainty representation framework. It has strong connections with belief functions (Dubois & Prade, 1982; Kampé de Fériet, 1982), random sets (Goodman, 1982; Pei-Zhuang & Sanchez, 1982; Sales, 1982), imprecise probabilities (Dubois & Prade, 1992; De Cooman & Aeyels, 1999) or propositional logic (Benferhat et al., 1999). For a concise but thorough overview of possibility theory, the reader is referred to (Dubois & Prade, 2015) but AppendixB already provides deeper insights into this framework and as to why it is particularly relevant for classifier aggregation tasks.

Possibility theory is a widely used framework in symbolic artificial intelligence. It allowed the derivation of new propositional and/or modal logics in which the level of uncertainty of logical propositions can be assessed (Prade (1991); Dubois et al. (2017)). This has applications in logic programming (Alsinet et al. (2008)), automated reasoning (Dubois et al. (1994)) or expert systems (Shenoy (1992)). A popular class of non-probabilistic graphical models relying on ordinal condition functions can also be revisited as a possibilistic model, see Amor et al. (2018) for recent developments in this field. It has also been used in other branches of artificial intelligence such as information fusion (Destercke et al. (2008)) or machine learning (Serrurier & Prade (2015); Hüllermeier (2002)).

In this paper, we adopt a knowledge based system view of this theory. In this regard classifier predictions are expert knowledge to which a degree of belief is attached in the form of possibility distributions. Following a normative approach, experts are reconciled by designing a conjunctive rule that must obey the desirable properties presented in the previous section.

3.1. Possibility theory basics

A possibility distribution π maps each element of Ω to the unit interval $[0; 1]$ whose extreme values correspond respectively to impossible and totally possible epistemic states. If $\pi(y) = 1$ then this class label is totally possible (meaning that we have no evidence against y). If $\pi(y) = 0$ then y is ruled out as a possible class label.

Given a subset A of Ω , a possibility measure Π is given by:

$$\Pi(A) = \max_{y \in A} \pi(y) \quad (4)$$

which means that the possibility of a subset A is equal to the maximum possibility degree in this subset. A possibility measure is thus maxitive: $\Pi(A \cup B) = \max(\Pi(A), \Pi(B))$ as opposed to probability measures which are summative. Observe that this property accounts for the fact that the possibility distribution is enough information to compute the possibility measure of any subset.

3.2. From classifier confusion matrices to possibility distributions

If one normalizes the j^{th} column of the confusion matrix $\mathbf{M}^{(k)}$, then we obtain an estimate of the probability distribution $p(Y = \omega_j | c_k = \omega_j)$. So, if the k^{th} classifier predicts ω_j for some input \mathbf{x} , we can adopt these frequentist probabilities as our beliefs on the class label of \mathbf{x} . But unless, unrealistic conditional independence assumptions³ are made, probabilistic calculus rules will not easily allow to combine beliefs arising from several classifiers.

As an alternative to this approach, we propose to build a possibility distribution from $p(Y = \omega_j | c_k = \omega_j)$ as information fusion in the possibilistic framework can mitigate dependency issues and does not lead to intractable computations. To cast the problem in the possibilistic framework, we use Dubois and Prade transform (DPT) (Dubois & Prade, 1982). For some arbitrary probability distribution p on Ω , let per denote a permutation on Ω such that probability masses of p are sorted in descending order, i.e. if $p' = p \circ \text{per}$ then $p'(\omega_1) \geq p'(\omega_2) \geq \dots \geq p'(\omega_\ell)$. The (unique) possibility distribution π arising from p through DPT is given by

$$\pi(\text{per}(\omega_i)) = \begin{cases} 1 & \text{if } i = 1 \\ \pi(\text{per}(\omega_{i-1})) & \text{if } i > 1 \text{ and } p'(\omega_i) = p'(\omega_{i-1}) \\ \sum_{q=i}^{\ell} p'(\omega_q) & \text{otherwise} \end{cases} \quad (5)$$

If $p_{Y|c_k=\omega_j}$ denotes the distribution of class labels when the k^{th} classifier predicts ω_j , the corresponding possibility distribution is denoted by $\pi_{k|j} = \text{DPT}\{p_{Y|c_k=\omega_j}\}$. For each input \mathbf{x} , the K classifier predictions are turned into K expert opinions in the form of possibility distributions $(\pi_{k|\text{ind}(c_k(\mathbf{x}))})_{k=1}^K$ where $\text{ind}(\omega_j) = j$.

3.3. Aggregation of possibility distributions

Formally speaking, any K -ary operator on the set of possibility distributions is an admissible combination operator. Triangular norms, or t-norms, are instrumental to yield well defined aggregation operators for possibility distributions. A t-norm $\mathcal{T} : [0; 1]^2 \rightarrow [0, 1]$ is a commutative and associative mapping therefore it is easy to build a K -ary version of it using successive pairwise operations:

³These assumptions and the corresponding probabilistic approach are described in section 4.

$$\mathcal{T}(a_1, \dots, a_K) = \mathcal{T}(a_1, \mathcal{T}(\dots, \mathcal{T}(a_{K-1}, a_K))),$$

for any $(a_k)_{k=1}^K \in [0; 1]^K$.

Moreover, a t-norm has 1 as neutral element, 0 as absorbing element and possesses the monotonicity property which reads: for any $a, b, c, d \in [0; 1]$ such that $a \leq c$ and $b \leq d$, then $\mathcal{T}(a, b) \leq \mathcal{T}(c, d)$. Finally, a t-norm is upper bounded by the minimum of its operands.

To combine possibility distributions using a t-norm, we can simply apply a t-norm elementwise. For instance, if $\pi_{kk'}$ is the aggregated possibility distribution obtained by applying a t-norm to distributions π_k and $\pi_{k'}$, then

$$\pi_{kk'}(y) = \mathcal{T}(\pi_k(y), \pi_{k'}(y)), \forall y \in \Omega.$$

We will use the same t-norm symbol to stand for the overall combination of possibility distributions and we will write $\pi_{kk'} = \mathcal{T}(\pi_k, \pi_{k'})$. Examples of t-norms are elementwise multiplication \mathcal{T}_\times and elementwise minimum \mathcal{T}_\wedge .

Decision making based on maximum expected utility is also justified using non-additive measures (capacities) (Gilboa, 1987) such as possibility measures. Consequently, the possibilistic aggregated classifier, denoted c_{ens} , is given by

$$c_{\text{ens}}(\mathbf{x}) = \arg \max_{y \in \Omega} \pi_{\text{ens}}(y), \quad (6)$$

$$\text{with } \pi_{\text{ens}} = \mathcal{T}(\pi_{1|\text{ind}(c_1(\mathbf{x}))}, \dots, \pi_{K|\text{ind}(c_K(\mathbf{x}))}). \quad (7)$$

Algorithm 1 explains what computations should be anticipated as part of a training phase and Algorithm 2 summarizes how an input \mathbf{x} class label is predicted at test time. The procedure corresponding to these algorithms is referred to as Scalable POSSibilistic Classifier Combination (**SPOCC**). Note that there may be several class labels maximizing π_{ens} therefore the aggregated classifier prediction $c_{\text{ens}}(\mathbf{x})$ may be set-valued. Working with set-valued predictions is out of the scope of this paper and will be considered in future works. In the advent of a class label tie, and for any probabilistic, possibilistic or deterministic aggregation approach, one of these labels is chosen at random.

3.4. Adaptive aggregation w.r.t. dependency

The predictions of an ensemble of individual classifiers are usually significantly dependent because they are trained to capture the same bound between inputs and class labels. So if classifiers are at least weak classifiers, they will often produce identical predictions. More importantly, from an information fusion standpoint, if a majority of the classifiers are highly dependent and have a larger error rate than the remaining ones, they are likely to guide the ensemble toward their level of performances making classifier fusion counter-productive.

In the approach introduced in this paper, it is possible to mitigate dependency negative impact by choosing an idempotent t-norm such as elementwise minimum \mathcal{T}_\wedge . Indeed, in the worst dependency case, classifier c_k is a copy of classifier $c_{k'}$ therefore they have an unjustified weight in the ensemble predictions. But if two individual classifiers are identical they will also yield identical possibility distributions and if these latter are combined using \mathcal{T}_\wedge , then these two classifiers will be counted as one. This is exactly the spirit of property (c).

Two difficulties arise from this quest for robustness w.r.t. classifier redundancy:

Data: validation set \mathcal{D}_{val} , classifiers $(c_k)_{k=1}^K$, number of class labels ℓ .

for $k \in \{1, \dots, K\}$ **do**

 Compute confusion matrix $\mathbf{M}^{(k)}$ as in (1).

for $j \in \Omega$ **do**

 Compute conditional probability estimates

$$\hat{p}_{\text{mle}}(Y = \omega_i | c_k = \omega_j) \leftarrow \frac{M_{ij}^{(k)}}{\sum_{i'} M_{i'j}^{(k)}}, \forall i \in \Omega.$$

 Compute possibility distribution using (5)

$$\pi_{k|j} \leftarrow \text{DPT} \{ \hat{p}_{\text{mle}}(Y = \cdot | c_k = \omega_j) \}.$$

end

end

Return possibility distributions $(\pi_{k|j})_{\substack{1 \leq k \leq K \\ 1 \leq j \leq \ell}}$.

Algorithm 1: SPOCC - training phase

Data: input \mathbf{x} , classifiers $(c_k)_{k=1}^K$, possibility distributions $(\pi_{k|j})_{\substack{1 \leq k \leq K \\ 1 \leq j \leq \ell}}$ and t-norm \mathcal{T} .

for $k \in \{1, \dots, K\}$ **do**

 | Compute individual classifier prediction $j_k \leftarrow \text{ind}(c_k(\mathbf{x}))$.

end

Compute $\pi_{\text{ens}} \leftarrow \mathcal{T}(\pi_{1|j_1}, \dots, \pi_{K|j_K})$.

Return $c_{\text{ens}}(\mathbf{x}) \leftarrow \arg \max_{y \in \Omega} \pi_{\text{ens}}(y)$.

Algorithm 2: SPOCC - test phase

- (i) It is not recommended to systematically use an idempotent combination mechanism because it is also possible that two poorly dependent classifiers yield identical possibility distributions in which case it appears justified that their common prediction impacts on the ensemble aggregated decision.
- (ii) There are different levels of dependency among subsets of individual classifiers therefore, using a single t-norm to jointly aggregate them is not the best option.

To address the first issue, we propose to use the following parametric family $(\mathcal{T}_\lambda)_{\lambda \in [1; +\infty)}$ of t-norms:

$$\mathcal{T}_\lambda(a_1, a_2) = e^{-(|\log a_1|^\lambda + |\log a_2|^\lambda)^{\frac{1}{\lambda}}}, \forall a_1, a_2 \in [0; 1]. \quad (8)$$

This family is known as Aczel-Alsina t-norms and is such that $\mathcal{T}_1 = \mathcal{T}_\times$ and $\mathcal{T}_\infty = \mathcal{T}_\wedge$. We can thus tune λ all the higher as the level of dependence between classifiers is high.

To assess the dependence level λ among two classifiers c_k and $c_{k'}$, we use the following definition

$$\kappa(c_k, c_{k'}) = 1 - \exp\left(-\frac{1}{n_{\text{val}}} \left| \log\left(\frac{\mathcal{L}_0}{\mathcal{L}_1}\right) \right|\right), \quad (9)$$

where $\frac{\mathcal{L}_0}{\mathcal{L}_1}$ is the likelihood ratio of the independence model over the joint model. These likelihoods are given by

$$\mathcal{L}_0 = \prod_{i=1}^{n_{\text{val}}} \hat{p}_{\text{mle}}(c_k(\mathbf{x}^{(i)})) \hat{p}_{\text{mle}}(c_{k'}(\mathbf{x}^{(i)})), \quad (10)$$

$$\text{and } \mathcal{L}_1 = \prod_{i=1}^{n_{\text{val}}} \hat{p}_{\text{mle}}(c_k(\mathbf{x}^{(i)}), c_{k'}(\mathbf{x}^{(i)})). \quad (11)$$

These likelihoods are computed using all training examples contained in the validation set \mathcal{D}_{val} . The probabilities involved in the computation of \mathcal{L}_0 are the maximum likelihood estimates of the parameters of the multinomial marginal distributions $p(c_k(X))$ and $p(c_{k'}(X))$ respectively. The probabilities involved in the computation of \mathcal{L}_1 are the maximum likelihood estimates of the parameters of the multinomial joint distribution $p(c_k(X), c_{k'}(X))$.

The definition of the dependence level κ can be extended to more than two classifiers but this will turn out to be unnecessary because we will use hierarchical agglomerative clustering (Ward Jr, 1963) (HAC) to address issue (ii). HAC will produce a dendrogram \mathcal{G} , i.e. a t-norm computation binary tree. Each leaf in this tree is in bijective correspondence with one of the possibility distributions π_k induced by a classifier. There are thus K leaves in \mathcal{G} . Furthermore, each non-leaf node in the tree stands for a t-norm operation involving two operands only. Consequently, each non-leaf node has exactly two children and there $K - 1$ such nodes, one of them being the root node. Figure 1 gives an illustrative example of a dependence dendrogram allowing to compute the aggregated possibility distribution.

HAC relies on a classifier dissimilarity matrix \mathbf{D} . In our case, entries of this $K \times K$ matrix are simply given by $D_{kk'} = 1 - \kappa(c_k, c_{k'})$.




Figure 1: Example of a dendrogram for $K = 5$. Leaf nodes are at the bottom. For each of the four non-leaf nodes, a specific dependence level λ_a ($a \in \{1; 2; 3; 4\}$) must be determined to compute the aggregated possibility distribution.

The t-norm based possibility distribution aggregation method described in the above paragraphs is meant to replace the penultimate step of Algorithm 2 but most of the computations pertaining to this dependency adaptive aggregation can be done at training time. Indeed, the computation of the pairwise dependence levels and the dendrogram do not depend on the unseen example \mathbf{x} that we will try to classify at test time. For a minimal test phase computation time, we need to assign to each non-leaf node V_a of \mathcal{G} the appropriate dependence level λ_a (as illustrated in Figure 1) during the training phase. The corresponding array is denoted by $\boldsymbol{\lambda}[a] \mapsto \lambda_a$. The function $J_{\mathcal{G}; \boldsymbol{\lambda}}$ maps the set of possibility distributions $(\pi_k)_{1 \leq k \leq K}$ to the aggregated distribution π_{ens} by executing the computation graph \mathcal{G} and using the dependence levels contained in $\boldsymbol{\lambda}$. There are $K - 1$ hyperparameters in array $\boldsymbol{\lambda}$ that need to be tuned. They will be automatically set to appropriate values by heuristic search, see Appendix A for a presentation of this grid search based heuristic. It is noteworthy that this heuristic will use HAC to define clusters of classifiers, thereby reducing the number of hyperparameters to tune as this amounts to merge some nodes of the dendrogram and apply some t-norm to more than two possibility distributions.

3.5. Adaptive aggregation w.r.t. informational content

When the predictions delivered by classifier c_k are poorer than those of another classifier $c_{k'}$, it is instrumental to reduce the impact of c_k on the decisions issued by the ensemble. Regardless of the formal definition behind what are called "poor predictions", we propose to use the following mechanism to gradually fade classifier c_k out of the ensemble: for a given scalar $\alpha_k \in [0; 1]$, we update all conditional possibility distributions related to c_k as follows:

$$\pi_{k|j} \leftarrow (1 - \alpha_k) \pi_{k|j} + \alpha_k, \forall j \in \Omega. \quad (12)$$

This mechanism is equivalent to an operation known as discounting (Shafer, 1976). When $\alpha = 0$, then classifier c_k influence on the ensemble is not reduced. When $\alpha_k = 1$, classifier c_k is discarded from the ensemble since we obtain constant one possibility distributions which are the neutral element of t-norms and the t-norm based aggregation method introduced in the previous subsection.

Obviously, we need to find a value of the discounting coefficient tailored for each classifier and in line with what poor predictions are meant to be. Again, it is tempting to set these K hyperparameters using grid search but the corresponding complexity calls for a more subtle strategy. Similarly as for dependency hyperparameters, we will resort to a heuristic search.

Among other possibilities, our solution consists in binding the discounting rates together using the following formula:

$$\alpha_k = 1 - \left(\frac{1 - \hat{r}[c_k]}{1 - \min_{k'} \hat{r}[c_{k'}]} \right)^\rho, \quad (13)$$

where \hat{r} is the estimated error rate on the validation set and $\rho \in [0; +\infty]$ is a hyperparameter to tune by grid search. Using the above equation, the best base classifier is not discounted and we have $\hat{r}[c_k] \leq \hat{r}[c_{k'}] \Rightarrow \alpha_k \leq \alpha_{k'}$.

3.6. Fully adaptive aggregation

The fully adaptive version (w.r.t. both dependence and informative content) of SPOCC is referred to as **adaSPOCC**. The corresponding training and test phases are described in Algorithm 3 and 4 respectively. A python implementation can be downloaded at this link.

Data: validation set \mathcal{D}_{val} , classifiers $(c_k)_{k=1}^K$, number of class labels ℓ .

Execute SPOCC - training phase (algorithm 1)

```

for  $k \in \{1, \dots, K\}$  do
  | for  $k' \in \{k, \dots, K\}$  do
  | | Compute the dissimilarity  $D_{kk'} \leftarrow 1 - \kappa(c_k, c_{k'})$  using (9).
  | | Assign  $D_{k'k} \leftarrow D_{kk'}$ .
  | end
end

```

end

Obtain dendrogram \mathcal{G} by applying HAC to dissimilarity matrix \mathbf{D} .

Apply heuristic to set parameters $\lambda_a \in \boldsymbol{\lambda}$ (see AppendixA).

Compute parameters $(\alpha_k)_{k=1}^K$ as in (13).

Update all conditional possibility distributions as in (12).

Return possibility distributions $(\pi_{k|j})_{\substack{1 \leq k \leq K \\ 1 \leq j \leq \ell}}$, dendrogram \mathcal{G} , array $\boldsymbol{\lambda}$.

Algorithm 3: adaSPOCC - training phase

Data: input \mathbf{x} , classifiers $(c_k)_{k=1}^K$, possibility distributions $(\pi_{k|j})_{\substack{1 \leq k \leq K \\ 1 \leq j \leq \ell}}$, dendrogram \mathcal{G} , array $\boldsymbol{\lambda}$.

```

for  $k \in \{1, \dots, K\}$  do
  | Compute individual classifier prediction  $j_k \leftarrow \text{ind}(c_k(\mathbf{x}))$ .
end

```

$\pi_{\text{ens}} \leftarrow J_{\mathcal{G}; \boldsymbol{\lambda}}(\pi_{1|j_1}, \dots, \pi_{K|j_K})$ (computation graph execution).

Return $c_{\text{ens}}(\mathbf{x}) \leftarrow \arg \max_{y \in \Omega} \pi_{\text{ens}}(y)$.

Algorithm 4: adaSPOCC - test phase

3.7. Properties of the possibilistic ensemble

In this paper, we adopt a normative approach for the selection of a classifier decision aggregation mechanism. In this subsection, we give sketches of proofs showing that robustness properties (a) to (c) hold for adaSPOCC asymptotically:

- Property (a): if c_k is a random classifier then when $n_{\text{val}} \rightarrow \infty$, each conditional distribution $p_{Y|c_k=\omega_j}$ converges to a uniform distribution so DPT turns it into a constant one possibility distribution, which is the neutral element of \mathcal{T}_λ .
- Property (b): let c_k denote an adversarial classifier, i.e. $r[c_k] > \frac{\ell-1}{\ell}$. (ada)SPOCC uses the following rectified classifier $c_k^{(\text{rec})} = h \circ c_k$ defined as

$$c_k^{(\text{rec})}(\mathbf{x}) = \arg \max_{y \in \Omega} p(Y = y | c_k(\mathbf{x})). \quad (14)$$

We have

$$1 - r[c_k^{(\text{rec})}] = \sum_{y' \in \Omega} p(Y = y' | c_k^{(\text{rec})} = y') p(c_k^{(\text{rec})} = y'). \quad (15)$$

Moreover, we can write

$$p(Y = y' | c_k^{(\text{rec})} = y') = \sum_{y'' \in \Omega} p(Y = y' | c_k^{(\text{rec})} = y', c_k = y'') p(c_k = y'' | c_k^{(\text{rec})} = y'). \quad (16)$$

Given the definition of $c_k^{(\text{rec})}$ we know that $p(c_k = y'' | c_k^{(\text{rec})} = y') = 0$ if $y'' \notin h^{-1}(y')$. The definition also gives

$$p(Y = y' | c_k^{(\text{rec})} = y', c_k = y'') = \max_{y \in \Omega} p(Y = y | c_k = y''). \quad (17)$$

The maximal probability value of a discrete variable is always greater or equal than $\frac{1}{\ell}$ therefore

$$p(Y = y' | c_k^{(\text{rec})} = y') \geq \frac{1}{\ell} \sum_{y'' \in h^{-1}(y')} p(c_k = y'' | c_k^{(\text{rec})} = y') \quad (18)$$

$$\geq \frac{1}{\ell} p(c_k \in h^{-1}(y') | c_k^{(\text{rec})} = y'). \quad (19)$$

Again, given the definition of $c_k^{(\text{rec})}$ we know that $p(c_k \in h^{-1}(y') | c_k^{(\text{rec})} = y') = 1$. Since c_k is not the random classifier, at least one of the conditional distributions $p_{Y|c_k}$ is not uniform in which case the inequality is strict. We thus obtain $1 - r[c_k^{(\text{rec})}] > \frac{1}{\ell}$.

Finally, when $n_{\text{val}} \rightarrow \infty$, if $c_k(\mathbf{x}) = y$ and $c_k^{(\text{rec})}(\mathbf{x}) = y'$, the y^{th} column of $\mathbf{M}^{(k)}$ will be identical to the y'^{th} column of the confusion matrix of $c_k^{(\text{rec})}$ so they will be mapped to identical possibility distributions.

- Property (c): when $n_{\text{val}} \rightarrow \infty$, the likelihood ratio appearing in (9) writes

$$\frac{\mathcal{L}_0}{\mathcal{L}_1} = \prod_{i=1}^{n_{\text{val}}} \frac{p(c_k(\mathbf{x}^{(i)})) p(c_{k'}(\mathbf{x}^{(i)}))}{p(c_k(\mathbf{x}^{(i)}, c_{k'}(\mathbf{x}^{(i)}))}. \quad (20)$$

If $c_{k'}$ is a copy of c_k then $p(c_k(\mathbf{x}^{(i)}, c_{k'}(\mathbf{x}^{(i)})) = p(c_k(\mathbf{x}^{(i)})) = p(c_{k'}(\mathbf{x}^{(i)}))$ and

$$\frac{\mathcal{L}_0}{\mathcal{L}_1} = \prod_{i=1}^{n_{\text{val}}} p(c_k(\mathbf{x}^{(i)})). \quad (21)$$

If c_k is not a constant function, then probabilities are smaller than one and $\frac{\mathcal{L}_0}{\mathcal{L}_1} \rightarrow 0$. The pair of classifiers $(c_k, c_{k'})$ will thus be detected as maximally dependent by HAC and they will be aggregated using $\mathcal{T}_1 = \mathcal{T}_\wedge$ hence property (c) holds in this case.

When c_k is a constant function, then both c_k and $c_{k'}$ will yield identical possibility distributions that are a Dirac function. In this case, the output of \mathcal{T}_\wedge will also be this Dirac function, meaning that idempotence is always true in these circumstances. Note that, however, the procedure described in 3.4 will fail to detect the dependency between c_k and $c_{k'}$. There are plenty of ways to thwart this issue as constant classifiers are not difficult to detect. In practice, we will use add-one Laplace smoothing to estimate probabilities $p(c_k)$ so we will never obtain a Dirac function as possibility distribution.

Properties (a) to (c) rely on asymptotic estimates of multinomial distribution parameters which, from the strong law of numbers, converge almost surely to their exact values. Consequently, the properties do not hold only in expectation or with high probability but systematically (when n_{val} is large).

Although properties (a) to (c) are not as strong as (3), adaSPOCC is a scalable aggregation technique as the number of parameters it requires to learn from the validation set is in $O(\ell^2 K)$ while the number of parameters to learn from \mathcal{D}_{val} in (Balakrishnan & Mojrshuibani, 2015) is in $O(\ell^{K+1})$ and is therefore doomed to overfit when K is large.

3.8. Incremental aggregation

When a new classifier c_{K+1} must be appended to the ensemble, it suffices to compute its corresponding confusion matrix $\mathbf{M}^{(K+1)}$ to be able to run SPOCC. All previously estimated parameters (confusion matrices and possibility distributions of the other classifiers) can be readily re-used.

Going incremental for adaSPOCC is not as straightforward as for SPOCC. A new coefficient α_{K+1} needs to be computed but this latter is deduced from $\mathbf{M}^{(K+1)}$ so this is not an issue. However, the matrix \mathbf{D} also needs to be updated by appending a new line and a new column to it, which makes $K - 1$ new entries to compute because \mathbf{D} is symmetric and its diagonal elements are irrelevant. Then, HAC must be re-run. To increase the level of incrementalism of adaSPOCC in this regard, it is possible to use an incremental clustering algorithm such as Menon et al. (2019). The newly coming classifier will be either appended to an existing cluster or a new cluster that solely contains c_{K+1} will be created. In the first case, the hyperparameters $\boldsymbol{\lambda}$ can be left unchanged. In

the second case, there is an additional node in the dendrogram and one additional hyperparameter must be estimated by grid search. Since we perform grid search for only one such parameter, this is obviously faster than the heuristic search described by Algorithm 5. In conclusion, adaSPOCC is also an incremental aggregation algorithm and all previously estimated parameters are also re-used without needing to be updated.

4. Experimental results

In this section, we present a number of experimental results allowing to prove the robustness of SPOCC and adaSPOCC as compared to other aggregation techniques. The section starts with results obtained when the base classifiers are trained on a synthetic dataset and are meant to highlight performance discrepancies in simple situations where robustness is required. Another set of experiments on real datasets are also presented to prove that the method is not only meaningful on toy examples.

4.1. General setup

Designing numerical experiments allowing to compare aggregation methods is not a trivial task. A crucial aspect consists in training a set of base classifiers that achieve a form of diversity (Wozniak et al., 2014) so that the fusion of their predictions has a significant impact on performances. Among other possibilities (Brown et al., 2005), we chose to induce diversity by feeding the base classifiers with different disjoint subsets of data points at training time. The subsets are not chosen at random but instead in a deterministic way allowing each base classifier to focus on some regions of the input space and thus learn significantly different decision frontiers.

Because the union of the validation sets is an i.i.d. sample of $p(X, Y)$ and aggregation methods have access to the predictions of each learner on this set union, well designed aggregation methods are able to restore high levels of performances even if base learners are trained from non-i.i.d. samples. This allows us to test if aggregation methods are relatively agnostic to the quality of the data used to train the base learners.

Each aggregation technique is fed only with the predictions of the base classifiers on the validation set in order to tune hyperparameters or learn the combination itself. Consequently, SPOCC and adaSPOCC are only compared to well established methods that use the same level of information. The benchmarked aggregation techniques are :

- classifier selection⁴ based on estimated accuracies of the base classifiers,
- weighted vote aggregation based on estimated accuracies of the base classifiers,
- exponentially weighted vote aggregation based on estimated accuracies of the base classifiers,
- naive Bayes aggregation,

⁴Selection can be regarded as a special type of fusion.

- Bayes aggregation,
- stacking.

In the exponentially weighted vote aggregation, accuracies are not directly used as vote weight (as in standard weighted vote aggregation) but are mapped to weights using a softmax function. This function has a positive temperature hyperparameter that regulates the assignment of weights. When this parameter is zero, then we retrieve unweighted vote aggregation whereas when it is very large, then we retrieve classifier selection.

Bayes aggregation relies on (2). The conditional distributions involved in this equation are learned from the validation set. Naive Bayes aggregation uses conditional independence assumptions that allow to factorize probabilities as

$$p(y|\mathbf{c}(\mathbf{x})) \propto p(y) \prod_{k=1}^K p(c_k(\mathbf{x})|y). \quad (22)$$

The conditional independence assumptions are not realistic but yield a model with far less parameters to learn as compared to Bayes aggregation.

For each of the estimated probabilities involved in the mechanism behind SPOCC, adaSPOCC, Bayes or naive Bayes aggregation, we perform add-one-Laplace smoothing to avoid computational issues related to zero probabilities. The chosen t-norm for SPOCC is $\mathcal{T}_5 \approx \mathcal{T}_\wedge$.

Finally, we also train a softmax regression to map classifier predictions to the true class labels. This approach belongs to a methodology known as stacking (Wolpert, 1992). An L_2 regularization term is added to the cross-entropy loss. A positive hyperparameter regulates the relative importance of the regularization term.

All hyperparameters are tuned automatically using a cross-validated grid search on the validation set. For each hyperparameter, the grid contains 100 points. When the hyperparameter is unbounded, a logarithmic scale is used to design the grid.

The statistical significance of the reported results are given in terms of 95% confidence intervals estimated from bootstrap sampling. When the accuracies of two aggregation methods have overlapping confidence intervals, the performance discrepancy is not significant. A companion python implementation of adaSPOCC and benchmark methods can be downloaded at [this link](#)

4.2. Synthetic Data

In this subsection, we use a very simple generating process to obtain example/label pairs. Data points are sampled from four isotropic Gaussian distributions. The centers of these Gaussian distributions are located at each corner of a centered square in a 2D input space ($d = 2$). The standard deviations of each of the distribution is 1. There are $\ell = 2$ possible class labels: $\Omega = \{\omega_1; \omega_2\}$. Points such that x_1 and x_2 are both positive belong to ω_1 . Points such that x_1 and x_2 are both negative also belong to ω_1 . All the other points belong to ω_2 . Figure 2 shows one such dataset obtained from this generating process.

In this series of experiments, the dataset has $n = 200$ points and is divided in four overlapping subsets as depicted in Figure 3. Then a randomly selected




Figure 2: Synthetic dataset obtained from four Gaussian distributions. Examples belonging to class ω_1 are in blue while those belonging to class ω_2 are in cyan. Optimal decision frontiers are in magenta.




Figure 3: Subsets of the data seen respectively by c_1 to c_4

portion of 80% of each such subset is used to train one of the base classifier. The remaining 20% are used for the validation set. Each base classifier c_k sends the corresponding set of prediction/label pairs $\{(c_k(\mathbf{x}), y)\}_{(\mathbf{x}, y) \in \mathcal{D}_{\text{val}}}$ to the aggregation method.

Since we have access to the data generating process, the test set is dynamically created. We sample test points until the observed accuracies of all the tested methods are with probability 0.95 in their respective Clopper-Pearson confidence intervals of half-width 0.2%. The whole procedure is repeated 100 times. The averaged test errors in this case are thus a good approximation of the generalization errors of the tested methods.

Given the shape of optimal decision frontiers, the base classifiers trained in this subsection are decision trees with a maximal depth of two.

4.2.1. Robustness w.r.t. adversaries

Among other possibilities, adversarial predictions are simulated by sampling from a Bernoulli distribution $Z \sim \text{Ber}(\theta)$. Given $Z = 1$, the prediction of a base classifier is replaced with another (arbitrarily selected) class label that will not coincide with the classifier prediction. When $Z = 0$, the classifier prediction is unchanged. Consequently, an adversarial classifier built in this way from a base classifier with an error rate lower than random guess will achieve an error rate greater than random guess as $\theta \rightarrow 1$.

The evolution of the classification accuracy of the benchmarked aggregation methods as the number of adversaries grows can be witnessed on Figure 4.




Figure 4: Evolution of accuracy distributions (violin plots) for several aggregation methods w.r.t. the number of adversaries. SPOCC and adaSPOCC are in orange while other methods are in blue.

For simplicity, all adversaries are built from the same base classifier (c_1) with $\theta = 0.5$. Two methods cannot maintain the same level of performances when the number of adversaries increases: weighted vote ensemble and Bayes aggregation. For the weighted vote ensemble, this is explained by the fact that the number of misleading classifiers outnumber legitimate classifiers and start to obtain a majority of votes. For Bayes aggregation, the performances are degrading simply because of overfitting. Indeed, Bayes aggregation has a number of parameter to learn that is exponential in K while SPOCC and other methods have a number of parameters at most linear in K .

4.2.2. Robustness w.r.t. faults

Erroneous predictions are simulated by sampling from a Bernoulli distribution $Z \sim \text{Ber}(\theta)$. Given $Z = 1$, the prediction of a base classifier is replaced with an (arbitrarily selected) class label that will coincide with the classifier prediction with probability $\frac{1}{\ell}$. When $Z = 0$, the classifier prediction is unchanged. Consequently, a noisy classifier built in this way from a base classifier will achieve an error rate equal to $\frac{\ell-1}{\ell}$ (random guess) as $\theta \rightarrow 1$.

The evolution of the classification accuracy of the benchmarked aggregation methods as the number of noisy classifiers grows can be witnessed on Figure 5. For simplicity, all noisy classifiers are built from the same base classifier (c_1) with $\theta = 0.9$. Similarly as for adversarial classifiers, weighted vote ensemble




Figure 5: Evolution of accuracy distributions (violin plots) for several aggregation methods w.r.t. the number of noisy classifiers. SPOCC and adaSPOCC are in orange while other methods are in blue.

and Bayes aggregation cannot maintain the same level of performances when the number of perturbed classifiers increases. The same reasons are also behind these performance decays (majority of incorrect classifiers for the weighted vote ensemble and overfitting for Bayes aggregation).

4.2.3. Robustness w.r.t. informational redundancy

Redundancy in classifier predictions is simulated by adding several copies of one of the base classifiers (classifier c_1 in our experiments). As shown in Figure 6, this very simple setting allows to observe severe performance decays for the weighted vote ensemble, the exponentially weighted vote ensemble and the naive Bayes aggregation. Vote based ensembles are very sensitive to changes of majority. Naive Bayes aggregation is also sensitive to this phenomenon and suffers from its inability to capture dependency relations between the base classifiers.

Unlike the previous experiment, it can be noted that Bayes aggregation maintains the same level of performances as the number of clones of c_1 increases. Because clones will always produce identical predictions as c_1 , training the Bayes aggregation in these conditions is equivalent to learn from $K = 4$ base classifiers regardless how many copies of c_1 are added. However, should these copies be slightly perturbed, then we would observe the same overfitting issues as in the previous experiments.

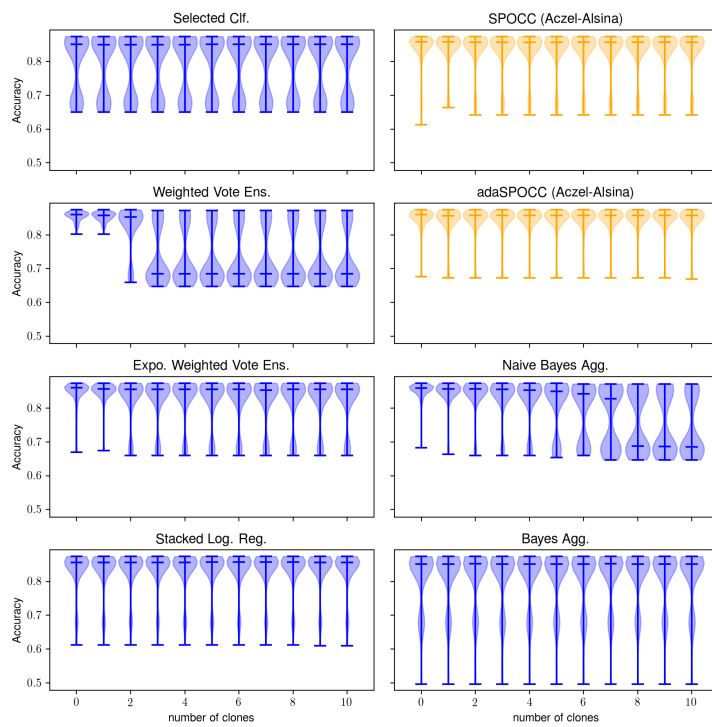


Figure 6: Evolution of accuracy distributions (violin plots) for several aggregation methods w.r.t. the number of copies of c_1 . SPOCC and adaSPOCC are in orange while other methods are in blue.

Table 1: Average performances of aggregation methods on the synthetic data. The first figure is the average accuracy followed by the semi-width of the 95% confidence interval width of this latter and the average standard deviation. Best accuracies (or those not statistically significantly different) are in bold characters.

Method	Adversaries	Faults	Redundancy	Global Average
Clf. Selection	79.26% ±0.52 std. 8.56%	79.25% ±0.50 std. 8.56%	79.25% ±0.51 std. 8.55%	79.25% ±0.28 std. 8.56%
Weighted Vote	81.47% ±0.25 std. 4.23%	81.89% ±0.23 std. 3.83%	76.46% ±0.52 std. 8.86%	79.94% ±0.22 std. 6.57%
Exp. Weighted Vote	84.32% ±0.23 std. 3.86%	84.35% ±0.23 std. 3.77%	82.75% ±0.39 std. 6.46%	83.81% ±0.17 std. 4.92%
Stacking	83.43% ±0.28 std. 4.75%	83.50% ±0.29 std. 4.82%	83.38% ±0.36 std. 5.89%	83.44% ±0.17 std. 5.18%
Naive Bayes Agg.	85.23% ±0.13 std. 2.25%	85.22% ±0.14 std. 2.25%	80.27% ±0.46 std. 8.12%	83.57% ±0.19 std. 5.55%
Bayes Agg.	66.20% ±0.77 std. 13.45%	66.36% ±0.79 std. 13.43%	81.16% ±0.51 std. 8.40%	71.24% ±0.46 std. 13.90%
SPOCC	84.42% ±0.23 std. 3.82%	84.42% ±0.23 std. 3.82%	83.52% ±0.34 std. 5.76%	84.10% ±0.16 std. 4.61%
adaSPOCC	84.01% ±0.23 std. 3.89%	84.01% ±0.23 std. 3.89%	84.54% ±0.52 std. 4.12%	84.16% ±0.13 std. 4.00%
Best base Clf.	82.65% ±0.41 std. 6.87%	82.65% ±0.41 std. 6.87%	82.65% ±0.41 std. 6.87%	82.65% ±0.41 std. 6.87%
Optimal Clf.	87.52% ± ≈ 0 std. ≈ 0%	87.52% ±0.52 std. ≈ 0%	87.52% ± ≈ 0 std. ≈ 0%	87.52% ± ≈ 0 std. ≈ 0%

4.2.4. Summarizing synthetic data experiment robustness results

In the previous paragraphs, we have seen which methods are tolerant to adversaries, faults and redundancy and scale well w.r.t. K . Only SPOCC, adaSPOCC and stacking seem to be robust w.r.t. each of these forms of difficulties. Beside robustness, their absolute performances also matter. Average⁵ performances are reported in Table 1 for each experiment as well as the global average on all experiments. We also provide as reference the optimal (Bayes) classifier accuracy as well as the performances of the best base classifier c_k , i.e. optimal selection.

In terms of accuracies, SPOCC or adaSPOCC are always the top 1 or top 2 aggregation approach. While the naive Bayes aggregation is slightly better than SPOCC or adaSPOCC in the two first series of experiments, it performs very significantly worse in the last one and it is outperformed on global average. Stacking and all other methods obtain worse (or sometimes comparable) results as compared to (ada)SPOCC. Moreover, observe that adaSPOCC achieves the smallest variance, meaning that its performances are more stable across dataset draws. Normalized confusion matrices corresponding to the global average on all experiments (last column of Table 1) are shown in Figure 7. It shows the distribution of error rates in terms of type I and type II errors. Let alone classifier selection (which achieves anyway poor general performances), adaSPOCC is the aggregation method with the smallest type I / type II error discrepancy.

⁵Averages w.r.t either the number of adversaries, noisy classifiers or clones respectively.

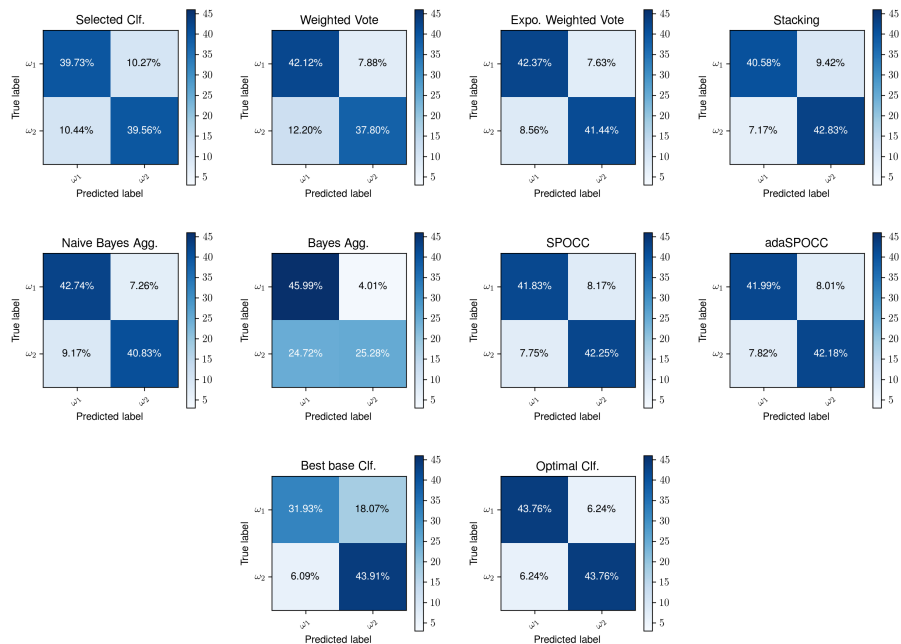


Figure 7: Normalized averaged confusion matrices for the reported results in the last column of Table 1.

4.2.5. Other experimental aspects

The main goal of this experimental section is to illustrate the robustness properties (a) to (c) that adaSPOOC possesses. The results reported in the previous paragraphs match this purpose. However, other aspects are also interesting to examine. In the following paragraphs, we investigate the behavior of the tested aggregation methods under two different circumstances:

- when the base classifiers are heterogeneous,
- when the dataset is imbalanced, meaning that $p(Y)$ is not uniform.

Heterogeneous ensemble of classifiers. An important advantage of the class label agnostic aggregation setting over others is that no assumption at all are made on the base classifiers and therefore any training algorithm can be used to derive them. To illustrate this ability, we reproduce the same experiment as in 4.2.3 with different base classifiers. Now, c_1 is trained using logistic regression, c_2 is a 5-nearest neighbor classifier, c_3 is an SVM with radial basis function as kernel while c_4 is a decision tree like before. Classifier c_1 will achieve smaller accuracy because a linear decision function underfits the data in this case. Since this classifier will be progressively duplicated, it will also check the ability of methods to cope with increasingly many lesser accurate classifiers.

The corresponding results are displayed in Figure 8. Since at least one member of the ensemble performs more poorly than in 4.2.3, all methods have their accuracy distributions eroded. However, the conclusions from the previous experiments are confirmed as the same methods achieve robustness to information

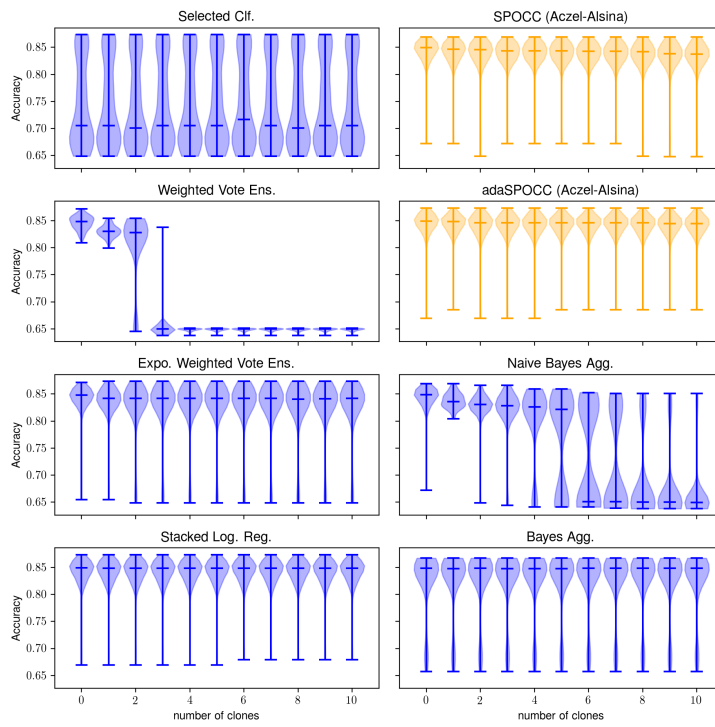


Figure 8: Evolution of accuracy distributions (violin plots) on an heterogeneous ensemble of 4 classifiers for several aggregation methods w.r.t. the number of copies of c_1 . SPOCC and adaSPOCC are in orange while other methods are in blue.

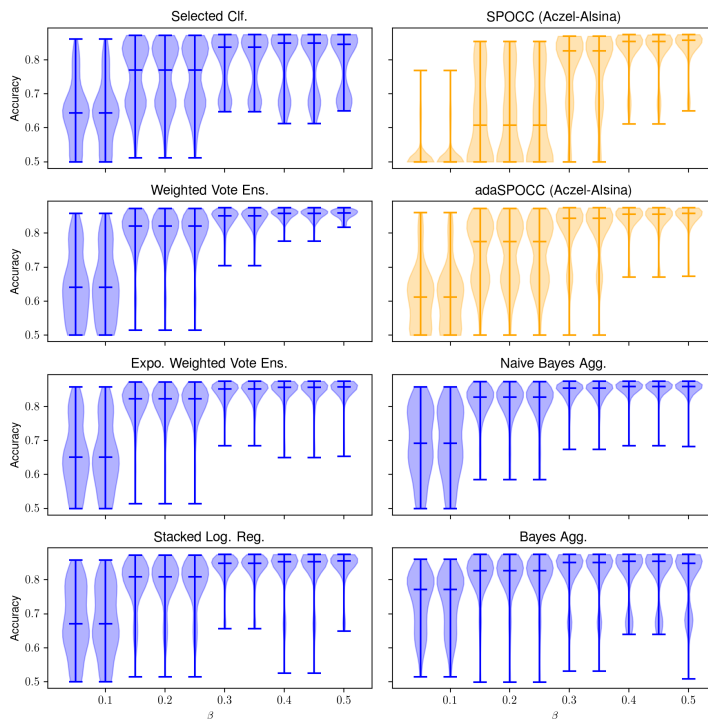


Figure 9: Evolution of accuracy distributions (violin plots) with class label imbalance for several aggregation methods w.r.t. β (first class label probability). SPOCC and adaSPOCC are in orange while other methods are in blue.

redundancy and corruption. It is also made clear that the ability of the aggregation to overcome these difficulties does not lie with the training algorithms employed to obtain the base classifiers.

Class label imbalance. In practice, it is common that the generative process underlying our data is such that the class label probability distribution is not uniform. In the previous set of experiments, such an imbalance was in place for the base learner but not for the aggregation methods. We now modify the generative process such $p(Y = \omega_1) = \beta$ and $p(Y = \omega_2) = 1 - \beta$. The level of imbalance is progressively reduced as β increases. The set of examined values is: $\beta \in \{0.05; 0.1; 0.15; \dots; 0.5\}$.

The corresponding results are displayed in Figure 9. For any value of β , this setting is very favorable to majority based methods. Indeed, if the decision tree training went alright, we should obtain approximately the following predictors

$$c_1(\mathbf{x}) = \begin{cases} \omega_2 & \text{if } x_1 < 0 \text{ and } x_2 > 0 \\ \omega_1 & \text{otherwise} \end{cases} \quad c_2(\mathbf{x}) = \begin{cases} \omega_1 & \text{if } x_1 > 0 \text{ and } x_2 > 0 \\ \omega_2 & \text{otherwise} \end{cases}$$

$$c_3(\mathbf{x}) = \begin{cases} \omega_2 & \text{if } x_1 > 0 \text{ and } x_2 < 0 \\ \omega_1 & \text{otherwise} \end{cases} \quad c_4(\mathbf{x}) = \begin{cases} \omega_1 & \text{if } x_1 < 0 \text{ and } x_2 < 0 \\ \omega_2 & \text{otherwise} \end{cases} .$$

We see that, for any \mathbf{x} , there are 3 classifiers out of 4 that yield correct predictions therefore majority voting based aggregations are expected to perform very well. In addition, we also have $r[c_i] \approx r[c_j]$ for any i and j so the weighted vote and exponentially weighted vote aggregation are nearly equivalent to majority voting. As we can see, these two methods perform very well when $\beta = 0.5$ and can tolerate imbalance up to $\beta = 0.15$. Below this value, there are two few points in the class ω_1 ⁶ for the decision tree to learn meaningful prediction rules and thus their aggregation (regardless of the method) is not meaningful either. Stacking can also easily learn a combination rule that mimics majority vote. It thus compares favorably to vote based methods.

The naive Bayes aggregation also works very well in this setting because $p(c_1|Y = \omega_1) \approx \mathbb{I}\{c_1 = \omega_1\}$, $p(c_2|Y = \omega_2) \approx \mathbb{I}\{c_2 = \omega_2\}$, $p(c_3|Y = \omega_1) \approx \mathbb{I}\{c_3 = \omega_1\}$ and $p(c_4|Y = \omega_2) \approx \mathbb{I}\{c_4 = \omega_2\}$. Consequently, the naive Bayes aggregation will easily rule out the classifier yielding an incorrect prediction. We can see that this method achieves comparable performances as compared to the vote based ones.

The Bayes aggregation is a gold standard because it infers the optimal decision rule as explained in 2.3. However it still has $\ell(\ell^K - 1) = 30$ parameters to learn from 40 data points in the validation set and will thus slightly overfit. It thus achieves worse accuracies than naive Bayes or vote based aggregation but seems to better handle extreme imbalance.

Because indicator functions (or Dirac masses) are fixed points of DPT and zero is the absorbing element of t-norms, SPOCC and adaSPOCC can enjoy the same type of information as the naive Bayes aggregation does. They indeed perform well when $\beta > 0.35$. However, they exhibit higher sensitivity to imbalance than other methods. As often, adaSPOCC appears to be more robust than SPOCC. We believe this is due to the fact that they are discriminative aggregation models in the sense that they do not rely on the whole data distribution but solely operate on the conditional distributions $p(Y|c_k)$. As the experiments on real data will show, adaSPOCC already works pretty well on several imbalanced datasets, however possible fixes for this limitation will be investigated in future works and are discussed in section 5.

4.3. Real Data

To upraise the ability of the benchmarked methods to be deployed in more realistic situations (such as decentralized learning), we also need to test them on sets of real data. Since this is essentially useful in a big data context, we chose eight from moderate to large public datasets. The specifications of these datasets are reported in Table 2.

Example entries from the 20newsgroup data set are word counts obtained using the term frequency - inverse document frequency statistics. We reduced the dimensionality of inputs using a latent semantic analysis (Deerwester et al., 1990) which is a standard practice for text data. We kept 100 dimensions. Also, as recommended, we stripped out each text from headers, footers and quotes which lead to overfitting. Besides, for the Wine and Avila datasets, the number of class labels is originally 10 and 12 respectively. We binarized these

⁶When $\beta = 0.1$, classifiers c_2 and c_4 have access (in expectation) to 8 data points in class ω_1 and only four such points when $\beta = 0.05$.

Table 2: Real dataset specifications

Name	Size n	Dim. d	Nbr. of classes ℓ	Data type	Class imbalance	Source
20newsgroup	18846	100 (after red.)	20	text	yes	sklearn
MNIST	70000	784	10	image	no	sklearn
Satellite	6435	36	6	image features	yes	UCI repo. (Statlog)
Wine	6497	11	2 (binarized)	chemical features	yes	UCI repo. (Wine Quality)
Spam	4601	57	2	text	yes	UCI repo. (Spam)
Avila	10430	10	2 (binarized)	layout features	yes	UCI repo. (Avila)
Drive	58509	48	11	current statistics	no	UCI repo. (Sensorless Drive Diagnosis)
Particle	130064	50	2	signal	yes	UCI repo. (MiniBooNE particle identification)

classification tasks because some classes have very small cardinalities which is problematic for our experimental design in which datasets are divided into several distinct subsets. Indeed, some subsets may possess no example at all of some classes which leads to imprecise labeling which is beyond the scope of this paper. To circumvent these acute class imbalance issues, classes were merged as follows:

- In the Wine data set, class labels are wine quality scores. Two classes are obtained by comparing scores to a threshold of 5.
- In the Avila dataset, class labels are middle age bible copyist identities. The five first copyists are grouped in one class and the remaining ones in the other class.

Unlike synthetic data sets, we need to separate the original dataset into a train set and a test set. To avoid a dependency of the reported performances w.r.t train/test splits, we perform 2-fold cross validation (CV). Also, we shuffled at random examples and repeated the training and test phases 100 times.

To induce diversity in the base classifiers, we separated the training data into 6 distinct pieces using the following procedure: for each data set, for each class,

1. apply principal component analysis to the corresponding data,
2. project this data on the dimension with highest eigenvalue,
3. sort the projected values and split them into 6 subsets of cardinality $n_i/6$ where n_i is the proportion of examples belonging to class ω_i .

We argue that this way of splitting data leads to challenging fusion tasks because some base classifiers may see data that are a lot easier to separate than it should and will consequently not generalize very well. Actually, the training data to which classifier c_k has access is a non-i.i.d. sample of the distribution of (X, Y) .

We used softmax regression with an L_2 regularization term to train the base classifiers. The regularization hyperparameter is set to default (i.e. 1.0).

To make sure that robustness observations from the previous subsection are confirmed on real data, we also add two noisy classifiers to the ensemble. Both noisy classifiers are built from c_1 with $\theta = 0.01$ therefore they are perturbed copies of c_1 .

Table 3: Classification accuracies (with bootstrap confidence intervals and standard deviations) for several real data sets when 6 base classifiers were trained separately on disjoint subsets of the datasets. Datasets were split in a non-i.i.d. way using a PCA based protocol. Two slightly noisy copies of c_1 were added so there are $K = 8$ base classifiers to aggregate.

Method	20newsgroup	MNIST	Satellite	Wine	Spam	Avila	Drive	Particle
Clf. Selection	40.86% ±0.12% std. 0.83%	70.86% ±0.08% std. 0.79%	75.23% ±0.15% std. 1.54%	64.92% ±0.26% std. 2.60%	87.31% ±0.18% std. 1.83%	60.56% ±0.34% std. 3.48%	51.55% ±0.31% std. 3.18%	82.11% ±0.15% std. 1.56%
Weighted Vote	38.67% ±0.15% std. 1.59%	67.48% ±0.06% std. 0.61%	75.34% ±0.08% std. 0.85%	65.60% ±0.16% std. 1.62%	85.22% ±0.20% std. 2.05%	60.77% ±0.35% std. 3.60%	50.72% ±0.33% std. 3.38%	79.35% ±0.24% std. 2.50%
Exp. Weighted Vote	42.17% ±0.08% std. 0.83%	75.10% ±0.06% std. 0.61%	77.88% ±0.10% std. 1.01%	65.63% ±0.18% std. 1.81%	88.50% ±0.14% std. 1.42%	62.30% ±0.35% std. 3.60%	57.01% ±0.46% std. 4.67%	82.64% ±0.12% std. 1.29%
Stacking	19.01% ±0.15% std. 1.58%	37.82% ±0.24% std. 2.43%	61.72% ±0.10% std. 1.04%	64.65% ±0.26% std. 2.76%	89.67% ±0.13% std. 1.28%	65.58% ±0.12% std. 1.24%	37.10% ±0.28% std. 2.78%	83.75% ±0.11% std. 1.12%
Naive Bayes Agg.	39.32% ±0.43% std. 4.43%	76.11% ±0.17% std. 1.69%	78.06% ±0.09% std. 0.87%	64.21% ±0.21% std. 2.16%	87.83% ±0.23% std. 2.34%	60.76% ±0.33% std. 3.42%	66.39% ±0.60% std. 6.05%	80.10% ±0.30% std. 2.90%
Bayes Agg.	Intract.	Intract.	66.07% ±0.17% std. 1.65%	63.89% ±0.27% std. 2.77%	88.90% ±0.14% std. 1.42%	65.39% ±0.12% std. 1.26%	Intract.	83.29% ±0.11% std. 1.12%
SPOCC	35.99% ±0.19% std. 1.93%	78.00% ±0.22% std. 2.29%	77.55% ±0.10% std. 1.00%	63.60% ±0.10% std. 1.07%	86.74% ±0.34% std. 3.32%	61.91% ±0.33% 3.27%	66.86% ±0.53% std. 5.48%	73.31% ±0.23% std. 2.35%
adaSPOCC	41.19% ±0.10% std. 1.03%	79.13% ±0.26% std. 2.66%	78.59% ±0.08% std. 0.77%	64.92% ±0.25% std. 2.51%	89.26% ±0.14% std. 1.40%	63.33% ±0.32% 3.21%	67.75% ±0.47% std. 4.79%	82.13% ±0.14% std. 1.47%
Best base Clf.	42.11% ±0.06% std. 0.66%	70.89% ±0.06% std. 0.66%	75.82% ±0.08% std. 1.25%	65.66% ±0.16% std. 1.42%	87.73% ±0.09% std. 1.27%	62.59% ±0.11% std. 1.78%	52.06% ±0.24% std. 2.63%	82.33% ±0.12% std. 1.65%
Centralized Clf.	57.43% ±0.04% std. 0.39%	91.44% ±0.02% std. 0.12%	82.43% ±0.04% std. 0.44%	73.72% ±0.05% std. 0.53%	92.18% ±0.05% std. 0.50%	68.23% ±0.04% std. 0.47%	74.72% ±0.03% std. 0.35%	88.56% ±0.25% std. 2.58%

Table 4: Maximal accuracy discrepancy w.r.t. the best approach. Max is taken over the 8 datasets.

Method	Clf. Selection	Weighted Vote	Exp. Weighted Vote	Stacking	Naive Bayes Agg.	Bayes Agg.	SPOCC	adaSPOCC
max. discr.	16.2%	17.03%	10.74%	41.31%	4.82%	12.52%	10.44%	2.25%

Average accuracies over random shuffles and CV-folds are given in Table 3 for $K = 8$ base classifiers. Train/validation split ratio are identical to the synthetic dataset case. A number of observations can be made based on these results:

- Classifier selection based on estimated accuracies is always significantly outperformed by some of the aggregation techniques which shows that the experimental protocol meets its purpose (providing a setting allowing to do better than base classifiers). Even "oracle" classifier selection (reported as best base classifier in Table 3) is outperformed in 6 datasets out of 8 and achieves comparable performance in the remaining two.
- adaSPOCC always obtain better results than SPOCC which indicates that it is safer to analyze classifier dependencies as well as estimated individual performances on real data.

- Memory occupation became problematic for Bayes aggregation whenever $\ell > 6$ as $\ell(\ell^K - 1)$ parameters need to be estimated and stored. It achieves unsurprisingly poor performances when $\ell = 6$ (Satellite dataset) confirming its inability to scale w.r.t. ℓ or K .
- adaSPOCC is one of the most efficient aggregation approach. It achieves the highest average rank (over the 8 datasets). adaSPOCC has average rank of 2.1 followed by the exponentially weighted vote which is in average the top 3 approach.
- adaSPOCC is robust in the sense that it achieves the minimal maximal discrepancy w.r.t. the best concurrent approach. Absolute values of maximal discrepancies (over the 8 datasets) w.r.t. the best approach are reported in Table 4.

The corresponding normalized confusion matrices are given in AppendixC. These matrices exhibit different patterns, thereby showing that the methods converge to significantly different aggregation strategies. Also, interestingly, the methods do not necessarily issue incorrect predictions under the same circumstances and a second stage of aggregation may take advantage of their respective strengths.

5. Conclusion

In this article, a new classifier aggregation technique is introduced. This technique relies on the framework of possibility theory. Conditional probabilities of class labels given a classifier prediction are estimated on a validation set and transformed in possibility distributions. For each input to be classified, the set of possibility distributions issued by the classifier predictions are regarded as a set of propositions that are conjunctively combined using a t-norm. The obtained method, called SPOCC, is scalable w.r.t. to both the number of class labels and the number of base classifiers. It is also incremental as extra-classifiers can be appended later to the ensemble without re-computing previously derived parameters.

An adaptive version of this method, called adaSPOCC is also introduced. It is proposed to perform hierarchical agglomerative clustering to identify subsets of classifiers which are not statistically independent. Each such cluster can thus be combined sequentially with different t-norms. T-norms are chosen from the Aczel-Alsina parametric family which allows to reduce the impact of redundant predictions when necessary. Moreover, the individual impact of a base classifier can also be regulated by setting discounting coefficients. When one such coefficient is set to one, the corresponding classifier is discarded from the fusion process. These coefficients as well as the t-norm parameters are automatically tuned using heuristic search monitoring the ensemble accuracy on the validation set.

The adaptive version of this non-probabilistic aggregation method possesses a number of nice statistical properties. These properties are well supported by several numerical experiments and clearly show its ability to tolerate adversaries, faults or information redundancy. In a series of experiments on a synthetic dataset and involving variable numbers of (respectively) noisy, adversarial or redundant classifiers, adaSPOCC achieves the highest average accuracy by a

comfortable margin (≈ 2.5 times the 95% confidence interval width). In a second series of experiments on 8 real datasets and involving several (jointly) noisy, adversarial and redundant classifiers, adaSPOCC achieves the highest average accuracy rank. Thanks to its robustness, its worst performance is limited to a 2.25% misclassification error overhead as compared to the best method. This is the minimal such overhead among all other methods, the second smallest such overhead being 4.82% which is more than two times larger.

There are several future research tracks that we plan to investigate to further develop this contribution. One of them consists in investigating to what extent the approach is modular w.r.t. imprecise classifiers, i.e. classifiers that can only discriminate between subsets of class labels. Since possibility theory is natively compatible with set theory, this modularity seems not too challenging to achieve as opposed to many other concurrent approaches.

A second one consists in designing a more efficient way to set parameters λ than grid search (and the heuristic search presented in Appendix A). To replace grid search, we plan to try to fit π_{ens} to the one-hot encoding of the true label. It is already known that L_k norm based metrics are relevant distance for possibility distributions as illustrated in Jousselme & Maupin (2012). A projected gradient descent would allow us to find a relevant estimate of λ .

Another interesting question is to adapt the proposed classifier aggregation method to regression tasks. This will require to use possibilities on the real line for instance which may be computationally more demanding.

Another line of work consists in proposing an alternative version of adaSPOCC in which possibility distributions are built from classifier output scores instead of just their class label predictions. Classifier score are much more informative than class label predictions because they provide a ranked list of candidate solutions equipped with levels of confidence. On the downside, they require either not be agnostic on base learners or to perform a calibration step. Mapping scores to possibilities can be regarded as a regression task for which a dedicated portion of the validation set can be used for fitting.

Finally, while adaSPOCC can tolerate a certain level of class imbalance, its sensitivity to this phenomenon could be reduced in several ways. One can think of turning the class probability distribution into a possibility distribution and combine it to the aggregated possibility distribution issued by adaSPOCC. It is also possible to give different weights to the data points in the validation set and optimize the t-norm with respect to a corresponding weighted empirical risk instead of the unweighted version that we used in this article.

Appendix A. Heuristic search for dependency parameters

In this appendix, we explain how to set hyperparameters $(\lambda_a)_{a=1}^{K-1}$ which are necessary to execute the computation graph \mathcal{G} as part of adaSPOCC. Each hyperparameter λ_a regulates the level of dependency between operands aggregated using the Aczel-Alsina t-norm \mathcal{T}_{λ_a} . Each λ_a lives in $[1; +\infty]$ and we can use a logarithmic grid and the validation set to assess the impact of a given value of λ_a in terms of classification accuracy of the ensemble. However, resorting to baseline grid search has exponential complexity in K and we will thus employ a heuristic search to keep computation time at bay.

To cleverly browse possible values for $(\lambda_a)_{a=1}^{K-1}$, we can remark that HAC agglomerates classifiers from most dependent ones to least dependent ones. This

implies that if V_a is a child node of $V_{a'}$ in \mathcal{G} then $\lambda_a \geq \lambda_{a'}$. Consequently, instead of systematically visit all configurations, we will start by performing a grid search on the full grid for the lowest nodes in the hierarchy and freeze the corresponding hyperparameters. Then, we will move to their parent nodes and perform grid search only for smaller values. This sequential grid search has a maximal complexity equal to $K - 1$ times the cost of a 1D grid search.

Another trick allowing to improve the procedure consists in jointly setting a subset of hyperparameters. Clusters can be obtained from \mathcal{G} . This is actually the original intent behind HAC. For a given number of clusters N_c , clusters are obtained by thresholding cophenetic correlation coefficients between pairs of classifiers. More precisely, if a pair of classifiers have a cophenetic correlation distance below the threshold, they are considered to belong to different clusters. Starting with a sufficiently high value of the threshold so that all classifiers are in one unique cluster, the threshold is lowered until the constraint on N_c is violated, i.e. further lowering it yields $N_c + 1$ clusters. Note that the obtained clusters always correspond to non-overlapping branches of \mathcal{G} . This strategy is wrapped up by iterating on N_c , starting from $N_c = 2$ to K .

Algorithm 5 summarizes the proposed heuristic search for dependency parameters. The 1D grid search for one λ_a is performed on a predefined grid and the retained value is the one achieving highest accuracy of the classifier ensemble.

```

Initialize  $\lambda_a \leftarrow 1, \forall V_a \in \mathcal{G}$ .
for  $N_c$  from 2 to  $K$  do
    Obtain  $N_c$  clusters denoted by  $\mathcal{C}_1 \dots \mathcal{C}_{N_c}$ .
    for  $i$  from 1 to  $N_c$  do
        Find  $V_a$  s.t. its descendants contain  $\mathcal{C}_i$  and no other leaf node.
        Perform grid search jointly for  $\lambda_a$  and all  $\lambda_{a'}$  in correspondence
            with non-leaf descendants of  $V_a$ .
        Append node  $V_a$  and all its descendants to the set Treated.
    end
    Obtain list Remain of those nodes in  $\mathcal{G}$  which do not belong to Treated.
    Sort list Remain so that Remain[ $i$ ] cannot be an antecedent of
        Remain[ $j$ ] if  $i < j$ .
    for  $i$  from 1 to length of Remain do
        Perform grid search for  $\lambda_a$  where  $V_a = \text{Remain}[i]$  and subject to
             $\lambda_a \leq \lambda_{a'}$  for any  $V_{a'}$  that is a descendant of  $V_a$ .
    end
    Save in  $r_i$  the empirical error rate achieved with this value of  $\lambda$ .
    if  $N_c \geq 3$  and  $r_i \geq r_{i-1}$  then
        Revert to the value of  $\lambda$  obtained at the previous iteration.
        Stop looping.
    end
end
Return  $\lambda$ .

```

Algorithm 5: Heuristic search for dependency hyperparameters

Appendix B. Deeper insights into possibility theory

Possibility theory was introduced by Zadeh (1978) and further developed by Dubois & Prade (1988) with the motivation to offer a well-defined and formal mathematical representation for linguistic statements that permits handling imprecise or vague information. For instance, the word *cheap* can be given a large set of values according to everyone's subjective definition and context of cheapness. Possibility values can be interpreted as *degrees of feasibility* of event occurrence. An important difference with the way probabilities encode uncertainty is that high possibility values are non-informative while high probability values are. Indeed, a very high possibility for event A means that, should A occur or not, we would not be surprised. If A has a very high probability mass, then we would be surprised that A does not occur. Conversely, low possibility and probability values are both informative as they both indicated that an occurrence of A is unlikely.

Generally speaking, possibilities are a special class of imprecise probabilities which is a setting in which probability values can only be bracketed by two bounds. Indeed, the uncertainty of an event in the possibilistic framework is better upraised by a pair of values (possibility and necessity) which can be seen as probability bounds. In this regard, possibility theory is not at all an orthogonal theory to probabilities but provide a convenient language to grasp some aspects of uncertainty that lead to highly complex models in probabilistic language.

Appendix B.1. Main concepts

Possibility theory can be formalized as follows. Suppose events $B \subseteq \Omega$ can be sorted into one of the following families:

- \mathcal{T} the family of true propositions,
- \mathcal{F} the family of false propositions,
- and \mathcal{U} the family of undecided propositions.

Define two functions N and Π that respectively represent the certainty of truth and the possibility of truth. More formally,

$$N(B) = \mathbb{I}\{B \in \mathcal{T}\}, \quad (\text{B.1})$$

$$\Pi(B) = \mathbb{I}\{B \in \mathcal{T} \cup \mathcal{U}\}, \quad (\text{B.2})$$

where \mathbb{I} is the indicator function. Take two events B_1 and B_2 . If they are both certain then so is $B_1 \cap B_2$ and if at least of one of them is not certain then $B_1 \cap B_2$ is not certain either. We thus have

$$N(B_1 \cap B_2) = \min \{N(B_1); N(B_2)\}. \quad (\text{B.3})$$

Conversely, if at least either B_1 or B_2 is possible then so is $B_1 \cup B_2$, hence

$$\Pi(B_1 \cup B_2) = \max \{\Pi(B_1); \Pi(B_2)\}. \quad (\text{B.4})$$

Moreover, if B is impossible then B^c is surely true, hence

$$\Pi(B) = 1 - N(B^c). \quad (\text{B.5})$$

Now, by preserving the above logical rules but allowing N and Π to take values in $[0; 1]$ and to represent graded membership of families \mathcal{T} and $\mathcal{T} \cup \mathcal{U}$, then we obtain the framework of possibility theory.

In this framework, the function N is called necessity measure and the function Π is called possibility measure and it is immediate from (B.5) that they encode the same information. From (B.4), it is also obvious that the possibility measure of any subset B can be obtained from the possibility measures of singletons. The restriction of Π to singletons is denoted π and called possibility distribution. A possibility distribution has the same memory complexity as a probability distribution.

Example 1. Suppose $\Omega = \{\omega_1; \omega_2; \omega_3\}$. Suppose classifier c_1 is a 4-nearest-neighbor and class labels ω_1 and ω_2 receive two votes when c_1 is asked to predict the class label y of some input \mathbf{x} . Consequently, aggregation methods receive the following message from c_1 : $B_1 = \{\omega_1; \omega_2\} \in \mathcal{T}$. This piece of information is encoded in the following possibility distribution:

$$\begin{array}{cccc} y & \omega_1 & \omega_2 & \omega_3 \\ \pi & 1 & 1 & 0 \end{array}$$

The corresponding possibility and necessity measures are given by

$$\begin{array}{ccccccc} y & \{\omega_1\} & \{\omega_2\} & \{\omega_1; \omega_2\} & \{\omega_3\} & \{\omega_1; \omega_3\} & \{\omega_2; \omega_3\} & \Omega \\ N & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ \Pi & 1 & 1 & 1 & 0 & 1 & 1 & 1 \end{array}$$

The above example illustrates the ability of possibility distributions to easily encode set theoretic information. The same information can be encoded by an infinite set of probability distributions, i.e. the set of those distributions such that $p(\{\omega_1; \omega_2\}) = 1$. Observe that the probability distribution $p_0(\omega_1) = p(\omega_2) = \frac{1}{2}$ belongs to those distributions in question. The distribution p_0 does not only inform us that y is either ω_1 or ω_2 but also that there are evidence that these outcomes are equally probable. It is thus more informative than the statement $B_1 \in \mathcal{T}$.

Note that SPOCC does not exploit this aspect of possibility theory but this is mentioned as one of the perspectives for future work. In this paper, SPOCC leverages the flexibility of possibility theory when it comes to aggregate propositions in the form of possibility distributions.

Appendix B.2. Probability / Possibility transforms

In this article, the information issued by classifier c_k that we would like to encode as a possibility distribution corresponds to one column of matrix $\mathbf{M}^{(k)}$. After dividing by n_{val} these entries of $\mathbf{M}^{(k)}$, we obtain a maximum likelihood estimate of the conditional distribution $p(y|c_k(\mathbf{x}))$. It is therefore necessary to turn the information encoded by this probability distribution into a possibility distribution in order to exploit possibilistic aggregation mechanisms.

To turn a probability distribution p into a possibility one π , three desirable properties stated by Dubois et al. (2004) are

- *consistency*: $\forall A \subseteq \Omega, \Pi(A) \geq p(A)$ where Π is the possibility measure spanned by π . So Π is a probability upper bound.

- *preference preservation*: $\forall(\omega, \omega') \in \Omega^2, p(\omega) > p(\omega') \Leftrightarrow \pi(\omega) > \pi(\omega')$, so there is a form of compatibility between the preferences encoded by π and those encoded by p .
- *maximal specificity*: π is the most informative possibility distribution among those possibility distributions consistent and preserving preferences with p . Considering two possibility distribution $\pi^{(1)}$ and $\pi^{(2)}$, the possibility distribution $\pi^{(1)}$ is said to be more informative than $\pi^{(2)}$, denoted $\pi^{(1)} \preceq \pi^{(2)}$, if $\forall\omega, \pi^{(1)}(\omega) \leq \pi^{(2)}(\omega)$.

The only transformation verifying the above properties is DPT which is presented in 3.2.

Example 2. Suppose $\Omega = \{\omega_1; \omega_2; \omega_3\}$. Below are two examples of application of DPT to two conditional probability distributions:

y	ω_1	ω_1	ω_3	y	ω_1	ω_1	ω_3
$p(y c_1(\mathbf{x}) = y_3)$	0.1	0.15	0.75	$p(y c_2(\mathbf{x}) = y_2)$	0.2	0.6	0.2
$\pi_{1 3}$	0.1	0.25	1	$\pi_{2 2}$	0.4	1	0.4

Appendix B.3. Conjunctive aggregation of possibility distributions

In classical logic, the logical conjunction operator (logical AND) is true when all operands are true. When the operands are subsets, it thus amounts to taking the intersection of these subsets. Logical conjunction allows us to combine possibility distributions that are basically indicator functions of a set (as in example 1). When possibility values are in $]0; 1[$ (as in example 2), we need to define an aggregation operator that is a generalization of logical conjunction.

The parametric operator \mathcal{T}_λ presented in 3.3 and 3.4 that relies on Aczel-Alsina t-norms is one such generalization in the sense that if $\pi_1(\omega) = \mathbb{I}\{\omega \in B\}$ and $\pi_2(\omega) = \mathbb{I}\{\omega \in C\}$ then, $\mathcal{T}_\lambda(\pi_1, \pi_2)(\omega) = \mathbb{I}\{\omega \in B \cap C\}$.

Actually, the operator \mathcal{T}_λ achieves a more general form of conjunctivity which reads

$$\mathcal{T}_\lambda(\pi_1, \pi_2) \preceq \pi_1 \text{ and } \pi_2, \forall \pi_1, \pi_2. \quad (\text{B.6})$$

This means that the result of the aggregation of possibility distributions using \mathcal{T}_λ is always at least as informative than the most informative input distribution. This property comes from the fact that for any distributions π_1 and π_2 and any $\lambda < \lambda'$, we have

$$\mathcal{T}_\lambda(\pi_1, \pi_2) \preceq \mathcal{T}_{\lambda'}(\pi_1, \pi_2).$$

Remembering that $\mathcal{T}_\infty(\pi_1, \pi_2)$ is the entrywise minimum of π_1 and π_2 , it is obvious that $\mathcal{T}_\infty(\pi_1, \pi_2) \preceq \pi_1$ and π_2 and by transitivity (B.6) holds.

Example 3. Suppose $\Omega = \{\omega_1; \omega_2; \omega_3\}$. Below is a figure displaying the aggregated possibility distribution π_{ens} using \mathcal{T}_λ when one operand is distribution $\pi_{1|3}$ from example 2, a second operand is distribution $\pi_{2|2}$ from example 2 and the final and third operand is $\pi_{3|2} = \pi_{2|2}$.




Figure B.10: Evolution of possibility values of π_{ens} w.r.t. to λ

The above example illustrates that depending on the value of λ , the class label achieving maximal possibility is either a class label strongly supported by one classifier or moderately supported by a majority of classifiers. Since λ is tuned by a grid search based heuristic in adaSPOCC, the chosen aggregation policy is data driven and will maximize accuracy.

The key aspect of possibility theory that we exploit in the derivation of adaSPOCC is that the above described aggregation operator is flexible and has a number of parameters to learn that is linear in K the number of classifiers in the ensemble. In probability theory, aggregating $p(y|c_1)$ and $p(y|c_2)$ consists in inferring the distribution $p(y|c_1, c_2)$. As explained earlier and illustrated by the numerical experiments, direct inference leads to overfitting or is simply intractable. When simplifying conditional independence assumptions are made, the computations become tractable at the expense of accuracy.

In possibility theory, the philosophy behind aggregation operators such as \mathcal{T}_λ is fundamentally different. Aggregation operators are not derived through calculus rules but by logical rules. Desirable properties are stated and an operator satisfying them is derived (normative approach). This is a knowledge based system view, in which possibility distributions encode partial knowledge that must be combined in a principled way.

Appendix B.4. Decision making from possibilities


It is clear that based on the aggregated possibility distribution π_{ens} obtained from the ensemble of classifiers using Algorithm 2 or 4, the intuition calls for taking the argmax of this possibility distribution in order to produce a reliable prediction.

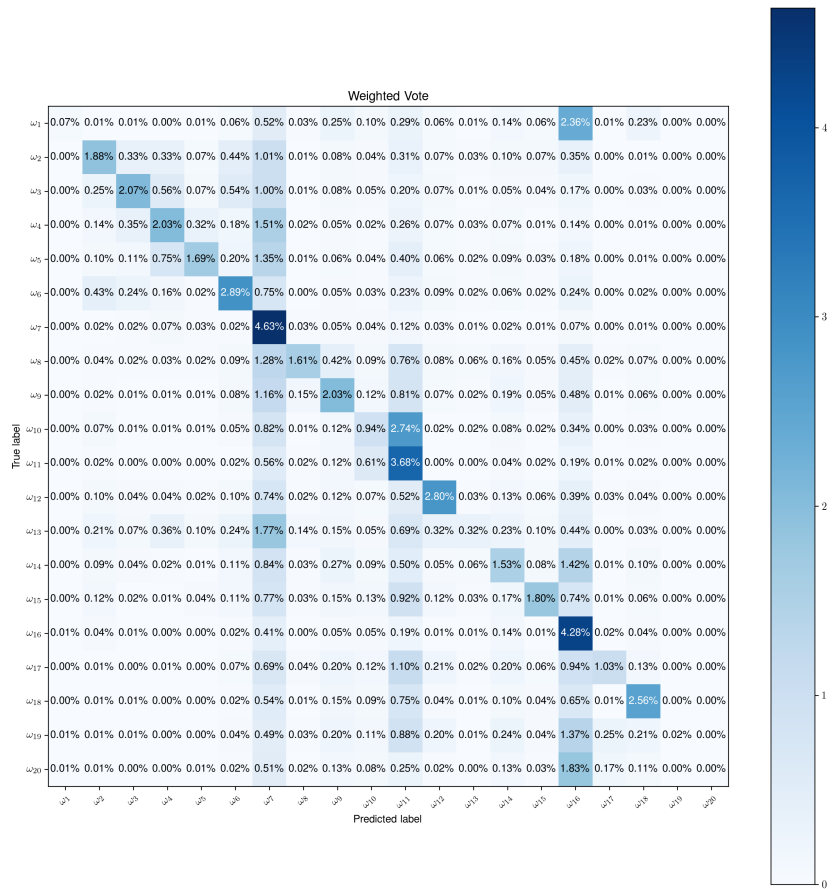
There are actually strong theoretical justifications behind this decision rule coming from decision theory. Indeed, a theorem by Gilboa (1987) explains that decisions based on non-additive measures (such as possibility measures) are in line with several decision theoretic axioms. These axioms are closely related to those of the famous theorem from Savage (1954) which holds for (finitely) additive probability distributions.


AppendixC. Normalized confusion matrices on real data

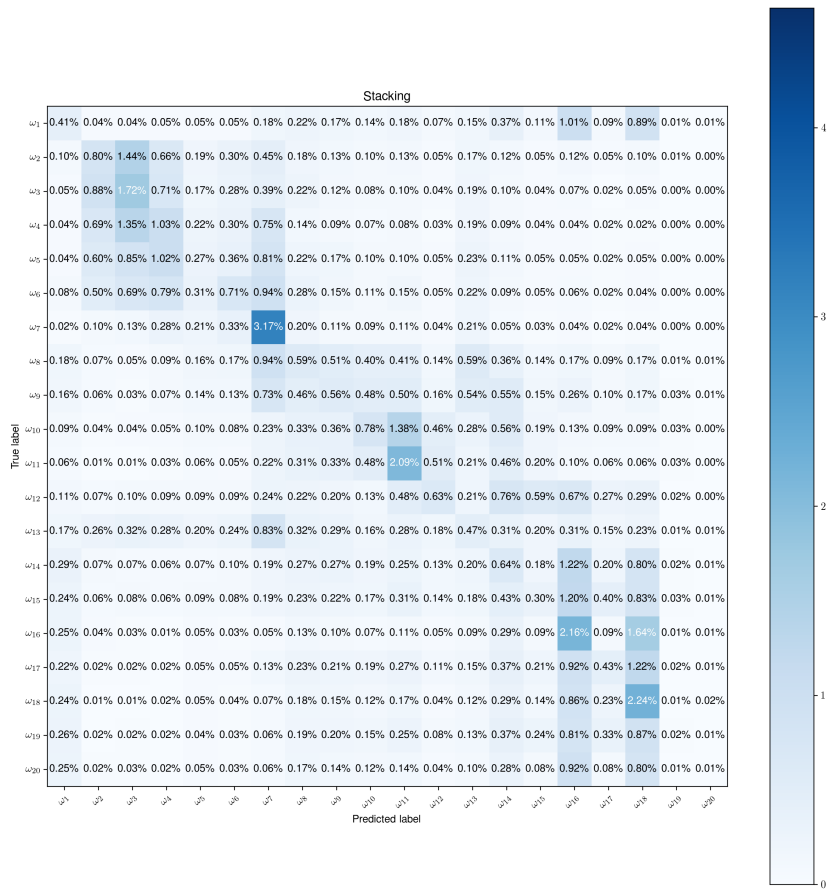
This section provides normalized confusion matrices corresponding to the experiments on real datasets exposed in 4.3 and summarized in Table 3. The matrices are grouped by dataset and for each of the 8 datasets there is one such matrix for each aggregation or reference method. The confusion matrix of the Bayes aggregation is missing for 20newsgroup, Drive and MNIST datasets because Bayes aggregation is intractable as soon as $\ell > 6$ with our equipment.


20newsgroup.










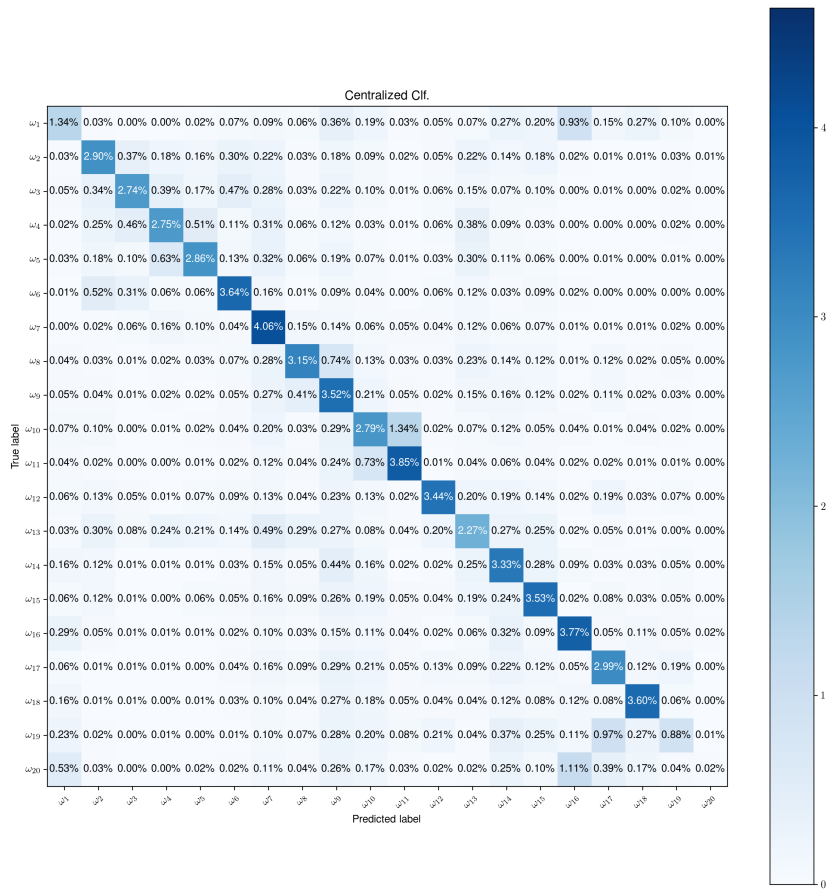








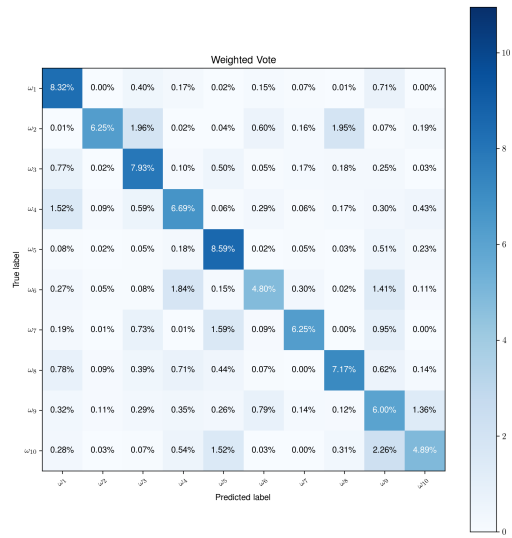










MNIST.















Satellite.



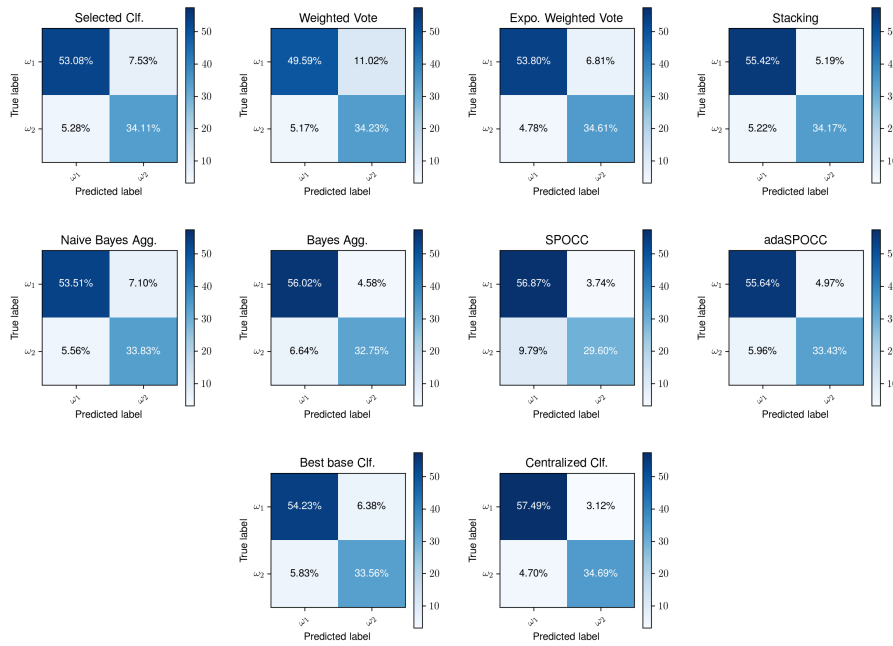


Wine.







Spam.





Avila.








Drive.












Particle.



References

Alsinet, T., Chesñevar, C. I., Godo, L., & Simari, G. R. (2008). A logic programming framework for possibilistic argumentation: Formalization and logical properties. *Fuzzy Sets and Systems*, 159(10), 1208–1228.

- Amor, N. B., Dubois, D., Gouider, H., & Prade, H. (2018). Possibilistic preference networks. *Information Sciences*, *460*, 401–415.
- Balakrishnan, N., & Mojirsheibani, M. (2015). A simple method for combining estimates to improve the overall error rates in classification. *Computational Statistics*, *30*(4), 1033–1049.
- Bella, A., Ferri, C., Hernández-Orallo, J., & Ramírez-Quintana, M. J. (2013). On the effect of calibration in classifier combination. *Applied intelligence*, *38*(4), 566–585.
- Benferhat, S., Dubois, D., Garcia, L., & Prade, H. (1999). Possibilistic logic bases and possibilistic graphs. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence* (pp. 57–64). Morgan Kaufmann Publishers Inc.
- Biau, G., Fischer, A., Guedj, B., & Malley, J. D. (2016). Cobra: A combined regression strategy. *Journal of Multivariate Analysis*, *146*, 18–28.
- Breiman, L. (1996). Bagging predictors. *Machine learning*, *24*(2), 123–140.
- Brown, G., Wyatt, J., Harris, R., & Yao, X. (2005). Diversity creation methods: a survey and categorisation. *Information Fusion*, *6*(1), 5–20.
- De Cooman, G., & Aeyels, D. (1999). Supremum preserving upper probabilities. *Information Sciences*, *118*(1-4), 173–212.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American society for information science*, *41*(6), 391–407.
- Destercke, S., Dubois, D., & Chojnacki, E. (2008). Possibilistic information fusion using maximal coherent subsets. *IEEE Transactions on Fuzzy Systems*, *17*(1), 79–92.
- Dubois, D., Foulloy, L., Mauris, G., & Prade, H. (2004). Probability-possibility transformations, triangular fuzzy sets, and probabilistic inequalities. *Reliable computing*, *10*(4), 273–297.
- Dubois, D., Lang, J., & Prade, H. (1994). Automated reasoning using possibilistic logic: Semantics, belief revision, and variable certainty weights. *IEEE transactions on knowledge and data engineering*, *6*(1), 64–71.
- Dubois, D., & Prade, H. (1982). On several representations of an uncertain body of evidence. *Fuzzy Information and Decision Processes*, (pp. 161–181).
- Dubois, D., & Prade, H. (1988). *Possibility theory: An approach to the computerized processing of information*. Plenum Press, New York.
- Dubois, D., & Prade, H. (1992). When upper probabilities are possibility measures. *Fuzzy sets and systems*, *49*(1), 65–74.
- Dubois, D., & Prade, H. (2015). Possibility theory and its applications: Where do we stand? In *Springer Handbook of Computational Intelligence* (pp. 31–60). Springer.

- Dubois, D., Prade, H., & Schockaert, S. (2017). Generalized possibilistic logic: foundations and applications to qualitative reasoning about uncertainty. *Artificial Intelligence*, 252, 139–174.
- Kampé de Fériet, J. (1982). Interpretation of membership functions of fuzzy sets in terms of plausibility and belief. In E. S. Madan M. Gupta (Ed.), *Fuzzy Information and Decision Processes* (pp. 93–98). North-Holland, Amsterdam.
- Gilboa, I. (1987). Expected utility with purely subjective non-additive probabilities. *Journal of Mathematical Economics*, 16(1), 65 – 88.
- Goodman, I. R. (1982). Fuzzy sets as equivalence classes of random sets. In R. Yager (Ed.), *Fuzzy Sets and Possibility Theory* (pp. pp. 327–343). Pergamon Press, Oxford.
- Guo, K., Xu, T., Kui, X., Zhang, R., & Chi, T. (2019). ifusion: Towards efficient intelligence fusion for deep learning from real-time and heterogeneous data. *Information Fusion*, 51, 215–223.
- Hoang, M., Hoang, N., Low, B. K. H., & Kingsford, C. (2019). Collective model fusion for multiple black-box experts. In *International Conference on Machine Learning* (pp. 2742–2750).
- Huang, Y. S., & Suen, C. Y. (1995). A method of combining multiple experts for the recognition of unconstrained handwritten numerals. *IEEE transactions on pattern analysis and machine intelligence*, 17(1), 90–94.
- Hüllermeier, E. (2002). Possibilistic induction in decision-tree learning. In *European Conference on Machine Learning* (pp. 173–184). Springer.
- Ji, J., Chen, X., Wang, Q., Yu, L., & Li, P. (2019). Learning to learn gradient aggregation by gradient descent. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19* (pp. 2614–2620). International Joint Conferences on Artificial Intelligence Organization.
- Jousselme, A.-L., & Maupin, P. (2012). Distances in evidence theory: Comprehensive survey and generalizations. *International Journal of Approximate Reasoning*, 53(2), 118 – 145. Theory of Belief Functions (BELIEF 2010).
- Kim, H.-C., & Ghahramani, Z. (2012). Bayesian classifier combination. In *Artificial Intelligence and Statistics* (pp. 619–627).
- Kim, T., & Ghosh, J. (2019). On single source robustness in deep fusion models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’ Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32* (pp. 4815–4826). Curran Associates, Inc.
- Lacoste, A., Marchand, M., Laviolette, F., & Larochelle, H. (2014). Agnostic bayesian learning of ensembles. In *International Conference on Machine Learning* (pp. 611–619).
- Li, L., Xu, W., Chen, T., Giannakis, G. B., & Ling, Q. (2019a). Rsa: Byzantine-robust stochastic aggregation methods for distributed learning from heterogeneous datasets. In *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 1544–1551). volume 33.

- Li, Y., Rubinstein, B., & Cohn, T. (2019b). Exploiting worker correlation for label aggregation in crowdsourcing. In *International Conference on Machine Learning* (pp. 3886–3895).
- Liu, H., & Zhang, L. (2019). Advancing ensemble learning performance through data transformation and classifiers fusion in granular computing context. *Expert Systems with Applications*, 131, 20–29.
- Loustau, S. (2008). Aggregation of svm classifiers using sobolev spaces. *Journal of Machine Learning Research*, 9(Jul), 1559–1582.
- Ma, X., Ji, C., Zhang, X., Wang, J., Li, J., & Li, K.-C. (2019). Secure multiparty learning from aggregation of locally trained models. In *International Conference on Machine Learning for Cyber Security* (pp. 173–182). Springer.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*.
- Menon, A. K., Rajagopalan, A., Sumengen, B., Citovsky, G., Cao, Q., & Kumar, S. (2019). Online hierarchical clustering approximations. *arXiv preprint arXiv:1909.09667*, .
- Pei-Zhuang, W., & Sanchez, E. (1982). Treating a fuzzy subset as a projectable random subset. In M. Gupta, & E. Sanchez (Eds.), *Fuzzy Information and Decision Processes* (pp. 213–220). North-Holland, Amsterdam.
- Prade, D. D.-H. (1991). Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proc. of IJCAI* (pp. 419–424). volume 91.
- Rigollet, P., & Tsybakov, A. B. (2012). Sparse estimation by exponential weighting. *Statistical Science*, (pp. 558–575).
- Sales, T. (1982). Fuzzy sets as set classes. *Stochastica*, 6(3), 249–264.
- Savage, L. J. (1954). The foundations of statistics. NY, John Wiley, (pp. 188–190).
- Schapire, R. E. (1990). The strength of weak learnability. *Machine learning*, 5(2), 197–227.
- Serrurier, M., & Prade, H. (2015). Entropy evaluation based on confidence intervals of frequency estimates: Application to the learning of decision trees. In *International Conference on Machine Learning* (pp. 1576–1584).
- Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton (NJ), USA: Princeton University press.
- Shenoy, P. P. (1992). Using possibility theory in expert systems. *Fuzzy Sets and Systems*, 52(2), 129–142.
- Tulyakov, S., Jaeger, S., Govindaraju, V., & Doermann, D. (2008). Review of classifier combination methods. In *Machine learning in document analysis and recognition* (pp. 361–386). Springer.

- Ward Jr, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301), 236–244.
- Wolpert, D. H. (1992). Stacked generalization. *Neural networks*, 5(2), 241–259.
- Wozniak, M., Grana, M., & Corchado, E. (2014). A survey of multiple classifier systems as hybrid systems. *Information Fusion*, 16, 3 – 17. Special Issue on Information Fusion in Hybrid Intelligent Fusion Systems.
- Yao, X., Huang, T., Zhang, R.-X., Li, R., & Sun, L. (2019). Federated learning with unbiased gradient aggregation and controllable meta updating. In *NeuRIPS Workshop on Federated Learning* (p. to appear).
- Zadeh, L. A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1, 3–28.