



HAL
open science

Importance Reweighting for Biquality Learning

Pierre Nodet, Vincent Lemaire, Alexis Bondu, Antoine Cornuejols, Adam Ouorou

► **To cite this version:**

Pierre Nodet, Vincent Lemaire, Alexis Bondu, Antoine Cornuejols, Adam Ouorou. Importance Reweighting for Biquality Learning. 2021 International Joint Conference on Neural Networks (IJCNN), Jul 2021, Shenzhen, China. pp.1-8, 10.1109/IJCNN52387.2021.9533349 . hal-03650243

HAL Id: hal-03650243

<https://hal.science/hal-03650243v1>

Submitted on 24 Apr 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.

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/344756926>

Importance Reweighting for Biquality Learning

Preprint · October 2020

CITATIONS

0

READS

55

4 authors:



Pierre Nodet

Orange Labs

7 PUBLICATIONS 10 CITATIONS

SEE PROFILE



Vincent Lemaire

Orange Labs

184 PUBLICATIONS 991 CITATIONS

SEE PROFILE



Alexis Bondu

Orange Labs, France, Châtillon

56 PUBLICATIONS 301 CITATIONS

SEE PROFILE



Antoine Cornuéjols

AgroParisTech

125 PUBLICATIONS 740 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



COCLICO (ANR Project) [View project](#)



Predictive Clustering [View project](#)

Importance Reweighting for Biquality Learning

Pierre Nodet

Orange Labs

AgroParisTech, INRAE

46 av. de la République
Châtillon, France

Vincent Lemaire

Orange Labs

2 av. P. Marzin

Lannion, France

Alexis Bondu

Orange Labs

46 av. de la République

Châtillon, France

Antoine Cornuéjols

UMR MIA-Paris

AgroParisTech, INRAE

Université Paris-Saclay

16 r. Claude Bernard

Paris, France

Adam Ouorou

Orange Labs

46 av. de la République

Châtillon, France

Abstract—The field of Weakly Supervised Learning (WSL) has recently seen a surge of popularity, with numerous papers addressing different types of “supervision deficiencies”, namely: poor quality, non adaptability, and insufficient quantity of labels. Regarding quality, label noise can be of different types, including completely-at-random, at-random or even not-at-random. All these kinds of label noise are addressed separately in the literature, leading to highly specialized approaches. This paper proposes an original, encompassing, view of Weakly Supervised Learning, which results in the design of generic approaches capable of dealing with any kind of label noise. For this purpose, an alternative setting called “Biquality data” is used. It assumes that a small trusted dataset of correctly labeled examples is available, in addition to an untrusted dataset of noisy examples. In this paper, we propose a new reweighing scheme capable of identifying noncorrupted examples in the untrusted dataset. This allows one to learn classifiers using both datasets. Extensive experiments that simulate several types of label noise and that vary the quality and quantity of untrusted examples, demonstrate that the proposed approach outperforms baselines and state-of-the-art approaches.

Index Terms—Supervised Classification, Weakly Supervised Learning, Biquality Learning, Trusted data, Label noise

I. INTRODUCTION

The supervised classification problem aims to learn a classifier from a set of labeled training examples in order to predict the class of new examples. In practice, conventional classification techniques may fail because of the imperfections of real-world datasets. Accordingly, the field of *Weakly Supervised Learning* (WSL) has recently seen a surge of popularity, with numerous papers addressing different types of “*supervision deficiencies*” [1], namely:

Insufficient quantity: when many training examples are available, but only a small portion is labeled, e.g. due to the cost of labelling. For instance, this occurs in the field of cyber security where human forensics is needed to label attacks. Usually, this issue is addressed by semi-supervised learning (SSL) [2] or active learning (AL) [3].

Poor quality labels: in this case, all training examples are labeled but the labels may be corrupted. This may happen when the labeling task is outsourced to crowd labeling. The Robust Learning to Label Noise (RLL) approaches address this problem [4], with three identified types of label noise: i) the *completely at random* noise which correspond to a uniform probability of label change ; ii) the *at-random* label noise

when the probability of label change depends upon each class, with uniform label changes within each class ; iii) the *not-at-random* label noise when the probability of label change varies over the input space of the classifier. This last type of label noise is recognized as the most difficult to deal with [5], [6].

Inappropriate labels: for instance, in Multi Instance Learning (MIL) [7] the labels are assigned to bags of examples, with positive label indicating that at least one example of the bag is positive. Some scenarios in Transfer Learning (TL) [8] imply that only the labels in the source domain are provided while the target domain labels are not. Often, these non-adapted labels are associated with slightly different learning tasks (*e.g. more precise and numerous classes are dividing the original categories*). Alternatively, non-adapted labels may characterize a differing statistical individual [9] (*e.g. a subpart of an image instead of the entire image*).

All these types of supervision deficiencies are addressed separately in the literature, leading to highly specialized approaches. In practice, it is very difficult to identify the type(s) of deficiencies with which a real dataset is associated. For this reason, we argue that it would be very useful to find a unified framework for Weakly Supervised Learning, in order to design generic approaches capable of dealing with any type of supervision deficiency.

In Section II of this paper, we present “*biquality data*”, an alternative WSL setting allowing a unified view of weakly supervised learning. A generic learning framework using the two training sets of biquality data (*the one trusted and the other one untrusted*) is suggested in Section III. We identify three possible ways of implementing this framework and consider one of them. This article proposes a new approach using example reweighting in Section IV. The effectiveness of this new approach in dealing with different types of supervision deficiencies, without a priori knowledge about them, is demonstrated through experiments with real datasets in Sections V and VI. Finally, perspectives and future works are discussed in Section VII.

II. BIQUALITY DATA

This section presents an alternative setting called “*Biquality Data*” which covers a large range of supervision deficiencies and allows for unifying the WSL approaches. The interested

reader may find a more detailed introduction on WSL and its links with Biquality Learning in [10].

Learning using biquality data has recently been put forward in [11]–[13] and consists in learning a classifier from two distinct training sets, one trusted and the other untrusted. The initial motivation was to unify semi-supervised and robust learning through a combination of the two. We show in this paper that this scenario is not limited to this unification and that it can cover a larger range of supervision deficiencies as demonstrated with the algorithms we propose and the obtained results.

The trusted dataset D_T consists of pairs of labeled examples (x_i, y_i) where all labels $y_i \in \mathcal{Y}$ are supposed to be correct according to the true underlying conditional distribution $\mathbb{P}_T(Y|X)$. In the untrusted dataset D_U , examples x_i may be associated with incorrect labels. We note $\mathbb{P}_U(Y|X)$ the corresponding conditional distribution.

At this stage, no assumption is made about the nature of the supervision deficiencies which could be of any type including label noise, missing labels, concept drift, non-adapted labels... and more generally a mixture of these supervision deficiencies.

The difficulty of a learning task performed on biquality data can be characterised by two quantities. First, the ratio of trusted data over the whole data set, denoted by p :

$$p = \frac{|D_T|}{|D_T| + |D_U|} \quad (1)$$

Second, a measure of the quality, denoted by q , which evaluates the usefulness of the untrusted data D_U to learn the trusted concept $\mathbb{P}_T(Y|X)$, where $q \in [0, 1]$ and 1 indicates high quality. For example in [13] q is defined using a ratio of Kullback-Leibler divergence between $\mathbb{P}_T(Y|X)$ and $\mathbb{P}_U(Y|X)$.

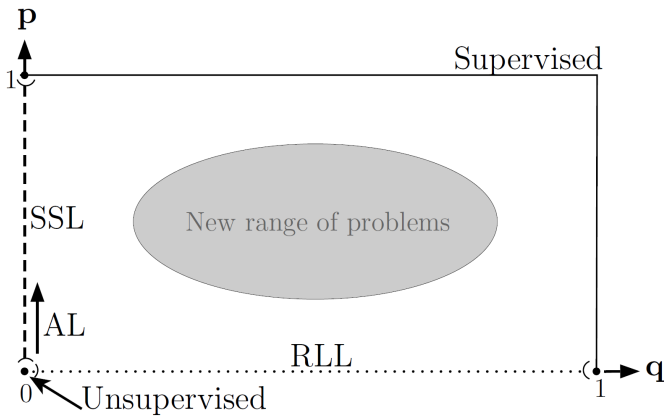


Fig. 1. The different learning tasks covered by the biquality setting, represented on a 2D representation.

The biquality setting covers a wide range of learning tasks by varying the quantities q and p (as represented in Figure 1):

- When $(p = 1 \text{ OR } q = 1)$ ¹ all examples can be trusted. Thus, this setting corresponds to a standard **supervised learning (SL)** task.

¹ $p = 1 \implies D_U = \emptyset \implies q = 1$

- When $(p = 0 \text{ AND } q = 0)$, there is no trusted examples and the untrusted labels are not informative. We are left with only the inputs $\{x_i\}_{1 \leq i \leq m}$ as in **unsupervised learning (UL)**.
- On the vertical axis defined by $q = 0$, except for the two points $(p, q) = (0, 0)$ and $(p, q) = (1, 0)$, the untrusted labels are not informative, and trusted examples are available. The learning task becomes **semi-supervised learning (SSL)** with the untrusted examples as unlabeled and the trusted as labeled.
- An upward move on the vertical axis, from a point $(p, q) = (\epsilon, 0)$ characterized by a low proportion of labeled examples $p = \epsilon$, to a point $(p', 0)$, with $p' > p$, corresponds to **Active Learning**, when an oracle provides new labels asked by a given strategy. The same upward move can also be realized in **Self-training** and **Cotraining** [14], where unlabeled training examples are labeled using the predictions of the current classifier(s).
- On the horizontal axis defined by $p = 0$, except for the points $(p, q) = (0, 0)$ and $(p, q) = (0, 1)$, only untrusted examples are provided, which corresponds to the range of learning tasks typically addressed by **Robust Learning to Label noise (RLL)** approaches.

Only the edges of Figure 1 have been envisioned in previous works – i.e. the points mentioned above – and a whole new range of problems are addressed in this paper. Moreover, biquality learning may be used to tackle tasks belonging to WSL, for instance:

- Positive Unlabeled Learning (PUL) [15] where only positive (trusted) and unlabeled instances are available, the later which can be considered as untrusted.
- Self Training and Cotraining [14] could be addressed at the end of the self labeling process: the initial training set is then the trusted dataset, and all self-labeled examples can be considered as the untrusted ones.
- Concept drift [16]: when a concept drift occurs, all the examples used before a detected drift may be considered as the untrusted examples, while the examples available after it are viewed as the trusted ones, assuming a perfect labeling process.
- Self Supervised Learning system as illustrated by Snorkel [17] or Snuba [18]: the small initial training set can be trusted, whereas all examples automatically labeled using the labeling functions may be considered as untrusted.

As can be seen from the above list, the Biquality framework is quite general and its investigation seems a promising avenue to unify different aspects of the Weakly Supervised Learning. A main contribution of this paper is to suggest one generic framework for achieving biquality learning and thus covering many facets of WSL. This is presented in the next section. This framework will be then applied in the experiments part of this paper to the problem of label noise.

III. BIQUALITY LEARNING

Learning the true concept² $\mathbb{P}_T(Y|X)$ on $D = D_T \cup D_U$ means minimizing the risk R on D with a loss L for a probabilistic classifier f :

$$\begin{aligned} R_{D,L}(f) &= \mathbb{E}_{D,(X,Y) \sim T}[L(f(X), Y)] \\ &= \mathbb{P}(X \in D_T) \mathbb{E}_{D_T,(X,Y) \sim T}[L(f(X), Y)] \\ &\quad + \mathbb{P}(X \in D_U) \mathbb{E}_{D_U,(X,Y) \sim T}[L(f(X), Y)] \end{aligned} \quad (2)$$

where $L(\cdot, \cdot)$ is a loss function, from $\mathbb{R}^{|\mathcal{Y}|} \times \mathcal{Y}$ to \mathbb{R} since $f(X)$ is a vector of probability over the classes. Since the true concept $\mathbb{P}_T(Y|X)$ cannot be learned from D_U , the last line of Equation 2 is not tractable as it stands. That is why we propose a **generic formalization** based on a mapping function g that enables us to learn the true concept from the modified untrusted examples of D_U . Equation 2 becomes:

$$\begin{aligned} R_{D,L}(f) &= \mathbb{P}(X \in D_T) \mathbb{E}_{D_T,(X,Y) \sim T}[L(f(X), Y)] \\ &\quad + \lambda \mathbb{P}(X \in D_U) \mathbb{E}_{D_U,(X,Y) \sim U}[g(L(f(X), Y))] \end{aligned} \quad (3)$$

In Equation 3, the parameter $\lambda \in [0, 1]$ reflects the quality of the untrusted examples of D_U modified by the function g . This time, the last line is tractable since it consists of a risk expectancy estimated over the training examples of D_U which follows the untrusted concept $\mathbb{P}_U(Y|X)$, modified by the function g .

Accordingly, the estimation of the expected risk requires to learn three items: g , λ and then f . To learn g , a mapping function between the two datasets, both D_T and D_U are used. Then, either λ is considered as a hyper parameter to be learned using D_T or λ is provided by an appropriate quality measure and is considered as an input of the learning algorithm. Finally, f is learned by minimizing the risk R on D using the mapping g .

In this formalization, the mapping function g plays a central role. Not exhaustively, we identify three different ways of designing the mapping function. For each of these, a different function g' enters the definition of function g :

- The first option consists in **correcting the label** for each untrusted examples of D_U . The mapping function thus takes the form $g(L(f(X), Y)) = L(f(X), g'(Y, X))$, with $g'(Y, X)$ denoting the new corrected labels and $f(X)$ the predictions of the classifier.
- In the second option, the untrusted labels are used unchanged. The untrusted examples X are **moved** in the input space where the untrusted labels becomes correct with respect to the true underlying concept. The mapping function becomes $g(L(f(X), Y)) = L(f(g'(X)), Y)$, where $g'(X)$ is the “moved” input vector of the modified untrusted examples.
- In the last option, g' **weights** the contribution of the untrusted examples in the risk estimate. Accordingly, we have $g(L(f(X), Y)) = g'(Y, X)L(f(X), Y)$. In this case, the parameter λ may disappear from Equation 3 since it can be considered as included in the function g' .

Section IV considers in-depth the last option and proposes a new approach where g' acts as an Importance Reweighting for Biquality Learning.

IV. A NEW IMPORTANCE REWEIGHTING APPROACH FOR BIQUALITY LEARNING

To estimate the mapping function g' , we suggest to adapt the importance reweighting trick from the covariate shift literature [19] to biquality learning. This trick relies on reweighting untrusted samples by using the Radon-Nikodym derivative (RND) [20] of $\mathbb{P}_T(X, Y)$ in respect to $\mathbb{P}_U(X, Y)$ which is $\frac{d\mathbb{P}_T(X, Y)}{d\mathbb{P}_U(X, Y)}$. Contrary to the “covariate shift” setting, the biquality setting handles the same distribution $\mathbb{P}(X)$ in the trusted and untrusted datasets. However, the two underlying concepts $\mathbb{P}_T(Y|X)$ and $\mathbb{P}_U(Y|X)$ are possibly different due to a supervision deficiency. By using these assumptions and the Bayes Formula, we can further simplifying the reweighting function to the RND of $\mathbb{P}_T(Y|X)$ in respect to $\mathbb{P}_U(Y|X)$, $\frac{d\mathbb{P}_T(Y|X)}{d\mathbb{P}_U(Y|X)}$.

$$\begin{aligned} R_{(X,Y) \sim T,L}(f) &= \mathbb{E}_{(X,Y) \sim T}[L(f(X), Y)] \\ &= \int L(f(X), Y) d\mathbb{P}_T(X, Y) \\ &= \int \frac{d\mathbb{P}_T(X, Y)}{d\mathbb{P}_U(X, Y)} L(f(X), Y) d\mathbb{P}_U(X, Y) \\ &= \mathbb{E}_{(X,Y) \sim U} \left[\frac{\mathbb{P}_T(X, Y)}{\mathbb{P}_U(X, Y)} L(f(X), Y) \right] \\ &= \mathbb{E}_{(X,Y) \sim U} \left[\frac{\mathbb{P}_T(Y|X)\mathbb{P}(X)}{\mathbb{P}_U(Y|X)\mathbb{P}(X)} L(f(X), Y) \right] \\ &= \mathbb{E}_{(X,Y) \sim U} \left[\frac{\mathbb{P}_T(Y|X)}{\mathbb{P}_U(Y|X)} L(f(X), Y) \right] \\ &= \mathbb{E}_{(X,Y) \sim U} [\beta L(f(X), Y)] \\ &= R_{(X,Y) \sim U, \beta L}(f) \end{aligned} \quad (4)$$

Equation 4 shows that $\beta = \frac{\mathbb{P}_T(Y|X)}{\mathbb{P}_U(Y|X)}$ is an estimation of the mapping function g' , thanks to Section III estimating β is the last step before an actual Biquality Learning algorithm.

Algorithm: Importance Reweighting for Biquality Learning (IRBL)

Input: Trusted Dataset D_T , Untrusted Dataset D_U , Probabilistic Classifier Family \mathcal{F}

- 1 Learn $f_U \in \mathcal{F}$ on D_U
- 2 Learn $f_T \in \mathcal{F}$ on D_T
- 3 **for** $(x_i, y_i) \in D_U$, **where** $y_i \in [1, K]$ **do**
- 4 $\left[\hat{\beta}(x_i, y_i) = \left\langle \frac{f_T(x_i)}{f_U(x_i)} \right\rangle_{y_i} \right.$
- 5 **for** $(x_i, y_i) \in D_T$ **do**
- 6 $\left[\hat{\beta}(x_i, y_i) = 1 \right.$
- 7 Learn $f \in \mathcal{F}$ on $D_T \cup D_U$ with weights $\hat{\beta}$

Output: f

The proposed algorithm, Importance Reweighting for Biquality Learning (IRBL), aims at estimating β from D_T and

²For reasons of space, we denote $\mathbb{P}_T(Y|X)$ by T and $\mathbb{P}_U(Y|X)$ by U .

D_U whatever the unknown supervision deficiency. It consists of two successive steps. First a probabilistic classifier f_T is learned from the trusted dataset D_T and another probabilistic classifier f_U is learned from the untrusted dataset D_U . Thanks to their probabilistic nature each of them estimates $\mathbb{P}_T(Y|X)$ and $\mathbb{P}_U(Y|X)$ by a probability distribution over the set of the K classes. Thus we can estimate the weight β of an untrusted sample (x_i, y_i) by dividing the prediction of $f_T(x_i)$ by $f_U(x_i)$ for the y_i class (see line 4). The weight β for all trusted samples will be fixed to 1 (see line 6). Then a final classifier is learned from both datasets D_T and D_U with examples reweighted by $\hat{\beta}$.

Our algorithm is theoretically grounded, since it is asymptotically equivalent to minimizing the risk on the true concept using the entire data set (see proof in Equation 5).

$$\begin{aligned}
\hat{R}_{D, \hat{\beta}L}(f) &= \frac{1}{|D|} \sum_{(x_i, y_i) \in D} (\mathbb{1}_{(x_i, y_i) \in D_T} L(f(x_i), y_i) \\
&\quad + \mathbb{1}_{(x_i, y_i) \in D_U} \hat{\beta}(x_i, y_i) L(f(x_i), y_i)) \\
&= \frac{1}{|D_T| + |D_U|} \sum_{(x_i, y_i) \in D_T} L(f(x_i), y_i) \\
&\quad + \frac{1}{|D_T| + |D_U|} \sum_{(x_i, y_i) \in D_U} L(f(x_i), y_i) \hat{\beta}(x_i, y_i) \\
&= \frac{p}{|D_T|} \sum_{(x_i, y_i) \in D_T} L(f(x_i), y_i) \\
&\quad + \frac{1-p}{|D_U|} \sum_{(x_i, y_i) \in D_U} L(f(x_i), y_i) \hat{\beta}(x_i, y_i) \\
&= p \hat{R}_{D_T, L}(f) + (1-p) \hat{R}_{D_U, \hat{\beta}L}(f) \\
&\approx p \hat{R}_{D_T, L}(f) + (1-p) \hat{R}_{D_T, L}(f) \\
&\approx \hat{R}_{D_T, L}(f)
\end{aligned} \tag{5}$$

Proof in Equation 5 is an asymptotic result: in practice our algorithm relies on the quality of the estimation of $\mathbb{P}_T(Y|X)$ and $\mathbb{P}_U(Y|X)$ in order to be efficient. In the biquality setting they both could be hard to estimate because of the small size of D_T and the poor quality of D_U .

V. EXPERIMENTS

The aim of the experiments is to answer the following questions: i) is our algorithm properly designed and does it perform better than baselines approaches? ii) is our algorithm competitive with state-of-the-art approaches?

First, Section V-A presents the supervision deficiencies which are simulated in our experiments. They correspond to two different kinds of weak supervision, namely, Noisy label Completely at Random (i.e. not X dependent) and Noisy label Not at Random (i.e. X dependent). From the Frenay's taxonomy [4] the former is the easiest to deal with and the later is often considered as difficult to manage. Then, Section V-B consists of three parts: a presentation of the baseline competitors, a brief report on the state-of-the-art competitors, and a description of the set of classifiers used. Finally, Section

V-C describes the datasets used in the experiments, and the chosen criterion to evaluate the learned classifiers. For full reproducibility, source code, datasets and results are available at : <https://github.com/pierrenodet/irbl>.

A. Simulated supervision deficiencies

The datasets listed in Section V-C consist of a collection of training examples that are assumed to be correctly labeled, denoted by D_{total} . In order to obtain a trusted dataset D_T and an untrusted one D_U , each dataset is split in two parts using a stratified random draw, where p is the percentage for the trusted part. The trusted datasets are left untouched, whereas corrupted labels are simulated in the untrusted datasets by using two different techniques:

a) *Noisy Completely At Random (NCAR)*:: Corrupted untrusted examples are uniformly drawn from D_U with a probability r , and are assigned a random label that is also uniformly drawn from \mathcal{Y} .

In the particular case of binary classification problems, the conditional distribution of the untrusted labels is defined by Equation 6.

$$\forall y \in \mathcal{Y}, \mathbb{P}_U(Y = y|X) = \frac{r}{2} + (1-r)\mathbb{P}_T(Y = y|X) \tag{6}$$

Here, r controls the overall number of random labels and thus is our proxy for the quality: $q = 1 - r$.

b) *Noisy Not At Random (NNAR)*:: Corrupted untrusted examples are drawn from D_U with a probability $r(x)$ that depends on the instance value. To generate a instance dependent label noise, we design a noise that depends on the decision boundary of a classifier f_{total} learned from D_{total} . The probability of random label $r(x)$ should be high when an instance x is close to the decision boundary, and low when it is far. In our experiments, the probability outputs of f_{total} are used to model our label noise as follows:

$$\forall x \in \mathcal{X}, r(x) = 1 - \theta |1 - 2f_{total}(x)|^{\frac{1}{\theta}} \tag{7}$$

where $\theta \in [0; 1]$ is a constant that controls the overall number of random labels and thus is our proxy for the quality: $q = \theta$. The parameter θ influences both the slope (factor) and the curvature (power) of $r(x)$ to modify the area under curve of $r(x)$: $\mathbb{E}[r(x)]$.

For binary classification problems, the conditional distribution of the untrusted labels is defined by Equation 8.

$$\begin{aligned}
&\forall x \in \mathcal{X}, \forall y \in \mathcal{Y}, \\
&\mathbb{P}_U(Y = y|X = x) = \frac{r(x)}{2} + (1-r(x))\mathbb{P}_T(Y = y|X = x)
\end{aligned} \tag{8}$$

B. Competitors

a) *Baseline competitors*: The first part of our experiments consists of a sanity check which compares the performance of the proposed algorithm to the following baselines:

- *Trusted*: The final classifier f obtained with our algorithm should be better than a classifier f_T that learned only from the trusted dataset, insofar as untrusted data bring useful information about the trusted concept. At least, f should not be worse than using only trusted data.

- *Mixed*: The final classifier f should be better than a classifier f_{mixed} learned from both trusted and untrusted dataset, without correction. A biquality learning algorithm should leverage the information provided by having two distinct datasets.
- *Untrusted*: The final classifier should be better than a classifier f_U that learns only from the untrusted dataset if there are trusted labels. Using trusted data should improve the classifier final performances.

b) *State-of-the-art-competitors*: The second part of our experiments compares our algorithm with two state-of-the-art methods: (i) a method from the Robust Learning to Label noise (RLL) [21], [22] family and (ii) the GLC approach [12].

- *RLL*: In recent literature a new emphasis is put on the research of new loss functions that are conducive to better risk minimization in presence of noisy labels. For example, [21], [22] show theoretically and experimentally that when the loss function satisfies a symmetry condition, described below, this contributes to the robustness of the classifier. Accordingly, in this paper we train a classifier with a symmetric loss function as a competitor. This first competitor is expected to have good results on completely-at-random label noise described in Section V-A. A loss function L_s is said *symmetrical* if $\sum_{y \in \{-1, 1\}} L_s(f(x), y) = c$, where c is a constant and $f(x)$ is the score on the class y . This loss function is used on $D_T \cup D_U$.

- *GLC*: To the best of our knowledge, GLC [12] is among the best performing algorithm that can learn from bi-quality data. It has been successfully compared to many competing approaches. Like ours, it is a two steps approach which is simple and easy to implement.

In a first step, a model f_U is learned from the untrusted dataset D_U . Then it is used to estimate a transition matrix C of $\mathbb{P}_{U|T}(Y)$ by making probabilistic predictions with f_U on the trusted dataset D_T and comparing it to the trusted labels.

In a second step, this matrix is used to correct the labels from the untrusted dataset D_U when learning the final model f . Indeed f is learned with L on D_T and with $L(C^\top f(X), Y)$ on D_U .

c) *Classifiers*: First of all, the choice of classifiers was guided by the idea of comparing algorithms for biquality learning and not searching for the best classifiers. This choice was also guided by the nature of the datasets used in the experiments (see section V-C). Secondly our algorithm, as well as GLC, implies two learning phases. For both reasons and for simplicity, we decided to use Logistic Regressions (LR) for each phase. LR is known to be limited, in the sense of the Vapnik-Chervonenkis dimension [23] since it can only learn linear separations of the input space \mathcal{X} , which could underfit the conditional probabilities $\mathbb{P}(Y|X)$ on D_T and D_U and lead

to bad β estimations. But this impediment, if met, will affect equally all the compared algorithms. LR is also used for the RLL classifier using the Unhinged symmetric loss function.

To obtain reliable estimations of conditional probabilities $\mathbb{P}(Y|X)$, the outputs of all classifiers have been calibrated thanks to Isotonic Regression with the default parameters provided by scikit-learn [24].

Logistic Regression is always be used and learned thanks to SGD with a learning rate of 0.005, a weight decay of 10^{-6} during 20 epochs and a batch size of 24 with Pytorch [25].

C. Datasets

In industrial applications familiar to us, such as fraud detection, Customer Relationship Management (CRM) and churn prediction, we are mostly faced with binary classification problems. The available data is of average size in terms of the number of explanatory variables and involves mixed variables (numerical and categorical).

For this reason we limited in this paper the experiments to binary classification tasks even if our algorithm can address multi-class problems. The chosen tabular datasets, used for the experiments, have similar characteristics than those of our real applications.

TABLE I
BINARY CLASSIFICATION DATASETS USED FOR THE EVALUATION.
COLUMNS: NUMBER OF EXAMPLES ($|D|$), NUMBER OF FEATURES ($|\mathcal{X}|$),
AND RATIO OF EXAMPLES FROM THE MINORITY CLASS (MIN).

| name | $ D $ | $ \mathcal{X} $ | min | name | $ D $ | $ \mathcal{X} $ | min |
|----------|--------|-----------------|-----|----------|---------|-----------------|-----|
| 4class | 862 | 2 | 36 | ibnsina | 20,722 | 92 | 38 |
| ad | 3,278 | 1558 | 14 | zebra | 61,488 | 154 | 4.6 |
| adult | 48,842 | 14 | 23 | musk | 6,598 | 169 | 15 |
| aus | 690 | 14 | 44 | phishing | 11055 | 30 | 44 |
| banknote | 1372 | 4 | 44 | spam | 4,601 | 57 | 39 |
| breast | 683 | 9 | 35 | ijcnn1 | 141,691 | 22 | 9 |
| eeg | 1498 | 13 | 45 | svmg3 | 1284 | 4 | 26 |
| diabetes | 768 | 8 | 35 | svmg1 | 7,089 | 22 | 43 |
| german | 1000 | 20 | 30 | sylva | 145,252 | 108 | 6.5 |
| hiva | 42,678 | 1617 | 3.5 | web | 49,749 | 300 | 3 |

They come from different sources: UCI [26], libsvm³ and active learning challenge [27]. A part of these datasets comes from past challenges on active learning where high performances with a low number of labeled examples has proved difficult to obtain. For each dataset, 80 % of samples were used for training and 20% were used for the test. With this choice of datasets, a large range of the class ratio is covered: Australian is almost balanced while Web is really unbalanced. Also, the size varies significantly in number of rows or columns, with corresponding impact on the difficulty of the learning tasks.

VI. RESULTS

The empirical performance of our approach, is evaluated in two steps. First, we investigate the efficiency of the reweighting scheme and its influence on the learning procedure of the final classifier. Second, our approach is benchmarked against competitors to evaluate its efficiency in real tasks.

³<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

A. Behavior of the IRBL method

In order to illustrate the proposed reweighing scheme, we picked a dataset, here the “ad” dataset used with a ratio of trusted data $p = 0.25$, and examined the histogram of the weights assigned to each untrusted example either corrupted or not. The case of Random Label Completely at Random is chosen and the hardest case where all labels are at random $q = 0$ is considered.

Figure 2 shows the histogram of the weights assigned to each untrusted example either corrupted or not. It is clear that the proposed method is able to detect corrupted and non-corrupted labels from the untrusted dataset. Figure 3 confirms this behavior when varying the value of the quality. For a perfect quality, the distribution of the β is unimodal with a median equal to one and a very narrow inter quantile range, whereas, when the quality drops, the distribution of the β for the corrupted labels decreases to zero.

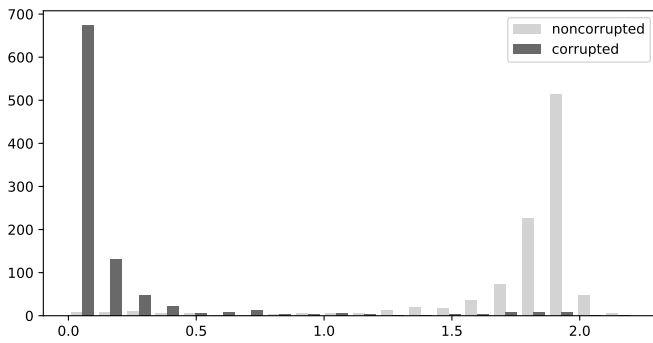


Fig. 2. Histogram of the β values on AD for $p = 0.25$ and $q = 0$ for NCAR for the corrupted and noncorrupted examples.

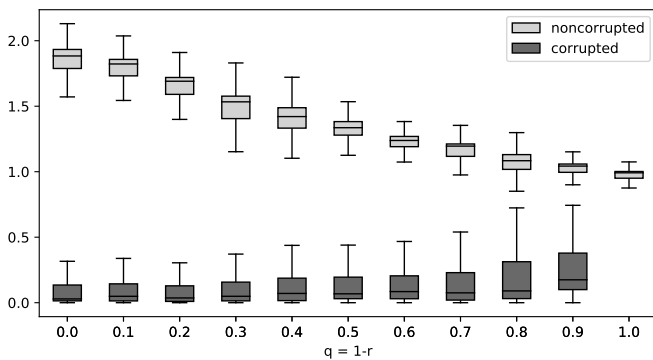


Fig. 3. Boxplot the β values on AD for $p = 0.25$ versus the quality, from $q = 0$ to $q = 1$ for NCAR.

It is equally interesting to look at the classification error when q , the quality of the untrusted data, varies. Figure 4 reports the performance for the proposed method and for the baseline competitors. It is remarkable that the performance of our algorithm, IRBL, remains stable when q decreases while the performance of the *mixed* and *untrusted* algorithms worsens. In addition, IRBL always obtains better performances than the trusted baseline.

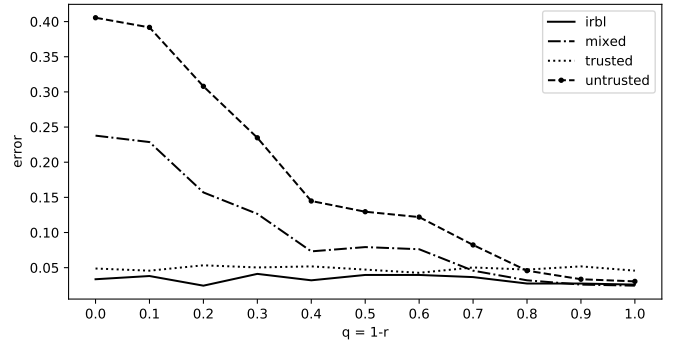


Fig. 4. Classification error on test set for IRBL against baselines on a full range of quality level (AD dataset, $p = 0.25$, NCAR).

B. Comparison with competitors

For a first global comparison, two critical diagrams are presented in Figures 5 and 6 which rank the various methods for the NCAR and NNAR label noise. The Nemenyi test [28] is used to rank the approaches in terms of mean accuracy, calculated for all values of p and q and over all the 20 data sets described in section V-C. The Nemenyi test consists of two successive steps. First, the Friedman test is applied to the mean accuracy of competing approaches to determine whether their overall performance is similar. Second, if not, the post-hoc test is applied to determine groups of approaches whose overall performance is significantly different from that of the other groups.

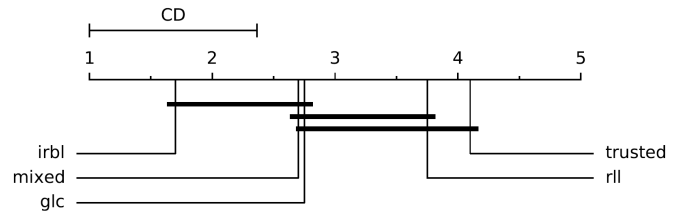


Fig. 5. Nemenyi test for the 20 datasets $\forall p, q$ for NCAR.

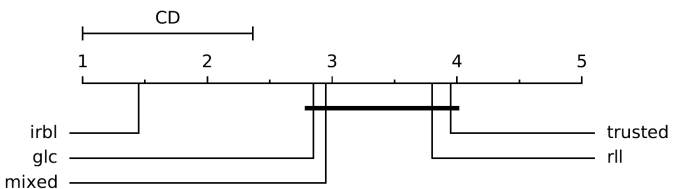


Fig. 6. Nemenyi test for the 20 datasets $\forall p, q$ for NNAR.

These figures show that the IRBL method is ranked first for the two kinds of label noise and provides better performance than the other competitors. Table II provides a more detailed perspective by reporting the mean accuracy and its standard deviation. These values are computed for different values of p over all qualities q and all datasets. This table also helps to see how the methods fare as compared to learning on perfect

TABLE II

MEAN ACCURACY (RESCALED SCORE TO BE FROM 0 TO 100) AND STANDARD DEVIATION COMPUTED ON THE 20 DATASETS $\forall q$ FOR (1) NCAR AND (2) NNAR. THE MEAN ACC WHEN USING ALL THE TRAINING DATA WITHOUT NOISE IS 88.65.

| | p | trusted | irbl | mixed | glc | rll |
|-----|------|------------------|------------------------------------|-------------------|------------------|------------------|
| (1) | 0.02 | 72.48 \pm 5.70 | 83.46 \pm 3.56 | 83.40 \pm 8.30 | 78.34 \pm 7.94 | 77.94 \pm 6.37 |
| | 0.05 | 78.50 \pm 4.33 | 84.94 \pm 2.24 | 83.85 \pm 7.35 | 81.19 \pm 5.15 | 77.97 \pm 6.44 |
| | 0.10 | 81.40 \pm 3.33 | 86.56 \pm 1.68 | 85.44 \pm 5.34 | 83.00 \pm 3.90 | 78.98 \pm 5.26 |
| | 0.25 | 85.61 \pm 2.39 | 87.96 \pm 1.18 | 86.99 \pm 2.80 | 86.27 \pm 2.03 | 79.86 \pm 2.61 |
| (2) | 0.02 | 72.48 \pm 5.70 | 82.93 \pm 3.18 | 81.30 \pm 10.05 | 77.55 \pm 7.78 | 75.47 \pm 9.47 |
| | 0.05 | 78.50 \pm 4.33 | 85.34 \pm 2.55 | 82.52 \pm 7.72 | 80.77 \pm 5.04 | 76.94 \pm 6.64 |
| | 0.10 | 81.40 \pm 3.33 | 86.82 \pm 1.45 | 84.44 \pm 5.14 | 83.22 \pm 4.10 | 77.95 \pm 4.51 |
| | 0.25 | 85.61 \pm 2.39 | 88.21 \pm 1.05 | 86.74 \pm 2.56 | 86.56 \pm 2.00 | 79.67 \pm 2.70 |
| | Mean | 79.50 \pm 3.94 | 85.71 \pm 2.11 | 84.33 \pm 6.16 | 82.11 \pm 4.74 | 78.10 \pm 5.50 |

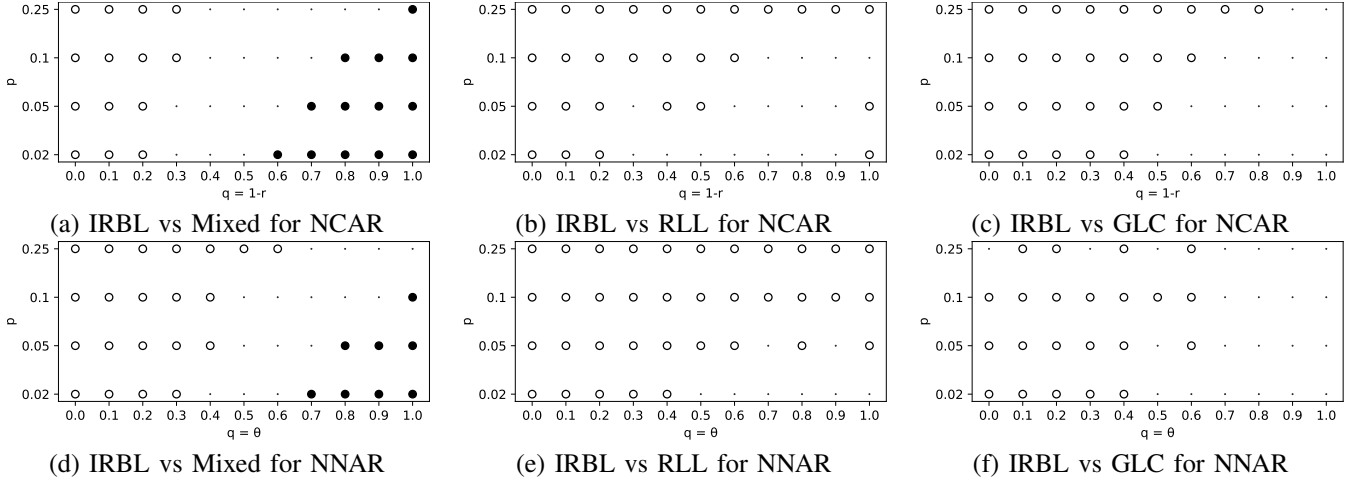


Fig. 7. Results of the Wilcoxon signed rank test computed on the 20 datasets. Each figure compares IRBL versus one of the competitors. Figures a, b, c are in the case of Noisy label Completely at Random and Figures d, e, f for the case of Noisy label Not at Random. In each figure “o”, “.” and “•” indicate respectively a win, a tie or a loss of IRBL compared to the competitors, the vertical axis is p and the horizontal axis is q .

data. Overall, IRBL obtains the best results and with a lower variability.

To get more refined results, the Wilcoxon signed-rank test [29] is used⁴. It enables us to find out under which conditions – i.e. by varying the values of p and q – IRBL performs better or worse than the competitors.

Figure 7 presents six graphics, each reporting the Wilcoxon test that evaluates our approach against a competitor, based on the mean accuracy over the 20 datasets. The two types of label noise (see Section V-A) correspond to the rows in Figure 7 and a wide range of q and p values are considered.

Thanks to these graphs we can compare in more details our method (IRBL) with the mixed methods, as well as with RLL and GLC. Regarding the mixed method, Figures 7 (a) and (b) return the results obtained versus varying values for p and q . For low quality values q , whatever is the value of p , IRBL is significantly better. For middle values of the quality there is no winner and for high quality values and low values of p , the mixed method is significantly better (this result seems to be observed in [12] as well). This is not surprising since at high

quality values, the mixed baseline is equivalent to learning with perfect labels.

These detailed results help us to understand why, in the critical diagram in Figure 5, although IRBL has a better ranking, it is not significantly better than the mixed method: mainly because of the presence of high quality value cases.

Regarding the competitors RLL and GLC, Figures 7(b), 7(c), 7(e) and 7(f) show that IRBL has always better or indistinguishable performances. Indeed, IRBL performs well regardless of the type of noise. This is an important result since it shows that we are able to deal not only with NCAR noise but also with instance dependent label noise (NNAR) which is more difficult. The method RLL gets more ties with IRBL on NCAR than on NNAR as expected. It is noteworthy that GLC has ties with IRBL when the quality is high whatever the label noise.

To sum up, the proposed method has been tested on a large range of types and strengths of label corruptions. In all cases, IRBL has obtained top or competitive results. Consequently, IRBL appears to be a method of choice for applications where biquality learning is needed. Moreover, IRBL has no user parameter and a low computational complexity.

⁴Here the test is used with a confidence level at 5 %.

VII. CONCLUSION

This paper has presented an original view of Weakly Supervised Learning and has described a generic approach capable of dealing with any kind of label noise. A formal framework for *biquality learning* has been developed where the empirical risk is minimized on the small set of trusted examples in addition to some appropriately chosen criterion using the untrusted examples. We identified three different ways to design a mapping function leading to three different such criteria within the *biquality learning* framework. We implemented one of them: a new Importance Reweighting approach for Biquality Learning (IRBL). Extensive experiments have shown that IRBL significantly outperforms state-of-the-art approaches, by simulating completely-at-random and not-at-random label noise over a wide range of quality and ratio values of untrusted data.

Future works will be done to extend experiments with multiclass classification datasets and other classifiers such as Gradient Boosted Trees [30]. An adaptation of IRBL to Deep Learning tasks with an online algorithm will be studied too.

REFERENCES

- [1] Z.-H. Zhou, "A brief introduction to weakly supervised learning," *National Science Review*, vol. 5, no. 1, pp. 44–53, 08 2017.
- [2] O. Chapelle, B. Scholkopf, and A. Zien, "Semi-supervised learning," *IEEE Transactions on Neural Networks*, vol. 20, no. 3, pp. 542–542, 2009.
- [3] B. Settles, "Active learning literature survey," University of Wisconsin-Madison Department of Computer Sciences, Tech. Rep., 2009.
- [4] B. Frénay and M. Verleysen, "Classification in the presence of label noise: a survey," *IEEE transactions on neural networks and learning systems*, vol. 25, no. 5, pp. 845–869, 2013.
- [5] J. Cheng, T. Liu, K. Ramamohanarao, and D. Tao, "Learning with bounded instance and label-dependent label noise," in *International Conference on Machine Learning (ICML)*, vol. 119. PMLR, 13–18 Jul 2020, pp. 1789–1799.
- [6] A. Menon, B. V. Rooyen, and N. Natarajan, "Learning from binary labels with instance-dependent corruption," *ArXiv*, vol. abs/1605.00751, 2016.
- [7] M.-A. Carbonneau, V. Cheplygina, E. Granger, and G. Gagnon, "Multiple instance learning: A survey of problem characteristics and applications," *Pattern Recognition*, vol. 77, p. 329–353, May 2018.
- [8] K. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," *Journal of Big data*, vol. 3, no. 1, p. 9, 2016.
- [9] D. Conte, P. Foggia, G. Percannella, F. Tufano, and M. Vento, "A method for counting people in crowded scenes," in *2010 7th IEEE International Conference on Advanced Video and Signal Based Surveillance*, 2010, pp. 225–232.
- [10] P. Nodet, V. Lemaire, A. Bondu, A. Cornuéjols, and A. Ouorou, "From Weakly Supervised Learning to Biquality Learning: an Introduction," in *In Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 2021.
- [11] M. Charikar, J. Steinhardt, and G. Valiant, "Learning from untrusted data," in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, 2017, p. 47–60.
- [12] D. Hendrycks, M. Mazeika, D. Wilson, and K. Gimpel, "Using trusted data to train deep networks on labels corrupted by severe noise," in *Advances in Neural Information Processing Systems 31*, 2018, pp. 10456–10465.
- [13] R. Hataya and H. Nakayama, "Unifying semi-supervised and robust learning by mixup," in *The 2nd Learning from Limited Labeled Data Workshop, ICLR*, 2019.
- [14] J. Zhao, X. Xie, X. Xu, and S. Sun, "Multi-view learning overview: Recent progress and new challenges," *Information Fusion*, vol. 38, pp. 43 – 54, 2017.
- [15] J. Bekker and J. Davis, "Learning from positive and unlabeled data: a survey," *Machine Learning*, vol. 109, pp. 719–760, 2020.
- [16] J. Gama, I. Zliobaite, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, 2014.
- [17] A. Ratner, S. H. Bach, H. Ehrenberg, J. Fries, S. Wu, and C. Ré, "Snorkel: Rapid training data creation with weak supervision," *The VLDB Journal*, vol. 29, no. 2, pp. 709–730, 2020.
- [18] P. Varma and C. Ré, "Snuba: Automating weak supervision to label training data," *Proc. VLDB Endow.*, vol. 12, no. 3, p. 223–236, Nov. 2018.
- [19] T. Liu and D. Tao, "Classification with noisy labels by importance reweighting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 3, p. 447–461, Mar 2016.
- [20] O. Nikodym, "Sur une généralisation des intégrales de m. j. radon," *Fundamenta Mathematicae*, vol. 15, no. 1, pp. 131–179, 1930. [Online]. Available: <http://eudml.org/doc/212339>
- [21] B. van Rooyen, A. Menon, and R. C. Williamson, "Learning with symmetric label noise: The importance of being unhinged," in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds., 2015, pp. 10–18.
- [22] N. Charoenphakdee, J. Lee, and M. Sugiyama, "On symmetric losses for learning from corrupted labels," in *International Conference on Machine Learning*, vol. 97, 2019, pp. 961–970.
- [23] V. N. Vapnik, "The nature of statistical learning theory," 1995.
- [24] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in python," *Journal of machine learning research*, vol. 12, pp. 2825–2830, 2011.
- [25] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, 2019, pp. 8024–8035.
- [26] D. Dua and C. Graff, "Uci machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [27] I. Guyon, "Datasets of the active learning challenge," University of Wisconsin-Madison Department of Computer Sciences, Tech. Rep., 2010.
- [28] P. Nemenyi, "Distribution-free multiple comparisons," *Biometrics*, vol. 18, no. 2, p. 263, 1962.
- [29] F. Wilcoxon, "Individual comparisons by ranking methods," *Biometrics Bulletin*, vol. 1, no. 6, pp. 80–83, 1945. [Online]. Available: <http://www.jstor.org/stable/3001968>
- [30] J. H. Friedman, "Greedy function approximation: a gradient boosting machine," *Annals of statistics*, pp. 1189–1232, 2001.