

Document downloaded from:

<http://hdl.handle.net/10251/61325>

This paper must be cited as:

Hernández Orallo, J. (2014). Probabilistic reframing for cost-sensitive regression. ACM Transactions on Knowledge Discovery from Data. 8(4):1-55. doi:10.1145/2641758.



The final publication is available at

<http://dx.doi.org/10.1145/2641758>

Copyright Association for Computing Machinery (ACM)

Additional Information

© ACM, 2014. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM Transactions on Knowledge Discovery from Data (TKDD), VOL. 8, ISS. 4, (October 2014) <http://doi.acm.org/10.1145/2641758>

# Probabilistic reframing for cost-sensitive regression

José Hernández-Orallo, Universitat Politècnica de València

Common day applications of predictive models usually involve a full use of the available contextual information. When the operating context changes, one may fine-tune the by-default (incontextual) prediction or may even abstain from predicting a value (a reject). *Global* reframing solutions, where the same function is applied to adapt the estimated outputs to a new cost context, is a possible solution here. An alternative approach, which has not been studied in a comprehensive way for regression in the knowledge discovery and data mining literature, is the use of a *local* (e.g., probabilistic) reframing approach, where decisions are made according to the estimated output *and* a reliability, confidence or probability estimation. In this paper, we advocate for a simple two-parameter (mean and variance) approach, working with a *normal* conditional probability density. Given the conditional mean produced by any regression technique, we develop lightweight ‘enrichment’ methods that produce good estimates of the conditional variance, which are used by the *probabilistic* (local) reframing methods. We apply these methods to some very common families of cost-sensitive problems, such as optimal predictions in (auction) bids, asymmetric loss scenarios and rejection rules.

Categories and Subject Descriptors: H.2.8 [**Database Management**]: Database Applications - Data mining; G.3 [**Mathematics of Computing**]: Probability and statistics - Correlation and regression analysis; I.2.6 [**Artificial Intelligence**]: Learning; I.5.1 [**Pattern Recognition**]: Models

General Terms: Performance, Experimentation, Reliability, Measurement

Additional Key Words and Phrases: Cost-sensitive regression, Conditional density estimation, Reframing, Reliability estimation in regression, Asymmetric loss, Calibration

## 1. INTRODUCTION

Consider a common case where a regression model has been built from some training data, and the model has to be deployed to new instances. If the context is the same for the new instances as it was for the training data, then the quality of the predictions will mostly depend on the observed quality of the model for the same context. However, if the context changes, the prediction given by the model may be suboptimal. For instance, if the model has been trained with a symmetric *loss function* but the deployment operating context involves an asymmetric loss function (where, e.g., underestimations have higher loss than overestimations), then predictions will need to be adjusted. In order to do this there are two options: (1) re-train or revise the model by using a possibly modified (e.g., oversampled) training data and the new loss function, or (2) use what we call a *reframing* function which takes the model and the operating context and outputs a new reframed prediction. The first option is not always possible since many regression methods are not cost-sensitive or cannot be (easily) adapted to work with different (possibly complex) loss functions. Also, in the cases where the first option is possible (the training data must be preserved indefinitely), the recurrent re-training (especially if the context changes for each instance) may have an important computational cost for large datasets or may be impaired by a context overfitting. The second option, reframing, uses a decision rule or mechanism that adapts the predictions of a model to another context. In regression, reframing has been used with global adaptation functions that modify all predictions using the same function (such as [Bansal et al. 2008; Zhao et al. 2011], who use the term ‘tuning’ instead).

---

Author’s addresses: José Hernández-Orallo ([jorallo@dsic.upv.es](mailto:jorallo@dsic.upv.es)). *Departament de Sistemes Informàtics i Computació. Universitat Politècnica de València, Spain.*

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© YYYY ACM 1556-4681/YYYY/01-ARTA \$15.00

DOI : <http://dx.doi.org/10.1145/0000000.0000000>

The kind of problem mentioned above has been profusely studied for classification, where the notion of operating context (or condition) is common and well understood. Some of the techniques and notions for addressing these cases are cost matrices, cost-sensitive classification [Elkan 2001], ROC analysis [Swets et al. 2000; Fawcett 2006a], threshold-choice methods [Hernández-Orallo et al. 2012; Hernández-Orallo et al. 2013], calibration [Cohen and Goldszmidt 2004; Bella et al. 2009] and, of course, the notions of soft classifiers (outputting a score or probability) versus the notion of crisp classifiers (just outputting a label). Certainly, there have also been a few efforts to find the parallel of these techniques for regression. However, most of them rely on a crisp view of the regression model, i.e., they work with regression models which *just* output a value. Examples of this are the Regression Error Curves [Bi and Bennett 2003], utility-based regression [Torgo 2005; Torgo and Ribeiro 2009], the definition of ranking measures [Rosset et al. 2007] and the use of transformation functions for regression which derive a *global* reframing (‘tuning’) that must be constantly (or polynomially) applied to the output of a *crisp* regression model [Bansal et al. 2008; Zhao et al. 2011]. None of these approaches reflects a clear parallelism of the usual approach for cost-sensitive classification. In fact, in classification, reframing is not commonly performed by using crisp classifiers (where some labels could be swapped in a stochastic way in order to achieve a desired class distribution or loss). Rather, reframing in classification is usually performed with soft classifiers, where the decision rule combines a rank or class probability estimation with the operating context in order to give a contextualised prediction. As a working hypothesis, we could think of a parallel of soft classification models in regression, known as *soft regression* models. Namely, whenever we consider a scoring classifier (or a ranker) in classification, which can sort their predictions by their reliability (at least in the binary case), we should consider a regression model which can sort their predictions by their reliability. Whenever we consider a probabilistic classifier, which in fact outputs a discrete distribution on the labels (a categorical distribution), we should consider a regression model which outputs a continuous distribution (e.g., a normal distribution), and not a single value. This correspondence is shown in Table I.

Table I. Correspondence between several types of classification and regression. Evaluation also depends on the kind of prediction. For instance, crisp prediction implies the comparison of the estimated output with the actual output, while probabilistic prediction implies the comparison of discrete distributions in classification ( $p(y|x)$  with  $\hat{p}(y|x)$ ) and the comparison of continuous distributions in regression ( $f(y|x)$  with  $\hat{f}(y|x)$ ).

	Classification	Regression
Crisp	a class label $\hat{c}(x)$	a numerical value $\hat{m}(x)$
Soft (Scoring)	a score for each class $\hat{s}_c(x)$	a numerical value $\hat{m}(x)$ and a reliability measure $\hat{r}(x)$ (e.g., confidence interval)
Soft (Probabilistic)	a categorical distribution (characterised by a conditional probability function $\hat{p}(y x)$ ).	a continuous distribution (characterised by a conditional density function $\hat{f}(y x)$ )

This correspondence suggests that the natural way of addressing context-sensitive problems in regression may be the use of soft regression models (as soft classification models are the natural way of addressing context-sensitive problems). We would like to explore regression techniques that not only output the estimated expected value for each instance  $x$ , i.e.  $\mathbb{E}(y|x)$  (also referred to as the conditional mean), but also accompany these predictions with an estimated error, reliability or density function. There are many approaches for this. One approach is to obtain the standard error for each prediction as calculated by each specific technique (e.g., linear regression) if the algorithm provides a way to obtain this value for each prediction (which is not always the case). A second approach is to estimate the “reliability of individual regression predictions” [Bosnić and Kononenko 2008], through sensitivity analysis, local averaging or other techniques, which can be applied to any regression method, as shown in [Bosnić and Kononenko 2009]. A third approach is *conformal prediction* [Shafer and Vovk 2008; Papadopoulos et al. 2002; Papadopoulos 2008; Papadopoulos et al. 2011], or any other method which derives a confidence interval. Finally, we can, of course, use conditional density (or distribution) estimation methods [Rosenblatt 1969; Hyndman et al. 1996;

Hall et al. 1999], which can derive the conditional probability density function of the dependent variable  $y$ , i.e.  $f(y|x)$ , by using kernel or distribution mixtures.

However, given these approaches for soft regression, none of them has been generally applied for context-sensitive problems, because either these proposals are inappropriate, or are much too complex. For instance, standard errors, reliability metrics and confidence intervals are useful to rank the predictions according to their reliability, or to address some tolerance issues, but they cannot be used to get a precise quantifiable magnitude of what the *expected loss* will be for an instance and a specific operating context. On the other hand, conditional density estimation looks like the appropriate setting for this, since we can (theoretically) calculate the expected loss (i.e., the risk) as an integral over all the possible values for the dependent variable, weighted by its density estimation. The problem is that it is not easy (especially for the data mining practitioner) to calculate this minimisation since the estimated density function may be non-monotonic, non-convex or not even continuous.

In this paper we propose a simple approach for soft regression. In most cases, it may be just sufficient to have a good estimation of a conditional *normal* (i.e., Gaussian) density function — which does not mean that the *true* conditional density function must be normal. Using a normal density function has several advantages. First, a normal distribution only needs two parameters, the mean (expected value) and the variance. This makes it possible to estimate these two parameters easily. It can be done from the regression methods themselves, from any estimation of the standard error, from a confidence interval and, of course, from a (more complex) density function or mixture thereof. Second, the variance can be used to rank predictions in a very straightforward way, as is done with reliabilities, but with a clear interpretation of the magnitudes. Third, and most importantly, we can work analytically with the normal distribution and smoothly derive the exact expression leading to the output that minimises the expected loss for many common loss functions.

In fact, we will see that there are extremely simple methods to estimate this variance which can be applied to any *crisp* regression technique. Some of these methods are just based on comparing the prediction for the training dataset with the actual value, disregarding the input domain. In this sense, these methods are closely related to calibration methods in classification, which aim at improving the estimated conditional probability such that it better matches the true conditional probability. However, we call them ‘enrichment’ methods since they preserve the original prediction mean, while only adding a second parameter, the variance, to form a more powerful and flexible soft regression model for context-sensitive applications.

Many common applications of regression where deployment contexts can change are then solved by this setting: cost-sensitive applications where we have asymmetric losses, screening applications where we need rejection rules to determine the examples for which no prediction will be issued, auction and retailing bids where prices (or other continuous variables) are chosen to obtain the maximum expected profit, situations where we want to derive the probability that two or more predictions are in the right order, etc.

In what follows, we analyse and derive the solutions for many of these problem families, using two-parameter regression models and deriving the optimal prediction for the corresponding loss function. For each of these families, we perform a complete set of experiments which show that our general approach is not worse than some specific solutions in the literature for some of these families, and is clearly better for others. As a result, the setting and methodology we introduce in this paper can be effectively (and easily) applied to a wide range of context-sensitive problems.

In brief, the goal of the paper is to show that two-parameter regression (as an estimation of conditional mean and variance or, more easily, as an enrichment of a crisp model by simple conditional variance estimations), followed by a probabilistic local reframing using a normal distribution, is a simple, general and powerful method which can successfully address many kinds of problems. Although relevant on their own, the enrichment techniques can be seen as *instrumental* for the reframing techniques, as the ultimate goal is to see that there are different ways of performing reframing in regression (we will introduce some new terminology for these different ways, mostly the distinction

between local reframing using soft regression models and global reframing using crisp regression models). Under this general view we will compare these different approaches experimentally on the three problem families and discuss the advantages and disadvantages of each approach.

The paper is organised as follows. Section 2 introduces some notation about regression models and loss functions, as well as relevant previous work. Section 3 defines different types of reframing and the optimal prediction expressions for probabilistic reframing. From here, section 4 re-states the objective of the paper in a more precise way, analyses the nature of the true conditional density and the estimated conditional density and introduces metrics to evaluate their divergence. The experimental methodology is defined and settled for the rest of the paper. Section 5 analyses, compares and makes a selection of several conditional density estimation methods and other methods that obtain soft regression models or that can be used as enrichment methods. Sections 6, 7 and 8 formally derive the optimal decision rules using an estimated conditional normal distribution for three families of loss functions: bids (such as those common in auctions, sales or other trading scenarios), asymmetric losses (absolute or squared) and rejection rules in regression. For the three cases, we perform a complete set of experiments using probabilistic reframing with different enrichment methods and compare them against global reframing methods. Section 9 makes a comprehensive analysis of results and applications, a summary of the contributions and the work ahead. Several appendices (which can be skipped on a first reading) complete the paper with some additional information for the datasets used in the experiments, more detailed results for some techniques and some proofs.

## 2. BACKGROUND

We start with some basic definitions and notation, followed by some related work. Then we introduce the key notion of reframing (and the distinction between global and local reframing).

### 2.1. Regression and conditional densities

Let us consider a multivariate input (or predictor) domain  $\mathbb{X} \subset \mathbb{R}^d$  and a univariate output (or response) domain  $\mathbb{Y} \subset \mathbb{R}$ . The domain space  $\mathbb{D}$  is then  $\mathbb{X} \times \mathbb{Y}$ . Labelled examples or instances are just pairs  $\langle x, y \rangle \in \mathbb{D}$ , and datasets are subsets of  $\mathbb{D}$ . Unlabelled examples are elements  $x \in X$ , sometimes represented as  $\langle x, ? \rangle$ . We denote by  $D_X$  and  $D_Y$  the projection of a dataset  $D$  for the input domain and output domain respectively. The true or target ‘model’ (the function we want to learn or approximate) will be represented as a function  $m : \mathbb{X} \rightarrow \mathbb{Y}$ . We usually consider that a labelled dataset is generated from  $m$  using some prior distribution on  $X$  (possibly with some degree of noise). Occasionally, we will use  $f(y|x)$  to denote the true model as a conditional density function, even if we consider  $m$  to be deterministic. In fact, when considering a dataset, with no repeated (or close) values of  $x$  for two different examples, we may just consider that  $f$  is actually a Dirac delta function (all the density mass falls over the true single value). This is clearly different from a normal distribution. In any case, we will work as if  $f$  could be any distribution.

A regression model is an estimation of the true model. A *crisp* (or *hard*) regression model  $\hat{m}$  is a function  $\hat{m} : \mathbb{X} \rightarrow \mathbb{Y}$ . Predictions (or estimations) are usually denoted by  $\hat{y}$ . The term *residual* is used for the difference between the estimated value and the true value, i.e.,  $\hat{y} - y$ . A *soft regression model* accompanies each prediction with a reliability, confidence or, more generally, a conditional probability density function  $\hat{f}(y|x)$  with  $y \in \mathbb{Y}$  and  $x \in \mathbb{X}$ . The corresponding cumulative distribution function is  $\hat{F}(y|x) = \int_{-\infty}^y \hat{f}(t|x) dt = \hat{p}(Y \leq y|x)$ , i.e. the probability of the output being lower or equal than  $y$  for an input  $x$ . The estimated expected value (conditional mean) is denoted by  $\hat{\mu}_{\hat{f}}(x) \triangleq \mathbb{E}_{\hat{f}}(y|x) = \int_{-\infty}^{\infty} y \hat{f}(y|x) dy$ . We denote its (conditional) standard deviation by  $\hat{\sigma}_{\hat{f}}(x)$ . We will drop the subindices when clear from the context. Note that the mean and the standard deviation are conditional, i.e., defined for one single example; these are not the mean and standard deviation of a distribution of examples (or a whole dataset).

The normal distribution will be represented as usual  $\mathcal{N}(\mu, \sigma^2)$ , with probability density function  $\phi_{\mu, \sigma^2}(\cdot)$  and cumulative distribution function  $\Phi_{\mu, \sigma^2}(\cdot)$ . For the standard probability density function

and the standard cumulative distribution function we will drop the subindices, and we will write  $\phi(\cdot)$  and  $\Phi(\cdot)$  respectively.

## 2.2. Cost-sensitive problems and loss functions

In context-sensitive learning, there are several features which describe a context, such as the data distribution, the costs of using some input variables and the loss of the errors over the output variables. In this paper, we focus on loss functions over the output. As we will see, by properly defining the loss function and its parameters we can analyse and address many problem families. Let us start with the definition of loss function:

*Definition 2.1.* A loss function is any function  $\ell : \mathbb{Y} \times \mathbb{Y} \rightarrow \mathbb{R}$  which compares elements in the output domain. For convenience, the first argument will be the estimated value, and the second argument the actual value, so its application is usually denoted by  $\ell(\hat{y}, y)$ .

Typical examples of loss functions are the absolute error ( $\ell^A$ ) and the squared error ( $\ell^S$ ), with  $\ell^A(\hat{y}, y) = |\hat{y} - y|$  and  $\ell^S(\hat{y}, y) = (\hat{y} - y)^2$ . These two loss functions are *symmetric*, i.e. for every  $y$  and  $r$  we have that  $\ell(y + r, y) = \ell(y - r, y)$ . They are also commutative, i.e., for every  $y_1$  and  $y_2$  we have that  $\ell(y_1, y_2) = \ell(y_2, y_1)$ .

While many methods use these generic loss functions (such as  $\ell^S$ ), most applications have different loss functions. For instance, the bounded absolute error ( $\ell_{BA, \beta}$ ) is defined as  $\ell_{BA, \beta}(\hat{y}, y) = \min(|\hat{y} - y|, \beta)$ , which is also symmetric and commutative. Another example is the bid loss function  $\ell_{\beta}^B(\hat{y}, y) = -\hat{y} + \beta$  if  $\hat{y} \leq y$  and 0 otherwise, which is clearly asymmetric. Note that some loss functions can be defined in terms of the residuals  $\hat{y} - y$ , but others cannot. In practice, there can be specialised loss functions for virtually any application domain.

While some previous works in the literature of regression techniques have focussed on re-designing some learning techniques to account for specific loss functions during training ([Crone 2002; Jino et al. 2010]), only a few have considered the problem as a post-hoc process, once the model has been learnt. A post-hoc process can be performed in cases where re-training with the new loss functions is not possible (because the regression technique is not cost-sensitive or because the training data is no longer available). It also has several advantages, such as model reuse and the possibility of applying the same methods to virtually any regression technique. This post-hoc process can be traced back to the seminal work by Granger [Granger 1969], showing that the optimal predictor for some asymmetric losses can be expressed as the conditional mean plus a constant bias term [Granger 1999]. However, it is recognised that solving this term is not always easy (or even possible in closed form) for many loss functions and density functions. Specific results have been studied for some particular loss functions, such as *Lin-Exp* (approximately linear on one side and exponential on the other side) and *Quad-Exp* (approximately quadratic on one side and exponential on the other side), which have general solutions with mild conditions [Zellner 1986]. Conversely, general closed-form solutions for *Lin-Lin* (asymmetric linear) and *Quad-Quad* (asymmetric quadratic) do not exist in general [Christoffersen and Diebold 1996] [Christoffersen and Diebold 1997]. In fact, even general non-closed-form expressions are not always possible unless some constraints are imposed, such as continuous loss functions, finite expected loss and particular properties on the moments of the density function [Elliott and Timmermann 2004]. In general, much of this work is restricted to continuous loss functions in time series or system reliability applications [Basu and Ebrahimi 1992] [Thompson and Basu 1996], but provide sufficient evidence that working with complex density functions is very problematic for general (and possibly discontinuous or non-convex) loss functions. This has motivated the appearance of other approaches which do not use any density estimation, such as the calculation of a global function which is applied to the outputs [Bansal et al. 2008; Zhao et al. 2011]. In this case, the restrictions come on the side of the loss function, which must be convex, and the requirement that the training set (technically, only the true values  $y$ ) must be preserved from the training to the deployment stage.

However, some other problems are not usually considered, not even as generalised loss functions [Granger 1999]. Examples of these context-sensitive problems are *rejection rules*, where we want that the model abstains from outputting a prediction for the most unreliable cases. None of the previous approaches has addressed this problem. In fact, *rejection rules* are common in classification [Ferri and Hernández-Orallo 2004; Pietraszek 2007], but rarely seen as cost-sensitive problems in regression. However, as we will see, these problems can be modelled with a loss function which sets a cost of rejection. Also, many regression problems used for product prescription, sale predictions and auction bids look towards finding the bid price, i.e., the appropriate quantity (or other negotiable feature) which has the maximum expected benefit [Bella et al. 2011]. Many of these problems can also be modelled by a loss function and, yet again, the predictions of the model can be fine-tuned for them.

As a result of this variety and diversity of problem families which can be modelled with loss functions or other kind of context information, we can integrate and generalise some of the existing (and new) *model adaptation* procedures into a more general term that we call *reframing*.

### 3. REFRAMING AND OPTIMAL PREDICTIONS

Given a loss function representing a particular context, the objective is to get predictions with low loss rather than predictions with low error. In order to do this, we do not train a model using this loss function (as risk minimisation approaches could do), because the loss function may not be known at the training stage or the regression technique may not be able to process loss information. Even if possible, re-training a model whenever the context changes is not a very efficient approach in terms of resources —for large datasets or when the loss function changes for each example (as happens in areas such as sales prediction where the loss function usually has some very volatile parameters about the current stocks, delays, etc.). In fact, if we recurrently re-train we may find problems in terms of validation, comprehensibility and reliability if the application requires stable, validated models, and not models which change (are re-trained) for each new instance.

#### 3.1. Reframing

As an alternative, we propose the use of *reframing* functions, which adapt the predictions of the original model to the context, represented by a loss function.

*Definition 3.1.* A reframing transformation is any method which produces a predicted output value given the input value  $x$ , the loss  $\ell$  and the estimated model  $\hat{f}$ .

$$r(x, \ell, \hat{f}) \rightarrow \hat{y} \quad (1)$$

where  $\hat{y}$  represents the reframed output.

Figure 1 shows the process of reframing graphically. Note that we do not impose any restriction on how the training data is obtained from the training context. Also, we do not assume that this data generation process has to be similar to the process generating the unlabelled data from the deployment context. In fact, the distributions of predictor  $x$  and response  $y$  will usually differ between contexts, as we will reflect in the experimental setting.

For those crisp regression techniques that do not produce a density, we can assume a delta Dirac function  $\hat{f}_m$ , whose mean is clearly the prediction point given by the model, i.e.,  $\hat{m}(x) = \hat{\mu}_{\hat{f}}(x)$ . We can still define reframing methods by expressing them as a *global* function which only depends on the loss function and the expected mean for each example, as follows.

$$r(x, \ell, \hat{f}) = R(\ell, \hat{\mu}_{\hat{f}}(x))$$

*Example 3.2.* For instance, given the bid loss function  $\ell_{\beta}^B$  with  $\beta = 0$ , we might consider the following global reframing transformation:

$$R_1(\ell_{\beta}^B, \hat{y}) = 0.8 \times \hat{y}$$

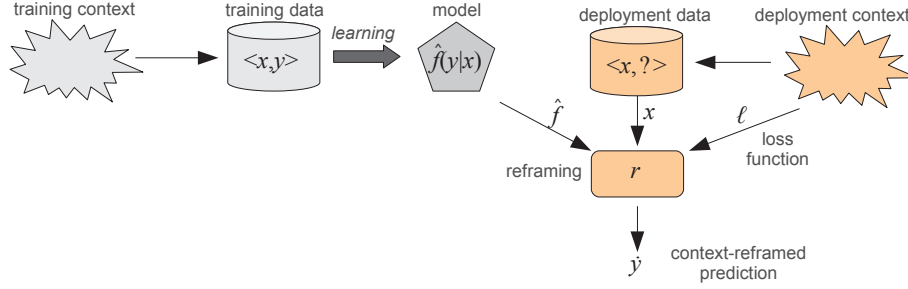


Fig. 1. Reframing process adapting predictions from one context to a different context.

which just systematically reduces predictions by a 20%. The rationale here (e.g., from the point of view of an auctioneer) is given by the fact that overestimations imply a 0 loss (no deal) and underestimations always imply a benefit (there is a deal, and we assume that  $y$  always represent positive prices, so giving a negative loss). So, a 20% reduction creates some margin which may produce higher overall benefit.

Alternatively, a local (probabilistic) reframing could be done using  $\hat{f}$ , if available. For instance, using the same bid loss function  $\ell_\beta^B$  with  $\beta = 0$ , a probabilistic reframing might be:

$$r_2(x, \ell_\beta^B, \hat{f}) = \hat{F}^{-1}(0.25)$$

where  $\hat{F}^{-1}$  is the quantile function for  $\hat{f}$ , i.e., inverse of the cumulative function  $\hat{F}$ . This means that we predict the value such that 25% of the expectancy for  $y$  is below that value.

Figure 2 shows the use of  $R_1$  and  $r_2$  above for two different instances.

In general, we can distinguish four kinds of reframing:

- Constant global reframing: all predictions are modified in the same way independently of  $\hat{y}$ , e.g. adding a constant  $s$  (i.e.,  $\hat{y} \leftarrow \hat{y} + s$ ). This constant  $s$  is called the *shift*.
- Non-constant global reframing: predictions are modified using a (e.g. polynomial) function of  $\hat{y}$ . While the shift is different for each example, it only depends on the prediction, and it can be considered a ‘global’ method, since it applies a global function.
- Non-probabilistic local reframing: predictions are modified by a transformation of  $\hat{y}$  using some reliability or confidence parameters. For instance, we could define a reframing method which only modifies (or rejects) the instances that are below a given reliability threshold or above a percentage of the confidence width.
- Probabilistic local reframing: the outputs are adapted according to a transformation over the conditional density function. If  $\hat{f}$  is a parametric distribution, we can just use the parameters as arguments for the transformation. For instance, if  $\hat{f}$  is a normal distribution for each example, then we can just define the reframing transformation in terms of the conditional mean  $\hat{y} = \hat{\mu}_{\hat{f}}(x)$  and the conditional standard deviation  $\hat{\sigma}_{\hat{f}}(x)$ .

The notion of shift (and reframing in general) for regression is closely related to similar procedures which are usually performed in classification. For instance, when we have a crisp classifier, we can (randomly) tweak some of the predictions, in order to favour one class against others according to a cost matrix. This view would correspond to the global reframing methods above. For soft classifiers, the choice of an appropriate threshold to convert scores into predictions would correspond to local reframing. In particular, when we work with calibrated classifiers and we try to find the optimal thresholds (as in [Lachiche and Flach 2003]) using a probabilistic setting, we have a scenario that is parallel to probabilistic local reframing.



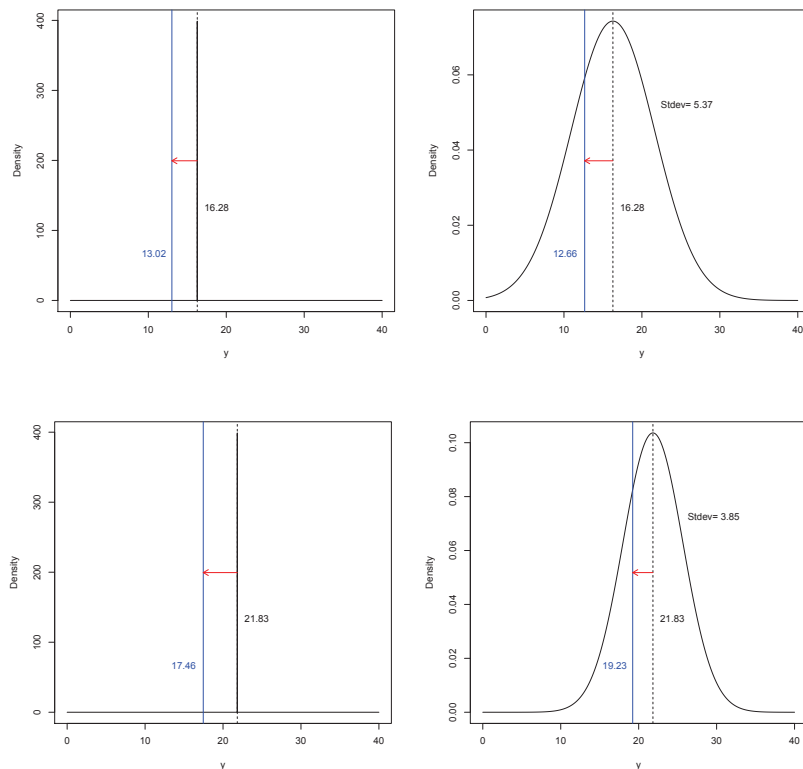


Fig. 2. Representation of how predictions are shifted according to different reframing methods. The top row corresponds to an instance whose expected value is 16.28. Without reframing, this would be the output value. The top-left plot shows how the prediction moves to 13.02 (applying the global reframing  $R_1$ , which multiplies the prediction by 0.8). The top-right plot shows a soft regression model with a normal conditional density function  $\hat{f}$  which gives a mean at 16.28 and standard deviation of 5.37 for this instance. The prediction moves to 12.66 (applying the local reframing  $r_2$ , using  $\hat{F}^{-1}(0.25)$ ). The bottom row shows a similar picture for a different instance where the expected value is 21.83. The soft regression model gives a standard deviation of 3.85 for this example. The reframing  $R_1$  leads to 17.46 (bottom left) while the reframing  $r_2$  leads to 19.23. As we see, the shift for the first instance (top) is greater for the local method (right), while the shift for the second instance (bottom) is greater for the global method (left).

In regression, approaches to global reframing rely on the calculation of a constant (or function) from the training set [Bansal et al. 2008; Zhao et al. 2011]. For instance, one simple method for global constant reframing is the calculation of the best constant shift for the training set given a loss function. One problem of these methods is that we do not always have the training (or a validation) dataset. Or even if we can keep the dataset, the loss function parameters may be different for each instance and the shift needs to be recalculated by exploring the whole training dataset all over again. Finally, this reframing may be problematic when the output distribution differs between the training dataset and the deployment dataset, because the global function is optimised for the training dataset.

On the contrary, local reframing does not have the above-mentioned problems, but requires, in the probabilistic case, an *accurate* conditional density estimation and a method to derive the *optimal* reframing in an analytical or numerical way. In what follows, we will focus on the notion of ‘optimal reframing’ in the probabilistic (local) case, since the notion of optimality for the non-probabilistic (global) case is more elusive (since reliability or confidence measures cannot be used, in general, to quantify the expected loss).

### 3.2. Optimal probabilistic reframing

It seems reasonable to think that better decisions can be made if we take the conditional density function  $\hat{f}(y|x)$  into account, rather than just the expected value  $\mathbb{E}_{\hat{f}}(y|x)$ , provided this density function is well-estimated.

The maximum density reframing is given by  $r^{max}(x, \ell, \hat{f}) = \arg \max_y \hat{f}(y|x)$ , which ignores the loss function, and just gives the point with maximum density. The *mean* (or void) reframing is given by  $r^{mean}(x, \ell, \hat{f}) = \mathbb{E}_{\hat{f}}(y|x) = \int_{-\infty}^{\infty} y \hat{f}(y|x) dy = \hat{\mu}_{\hat{f}}(x)$  which also ignores the loss function. For some density functions, e.g., a normal distribution, we have  $r^{max} = r^{mean}$ .

In general, we want to take the loss function  $\ell$  into account. For an unlabelled instance  $\langle x, ? \rangle$ , the expected loss (risk function) for prediction  $t$  is given by:

$$\mathcal{L}(x, t, \hat{f}, \ell) = \int_{-\infty}^{\infty} \ell(t, y) \hat{f}(y|x) dy \quad (2)$$

Then, the prediction with minimum expected loss is calculated by the following reframing transformation:

$$r^*(x, \ell, \hat{f}) \triangleq \arg \min_t \mathcal{L}(x, t, \hat{f}, \ell) = \arg \min_t \int_{-\infty}^{\infty} \ell(t, y) \hat{f}(y|x) dy \quad (3)$$

This equation says that the optimal prediction (ignoring the uncertainty of the estimation of  $\hat{f}$ ) for each example  $x$  depends on its estimated distribution and the loss function. Interestingly, the previous equation for  $r^*$  is independent of the data (marginal) distribution  $f(x)$ , which means that it can be applied to each individual instance without considering the rest. This is important, since the loss function  $\ell$  may even vary for different instances.

The question, now, is how to *solve* eq. (3). In some cases, if  $\ell$  and  $\hat{f}$  follow some properties, we can easily solve the equation. For instance, the following proposition gives the result for the easiest (well-known) case (the proof is in appendix G):

**PROPOSITION 3.3.** *If  $\hat{f}$  is symmetric<sup>1</sup> and  $\ell$  is symmetric and commutative then*

$$r^*(x, \ell, \hat{f}) = \hat{\mu}_{\hat{f}}(x) = \int_{-\infty}^{\infty} y \hat{f}(y|x) dy = r^{mean}(x, \ell, \hat{f})$$

*which states that the optimal prediction is given by the mean of the conditional density function.*

But in many applications,  $\ell$  is not symmetric. Also, for many density estimation methods,  $\hat{f}$  is not symmetric either. Only if  $\hat{f}$  is chosen to be a simple distribution (e.g., a normal distribution), the equation can be solved analytically for specific asymmetric loss functions, as we will see in the following sections. In the general case, however, the best prediction cannot be calculated in an analytical way and needs to be obtained by a numerical method, such as a Monte Carlo method, or any other method (e.g., hill-climbing) which can exploit the properties of particular cases for  $\ell$  and  $\hat{f}$ , such as (partial) monotonicity or convexity.

## 4. METHODOLOGY AND EXPERIMENTAL DESIGN

Once we have the ingredients, terminology and concepts, we can state our research goal more properly. Namely, in the rest of this paper, we will answer several questions. First, are there general and practical conditional density estimation methods which can be used effectively to convert an existing crisp regression model into a soft model suitable for reframing? In order to answer this question we need to explore techniques which can produce conditional density estimations for any crisp regression technique. We will focus on normal density estimations, because only two parameters are needed, mean and variance, and the former is already given by any regression model. This implies

<sup>1</sup>The notion of symmetry for the loss function has been defined above. The notion of symmetry for a distribution is the classical notion of symmetry relative to the mean.

that we will be able to derive the reframing transformation relatively smoothly for most loss functions. Consequently, the following section will be devoted to the experimental analysis of several old and new approaches to normal conditional density estimations. From this analysis we will be able to select some methods that will be used in subsequent sections.

The second, ultimate, question of this paper is whether probabilistic reframing methods based on these simple estimators are able to solve a broad set of cost-sensitive problem families, including bidding problems, asymmetric loss functions and rejection rules. We will devote a section to each of these families, we will derive the formal expressions for the reframing transformations and we will compare these methods with other previous *specific* methods in the literature addressing each of these problem families. We will see that probabilistic reframing is more general and effective.

#### 4.1. The true and estimated conditional densities. Reliability and calibration

Before addressing the above issues we need to better understand what kind of conditional density function we expect to find. If the target function is stochastic and/or the data is noisy then the true conditional function  $f(y|x)$  may follow any continuous distribution. The ideal solution would be to find  $\hat{f}(y|x) = f(y|x)$ . Clearly, if  $f(y|x)$  is not normal, there will always be some mismatch if we restrict  $\hat{f}$  to be normally distributed, as we do here. However, if the target function is deterministic, the true conditional function is a Dirac delta function. In this case, the idealistic solution would be to have a *perfect* model, with no variance. But for real situations this will produce no overlap between  $f(y|x)$  and  $\hat{f}(y|x)$ .

In practice, we do not know whether the true function is deterministic or not. Nonetheless, we usually work with datasets for which it is rarely the case that two examples with the same values for the independent variables  $x$  appear with two different values of  $y$ , or, if the dimensionality is high, the data may be so sparse that the true conditional density becomes more elusive. This means that  $f(y|x)$  is either unknown or a Dirac delta function, so trying to make  $f(y|x)$  and  $\hat{f}(y|x)$  match is not very useful. As a result, when understanding a soft regression model, we need to look for other interpretations of  $f$  and  $\hat{f}$ .

Regression models are usually trained (and evaluated) to minimise their mean squared error (*MSE*). This average of the squared residuals  $(y - \hat{y})^2$  is usually analysed in terms of the classical (see, e.g., [Hastie et al. 2009; Flach 2012]) bias-variance decomposition of *one* example for *all* possible datasets  $D$ :

$$\begin{aligned} \mathbb{E}_{\{D\}}[(y - \hat{y})^2] &= (\mathbb{E}_{\{D\}}[\hat{y}] - y)^2 + \mathbb{E}_{\{D\}}[(\hat{y} - \mathbb{E}_{\{D\}}[\hat{y}])^2] = (\mathbb{E}_{\{D\}}[\hat{y} - y])^2 + \mathbb{E}_{\{D\}}[(\hat{y} - \mathbb{E}_{\{D\}}[\hat{y}])^2] \\ &= (\text{Bias}_{\{D\}}(\hat{y}, y))^2 + \text{Var}_{\{D\}}(\hat{y}) \end{aligned}$$

where  $\mathbb{E}_{\{D\}}$  denotes the expected value for all possible datasets. The above decomposition is usually a way to understanding overfitting as high variance: the *predictions* vary very significantly when we change the dataset, i.e., when we move from training to test. Both components of the decomposition contribute to a mismatch between the true conditional density and the estimated density.

Nevertheless, this means that the squared error (i.e., the squared residual), if well estimated for the example, may serve as an approximation of the variance, especially if there is no systematic bias. This is the rationale under some of the methods of variance estimation that we will see in the following section, which interpret the variance of  $\hat{f}(y|x)$  as an estimation of the squared residuals. In other words, the conditional variance is closely related to the instance-specific *reliability* of the regression model.

Conditional density functions are usually more meaningful when  $\mathbb{X}$  has very few dimensions and we find dense areas from which  $\hat{f}$  can be well estimated. When the dimensionality is high, a possible approach for estimating a conditional variance or density (instead of considering multiple datasets) is to calculate partial conditional densities with subsets of the independent variables, for instance, by calculating  $\hat{f}(y|x_2 = a)$ , where  $x = \langle x_1, x_2, x_3 \rangle$ . Even in the deterministic case this will produce non-delta distributions, as  $y$  is not completely determined by subsets of  $x$ .

This leads to the most common approach: to estimate the conditional density (or just the conditional variance) by considering several examples, using a kernel or bin. By grouping several (similar) examples, we can estimate a variance from them, as done by several calibration methods in classification, such as binning. This will also be a usual approach in some of the methods in the following section.

It is important to note that low bias and high variance may be a sign of overfitting (of the crisp model) but does not say much about the overfitting of the conditional density estimation. Some models may have low bias and high variance and we may still be able to find good conditional density estimations that are able to estimate this high variance well. In other words, there are two aspects where we can have overfitting, as a crisp regression model (seen in the decomposition above) or as a soft regression model.

This suggests that the view of  $\hat{f}$  in terms of *calibration* may be more clarifying. In classification, we say that a model is calibrated [Cohen and Goldszmidt 2004; Bella et al. 2009] if the following property holds: given an estimated probability  $\hat{Pr}(c|x) = p$  for class  $c$  then the true class is  $c$  in a proportion of  $p$  cases. In order to apply this definition to deterministic functions, bins of equal or similar estimated probability are used to derive the proportion. Following the same rationale, the calibration of a soft regression model is the degree of concordance between the (aggregated) estimated density function of a group of examples with the (aggregated) true conditional density function of the group. Again, this rationale leads to the approaches based on bins of examples or kernels.

#### 4.2. Evaluation metrics for conditional density estimators

As discussed above, in practice, we do not have access to the true conditional density function. This means that we cannot compare  $f$  and  $\hat{f}$  directly and need to find other ways of evaluating the quality of  $\hat{f}$ . As we are considering a normal distribution for  $\hat{f}$ , we can distinguish three kinds of metrics, depending on what they can evaluate: (1) how good the conditional mean is, (2) how good the conditional variance is, and (3) how good the conditional density is (which is given by the qualities of the mean and the variance). We will introduce them in the following order: (1), (3) and (2).

Since we need to work with any possible base regression model, we present general measures instead of technique-specific measures for particular goodness-of-fit, parameter estimation or intrinsic variance estimation. In order to make results more commensurate and easier to compare, for all the measures which are not in the interval  $[0, 1]$ , we will apply the logistic function  $\Lambda(t) \triangleq \frac{1}{1+e^{-t}}$ . We will use the word ‘standardised’ to refer to this logistic normalisation. In some cases, we will apply the function  $1-t$  or other transformations to always get a decreasing  $[0, 1]$  scale (0 for very good estimations and 1 for very bad estimations).

Let us start with the evaluation of the conditional mean, which is easy as we do have the true conditional mean  $y$ . The evaluation of the conditional mean or expected value is usually measured by the mean squared error  $\frac{1}{|D|} \sum_{(x,y) \in D} (y - \hat{m}(x))^2$  over a dataset  $D$ , although other metrics are also common, such as the mean absolute error, several correlation indices, mean relative squared error, etc. We will use the *mean relative squared error* (*mrse*) to make the measure less dependent on the dataset and easy to compare with the constant (trivial) regression model (the model which always outputs the mean of the training dataset,  $\mu(D_y)$ ):

$$mrse(\hat{f}, D) \triangleq 2\Lambda \left( \frac{\sum_{(x,y) \in D} (y - \hat{m}(x))^2}{\sum_{(x,y) \in D} (y - \mu(D_y))^2} \ln 3 \right) - 1 \quad (4)$$

The factor  $\ln 3$  and the linear transformation makes that we get 0.5 if the error is the same as the constant (trivial) model, 0 for a perfect regression model and close to 1 for very bad estimations.

A more complex issue is to evaluate the quality of the conditional density. As mentioned above, we cannot compare it to the true conditional density, e.g., using the mean squared error for the

distributions, i.e.  $f(f(y|x) - \hat{f}(y|x))^2$ . This also happens for other distribution divergences, such as the KL-divergence. Alternatively, one common measure that does not rely on  $f(y|x)$  is the mean negative log-likelihood (*nll*), which is defined as  $\frac{1}{|D|} \sum_{\langle x,y \rangle \in D} -\ln(\hat{f}(y|x))$ . We will again use the logistic function, applied to the log-likelihood, i.e.:  $v \triangleq \Lambda(\ln(\hat{f}(y|x))) = \frac{1}{1+e^{-\ln(\hat{f}(y|x))}} = \frac{1}{1+\frac{1}{\hat{f}(y|x)}}$ . From here, we just switch to  $1-v$  to get 0 for very good estimations and 1 for very bad estimations, and derive the mean standardised likelihood (*mll*), as follows.

$$mll(\hat{f}, D) \triangleq \frac{1}{|D|} \sum_{\langle x,y \rangle \in D} 1 - \frac{1}{1 + \frac{1}{\hat{f}(y|x)}} = \frac{1}{|D|} \sum_{\langle x,y \rangle \in D} \frac{1}{1 + \hat{f}(y|x)} \quad (5)$$

The log-likelihood (or its logistic variant) evaluates, at the same time, the quality of the mean and the quality of the variance. If we want a measure of the latter only, a possibility might be the squared error between the residual and the standard deviation.

However, this measure also depends on how well the means are estimated, since when mean estimations are accurate (respectively inaccurate) the residuals are low (respectively high), and variances would tend to be low (respectively high) as well. So we are interested in a measure of the quality of the conditional variance as well. Since we do not have access to the true variance, an alternative is to calculate the variance ratio versus the squared residual, as this is a plausible interpretation of the conditional variance. If we denote the residual as  $res_{\hat{f}}(x,y) \triangleq (\hat{\mu}_{\hat{f}}(x) - y)$ , we can define the variance ratio as  $vr_{\hat{f}}(x,y) \triangleq \frac{\hat{\sigma}_{\hat{f}}(x)^2}{res_{\hat{f}}(x,y)^2}$ . The numerator is the estimated variance and the denominator is the squared residual. This ratio will be close to 1 if both quantities are similar. If both the numerator and the denominator are 0,  $vr_{\hat{f}}(x,y) = 1$  by definition. From here, the mean standardised variance ratio is given by the logistic function of the log ratio:

$$msvr(\hat{f}, D) \triangleq \frac{1}{|D|} \sum_{\langle x,y \rangle \in D} \left| 1 - 2\Lambda(\ln(vr_{\hat{f}}(x,y))) \right| = \frac{1}{|D|} \sum_{\langle x,y \rangle \in D} \left| 1 - 2\frac{1}{1 + vr_{\hat{f}}(x,y)} \right| \quad (6)$$

This measure is always between 0 and 1, with 0 being a perfect variance estimation (variance is always equal to the squared residual) and 1 being the worst variance estimation (variance being much higher or much lower than the squared residual). This last measure is entitled to be used as a metric of calibration, which tries to be independent of how good the conditional means are (possibly inherited from a crisp regression model).

### 4.3. Experimental Design

Now we have some criteria (the metrics above) to assess how good a soft regression model is, circumventing the true density function. They are independent of any loss function, and are used when we do not know (or want to ignore) the operating condition. Actually, this is what we will do in section 5 to choose some appropriate ways of obtaining good soft regression models. Of course, when we finally have (or want to use) a loss function, it is the overall loss which matters, and this is what we will use for evaluation in sections 6, 7 and 8, with several loss functions, using reframing.

Given the metrics and they have to be used, the final piece to be able to perform the evaluations of any of these sections is the definition of an appropriate experimental design for them. We will briefly describe the general experimental setting now, and we will let some other details for each specific section.

We will use forty datasets, as shown in tables XI and XII in appendix A. The first battery of datasets will be used for the experiments in section 5. We will use the other battery for the experiments in sections 6, 7 and 8. The reason for two different batteries is that we select the best conditional density estimation and enrichment methods in section 5 with the first battery, and we use these methods with a fresh and independent battery for the particular applications in the other sections.

In all the experiments we split the datasets without previously shuffling the datasets (i.e. the order is preserved before splitting them). We take 50% for dataset *Train* and 50% for dataset *Test*. This is one fold. Next, we randomly take a sample of 2/3 of the examples in *Train* for training, without replacement, and 2/3 of the examples in *Test* for testing, without replacement. We repeat this process 5 times. Next, we swap *Train* and *Test*, to have a second fold, and we repeat this process 5 times. So, in total, we get 10 train-test samples for each dataset. This configuration tries to mimic a realistic situation where the training and test distributions may differ. Note that shuffling the datasets would yield similar distributions for training and test (in terms for the predictors  $x$  and response  $y$ ), which is a quite uncommon scenario in practice (although relatively usual in data mining and machine learning research experiments). It is important to highlight that it is not our goal to obtain an estimation of how well each method will perform for each dataset under exactly the same distribution (where usual re-sampling methods such as 10-fold cross-validation would be appropriate), but to compare several methods on a realistic setting where the context between training and test can change, including the data distribution. In order to illustrate how training and test distributions differ, two extra columns (TrTeMD, TrTeKS) in the dataset tables **XI** and **XII** show the train-test relative means difference, calculated as  $\frac{|\mu_{Train} - \mu_{Test}|}{\sigma_{Train}}$ , and the train-test Kolmogorov-Smirnov statistic, respectively. Both are averaged for the two folds. The higher these values are the more dissimilar the training and test distributions are.

In order to assess the significance of the experimental results we will use a custom procedure, following [Japkowicz and Shah 2011] and [Flach 2012, ch.12], which in turn is mostly based on [Demšar 2006]. Since we will not have any baseline method, we will use a Friedman test to tell whether the difference between several methods is significant and then we will apply the Nemenyi post-hoc test. We agree with [García and Herrera 2008] that the Nemenyi test is a “very conservative procedure and many of the obvious differences may not be detected”, but we prefer to be conservative given our experimental setting and the use of a 0.95 confidence level. In some result tables we will show the means (even though in many cases they are not commensurate) and in some other tables we will show the average ranks (from which the Friedman and Nemenyi tests are calculated). We will also include the critical difference for the Nemenyi test, so we will be able to simply tell whether the difference between two algorithms is significant if the difference between their average ranks is greater than the critical difference.

## 5. NORMAL CONDITIONAL DENSITY ESTIMATION (NCDE): ENRICHMENT METHODS

A theoretically-optimal decision rule for a conditional density function and a loss function will only work if the conditional density function  $\hat{f}(y|x)$  is accurate. While there are many techniques for conditional density estimation (CDE, see appendix B), they may be inappropriate for cost-sensitive scenarios. First, as they focus on the *whole* conditional distribution, the estimated conditional *mean* given by these complex estimated conditional density functions is usually worse than the conditional mean output by many crisp regression methods. Second, CDE methods are usually slow. Third, the actual conditional density functions are rarely multi-modal (i.e., having several peaks or modes), and even if they are, it is not clear that adjusting many parameters to approximate this multi-modality will finally lead to the choice of a better (or even significantly different) optimal prediction for many loss functions. In the end, as mentioned in the previous section, some problems are deterministic and what we really want is an estimation of the residual (i.e., a reliability estimator) rather than (technically) a conditional density function for the output variable.

Instead of complex (usually non-parametric) CDE methods, one of the simplest, most common, parametric density functions is given by the normal (Gaussian) distribution. Estimating a normal distribution only requires the estimation of two parameters, the mean and the variance. It is important to clarify that the use of a normal conditional density  $\hat{f}(y|x) \sim \mathcal{N}$  does not entail —at all— that the output variable is distributed normally ( $\hat{f}(y) \sim \mathcal{N}$ ). Moreover, the use of an estimated normal conditional density  $\hat{f}(y|x)$  does not even mean that we *assume* that the true conditional density  $f(y|x)$  is normal. In fact, when we have an empirical dataset, we do not have information about the

true conditional distribution; we just have examples for which its actual distribution can be seen as a Dirac delta function. In other words, the use of a normal conditional density function follows practical considerations and can be seen (at most) as a representation of the model’s belief about how its uncertainty is distributed, i.e., a model of the *distribution of residuals*.

Consequently, in this section we will explore and develop normal conditional density estimation methods, or NCDE methods for short. This boils down to a soft regression model that, for every input instance  $x$ , just outputs two parameters:  $\hat{\mu}(x)$  and  $\hat{\sigma}(x)$ . The estimation of  $\hat{\mu}(x)$  is the goal of all (crisp and soft) regression methods. Consequently, we will focus below on the estimation of  $\hat{\sigma}(x)$ , comparing the results of several methods. The goal of this section is not to find the best estimator for  $\hat{\sigma}(x)$  as an isolated problem, but to find simple and general methods that work well when the conditional mean  $\hat{\mu}(x)$  is already given by any crisp regression technique.

### 5.1. Directly estimating the variance from the regression techniques

The first way of obtaining the mean and variance for each prediction is choosing a base regression technique which directly or indirectly is able to provide the variance (or a measure of standard error). In this paper, we will work with three common base regression techniques:

- Linear regression (*LR*): many implementations of linear regression can estimate the standard error for each predicted point,  $se(x)$ . If this is the case, we can just set  $\hat{\sigma}(x) = se(x)$ . The particular *LR* method we will use is ordinary least squares using the function `lm` of R [R Team et al. 2012; Torgo 2010] with default parameters.
- Nearest neighbours (*kNN*): in this case the variance is calculated as the variance of the actual  $y$  values for the  $k$ -closest elements. In particular, we use an unweighted  $k$ -nearest neighbours algorithm using the Euclidean distance (with all the attributes scaled by the function `scale` in R) with  $k = 10$ .
- Regression trees (*Tree*): in this case, one easy way of calculating the variance is to calculate the variance for the actual  $y$  values (in the training set) for each leaf of the tree. Then, for each new prediction on a new dataset, the variance will be given by the variance of the leaf where the example falls. We use the CART algorithm [Breiman et al. 1984] implemented by the function `tree` in the package `tree` in R with its default parameters<sup>2</sup>.

We will use these three *base* techniques throughout the rest of the paper. Table II shows the result for the three methods above using the three evaluation metrics (*mrse*, *msll*, and *msvr*) using their own variance estimation.

Interestingly, we can see that for some datasets one method is better than the rest for the conditional mean (evaluated by *mrse*), while it can be the worst for the conditional variance (evaluated by *msvr*). In general, we see that *LR* gives the worst estimations for the conditional mean (*mrse*) and variance (*msvr*), which then implies bad results for the density estimation (*msll*).

### 5.2. NCDE from conditional density, variance, reliability or confidence estimators

The procedure seen above is based on using the variance derived from the own regression technique. These techniques are crisp, and are not really designed to obtain good conditional variances or densities. Instead, ‘soft’ regression techniques (conditional density estimation, conditional variance estimation, reliability estimation and confidence estimation using conformal prediction) look more appropriate for deriving a *normal* conditional density estimator (*NCDE*) model. The use of these methods for *NCDE* would generally involve that we attempt a related, but different (and sometimes more complex) problem first, and then use some transformation or derivation from the soft model to the *NCDE* model. There is nothing against this, provided the results are good and the procedure does not become extremely difficult or inefficient. Below we see why these two criteria are not

<sup>2</sup>As the default parameter for this package keeps the tree unpruned, this method will lead to models that are significantly different from *LR*.

Table II. Three regression techniques using their own conditional variance estimation methods. Results use the datasets in Table XI, using the experimental methodology and the metrics in section 4.

	LR mrse	LR msll	LR msvr	kNN mrse	kNN msll	kNN msvr	Tree mrse	Tree msll	Tree msvr
1	0.32	0.81	0.63	0.36	0.81	0.48	0.43	0.82	0.57
2	0.58	0.79	0.59	0.46	0.76	0.49	0.25	0.69	0.53
3	0.12	0.54	0.52	0.28	0.68	0.60	0.22	0.70	0.50
4	0.15	0.72	0.55	0.39	0.84	0.56	0.32	0.81	0.55
5	0.10	0.81	0.82	0.11	0.63	0.52	0.12	0.63	0.51
6	0.55	0.74	0.63	0.20	0.65	0.57	0.18	0.65	0.55
7	0.97	0.90	0.54	0.43	0.77	0.57	0.46	0.76	0.52
8	0.14	0.73	0.70	0.14	0.63	0.55	0.15	0.61	0.51
9	0.07	0.98	0.95	0.32	0.98	0.89	0.29	1.00	0.94
10	0.31	0.74	0.55	0.32	0.63	0.52	0.19	0.53	0.50
11	0.19	0.74	0.57	0.18	0.76	0.61	0.14	0.69	0.65
12	0.93	0.89	0.66	0.54	0.78	0.47	0.48	0.71	0.56
13	0.03	0.63	0.60	0.50	0.90	0.61	0.44	0.89	0.64
14	0.46	0.83	0.77	0.36	0.73	0.51	0.23	0.65	0.51
15	0.50	0.73	0.68	0.54	0.73	0.60	0.60	0.73	0.56
16	0.49	0.77	0.58	0.32	0.73	0.55	0.24	0.68	0.58
17	0.96	0.97	0.84	0.43	0.80	0.50	0.37	0.77	0.51
18	0.33	0.75	0.67	0.36	0.76	0.54	0.43	0.77	0.54
19	0.35	0.71	0.60	0.37	0.70	0.52	0.45	0.69	0.56
20	0.60	0.83	0.68	0.20	0.68	0.50	0.27	0.70	0.56
Mean	0.41	0.78	0.66	0.34	0.75	0.56	0.31	0.72	0.57

met. A more detailed exploration is given in the appendices B, C, D and E, which also link to the literature.

Let us first review the most general approach, the direct estimation of a conditional density function  $\hat{f}(y|x)$ . Most conditional density estimation methods are designed to issue a complete model of the distribution, which is usually non-parametric. Appendix B describes this approach and shows how it can be adapted to get a normal conditional density. It also includes some experimental results which show that there is no improvement over the base techniques using their own conditional variance estimation methods. Also, general conditional estimation methods are very inefficient and cannot be used as a post-processing step for a crisp regression technique.

A second approach, conditional variance estimation (CVE), is much closer to our specific goal, and can be used to complement an existing crisp regression model by deriving a second parameter, the conditional variance, in order to make up a soft regression model. In fact, these methods can be understood as a post-processing step, which is applied to the whole training set, constructing a model of the residuals. Conditional variance estimation methods are explored in appendix C, but, again, results do not portray a clear advantage.

We have also explored some other methods based on reliability or confidence. Appendix D explores one of the best reliability estimation methods, CNK, included in a recent survey by Bosnić & Kononenko’s [Bosnić and Kononenko 2008] about reliability measures in regression. Also, it is meant to output an estimation of the standard deviation. However, the results are poor. Nonetheless, in appendix D we introduce a ‘correction’, known as KNC, which is just based on comparing the estimated mean with the closest  $k$  true values in the training (or validation) set. The results for KNC are better, which has spurred us to introduce a univariate version  $uKNC$  that we will see below, as an enrichment method. Finally, we explored conformal prediction in appendix E, which outputs confidence intervals, but the results were not better than the rest.

We will evaluate a selection of these methods at the end of this section.






Fig. 3. Enrichment methods convert a crisp regression model into a soft regression model by just comparing  $y$  with  $\hat{y}$ . The mean of the resulting conditional density function  $\hat{f}(y|x)$  is not altered (the original  $\hat{y}$  is kept). Only the second parameter of a normal distribution (the conditional variance) is added.

### 5.3. NCDE through enrichment methods

One of the problems of the previous methods is that they depend on the whole training set for estimating the conditional variance. This looks natural, since in order to get  $\hat{\sigma}(x)$ , we are supposed to need  $x$ . However, if we have a regression model, we already have  $\hat{y}$ , which actually carries information about the input value  $x$ . The basic idea of an ‘enrichment’ method is to derive  $\hat{\sigma}(x)$  from  $\hat{y}$  instead of deriving it from  $x$ . With the original  $\hat{y}$  and the newly derived  $\hat{\sigma}(x)$  we just have a *NCDE* model. Figure 3 shows this process of converting a crisp model into an enriched soft model. This univariate derivation can be performed in several different ways.

A first option is to estimate the residual  $u = y - \hat{y}$  (as a separate regression problem) and derive  $\hat{\sigma}(x)$  from it. This procedure, which does a univariate regression on the residuals given the outputs, is called *residual-based enrichment*, *RBE*. We detail the *RBE* procedure below:

**Definition 5.1.** Given an existing regression model ( $m_y$ ), a training or validation set  $T$ , and a (test) instance  $x$ , the *residual-based enrichment method (RBE)* is defined as follows:

- (1) Obtain  $\hat{y}_i = m_y(x_i)$  for each example  $\langle x_i, y_i \rangle \in T$ .
- (2) Calculate the residuals:  $u_i \leftarrow (y_i - \hat{y}_i)$ .
- (3) Apply a transformation function  $\theta$  to the residuals:  $v_i \leftarrow \theta(u_i)$ .
- (4) Train a regression model  $m_v$  for the dataset  $V = \{\langle \hat{y}_i, v_i \rangle\}_{i=1 \dots |T|}$ .
- (5) Obtain  $\hat{y} = m_y(x)$  and  $\hat{v} = m_v(\hat{y})$  for the example  $x$  to be predicted (in the test set).

So, for each example  $x$  in the test set, the estimated conditional mean for that example is  $\hat{\mu}(x) = \hat{y}$  and the estimated conditional standard deviation is  $\hat{\sigma}(x) = \theta^{-1}(\hat{v})$ . Note that steps 1 to 4 can be omitted if we just train and keep  $m_v$ .

The procedure is similar to the conditional variance estimation methods shown in appendix C, but we remove the dependency on  $x$  for the residual model. This procedure also resembles some calibration methods in classification. Platt’s method [Platt 1999] applies a univariate function (a sigmoid) to the outputs, in order to calibrate them.

In order to apply the *RBE* method, we only need to choose an appropriate transformation function  $\theta$  for step 3 and a regression technique for step 4. The transformation function  $\theta$  can be used at convenience to ensure that  $\hat{\sigma}(x)$  is always positive or to make an estimation of absolute or squared residuals. Several possibilities exist, but a natural choice is  $\theta(t) = t^2$ , if seen as a variance estimation method [Yu and Jones 2004; Wasserman 2006].

We explore the *RBE* method for the base techniques (*LR*, *kNN* and *Tree*) and two different regression techniques (*kNN* and *Tree*) for implementing step 4 in definition 5.1. Table III shows the

results. We see that the results show improvement for the base techniques that performed worst, with some slightly worse results on those that performed best.

Table III. Results (using the datasets in Table XI) for several base techniques (*LR*, *kNN* and *Tree*) with the residual-based enrichment (*RBE*) methods using *kNN* and *Tree* as regression techniques for implementing step 4 in definition 5.1. All the methods use  $\theta(t) = t^2$ . Results for *mrse* are not shown since they are equal to Table II.

	LR	LR	LR	LR	kNN	kNN	kNN	kNN	Tree	Tree	Tree	Tree
	ENR	ENR	ENR	ENR	ENR	ENR	ENR	ENR	ENR	ENR	ENR	ENR
	kNN	kNN	Tree	Tree	kNN	kNN	Tree	Tree	kNN	kNN	Tree	Tree
	msll	msvr	msll	msvr	msll	msvr	msll	msvr	msll	msvr	msll	msvr
1	0.79	0.53	0.79	0.53	0.81	0.52	0.82	0.52	0.84	0.69	0.83	0.63
2	0.76	0.56	0.79	0.54	0.77	0.52	0.78	0.52	0.72	0.61	0.68	0.52
3	0.77	0.82	0.81	0.84	0.68	0.60	0.71	0.61	0.75	0.61	0.70	0.49
4	0.71	0.56	0.73	0.56	0.82	0.47	0.80	0.53	0.86	0.69	0.83	0.60
5	0.63	0.54	0.64	0.55	0.63	0.52	0.63	0.53	0.74	0.67	0.64	0.51
6	0.71	0.61	0.69	0.59	0.67	0.59	0.67	0.59	0.72	0.67	0.64	0.56
7	0.98	0.99	0.96	0.98	0.77	0.55	0.76	0.60	0.78	0.56	0.75	0.53
8	0.64	0.54	0.64	0.54	0.63	0.56	0.63	0.54	0.70	0.67	0.60	0.53
9	0.95	0.91	0.96	0.92	0.99	0.92	0.99	0.93	1.00	0.98	0.99	0.92
10	0.69	0.44	0.71	0.44	0.67	0.53	0.68	0.58	0.69	0.73	0.53	0.56
11	0.71	0.49	0.74	0.57	0.75	0.63	0.76	0.63	0.73	0.74	0.68	0.65
12	0.88	0.64	0.86	0.57	0.77	0.47	0.79	0.50	0.76	0.67	0.73	0.61
13	0.71	0.72	0.66	0.65	0.92	0.66	0.89	0.59	0.89	0.68	0.90	0.62
14	0.74	0.48	0.74	0.47	0.73	0.52	0.73	0.52	0.71	0.63	0.66	0.52
15	0.72	0.59	0.71	0.58	0.74	0.62	0.73	0.60	0.76	0.63	0.73	0.58
16	0.77	0.60	0.79	0.58	0.72	0.54	0.73	0.55	0.73	0.65	0.70	0.57
17	0.96	0.87	0.99	0.93	0.79	0.55	0.79	0.55	0.79	0.64	0.78	0.56
18	0.74	0.55	0.74	0.58	0.75	0.53	0.75	0.56	0.75	0.54	0.76	0.56
19	0.69	0.52	0.70	0.54	0.70	0.51	0.71	0.54	0.74	0.64	0.69	0.55
20	0.79	0.60	0.83	0.61	0.69	0.52	0.69	0.52	0.78	0.68	0.69	0.57
Mean	0.77	0.63	0.77	0.63	0.75	0.57	0.75	0.58	0.77	0.67	0.73	0.58

Given that *enrichment* only requires a univariate regression technique, we can look for simpler and equally effective approaches, without the need of using a second regression technique, such as *kNN* and *Tree*. A single approach is *binning*, which just uses a sliding window over the estimated value  $\hat{y}$ . This approach resembles binning calibration in classification [Bella et al. 2009] More formally, the *BIN* method is defined as follows:

*Definition 5.2.* Given an existing regression model ( $m_y$ ), a training or validation set  $T$ , and a (test) instance  $x$ , the enrichment method *BIN* (using bin size  $k$ ) is defined as follows:

- (1) Obtain  $\hat{y}_i = m_y(x_i)$  for each example  $\langle x_i, y_i \rangle \in T$ .
- (2) Calculate the residuals:  $u_i \leftarrow (y_i - \hat{y}_i)$ .
- (3) Apply a transformation function  $\theta$  to the residuals:  $v_i \leftarrow \theta(u_i)$ .
- (4) Construct a dataset  $V = \{\langle \hat{y}_i, v_i \rangle\}_{i=1 \dots |T|}$ .
- (5) Sort  $V$  by  $\hat{y}_i$ .
- (6) Obtain  $\hat{y} = m_y(x)$  for the example  $x$  to be predicted (in the test set).
- (7) Construct the set  $W$  with the  $k/2$  values  $v_i$  in  $V$  immediately above  $\hat{y}$  and the  $k/2$  values  $v_i$  in  $V$  immediately below<sup>3</sup>.
- (8) Obtain  $\hat{v}$  as the mean of  $W$ .

<sup>3</sup>If there are not sufficient elements above or below we take as many as we can.

The estimated conditional mean is  $\hat{\mu}(x) = \hat{y}$  and the estimated conditional standard deviation is  $\hat{\sigma}(x) = \theta^{-1}(\hat{v})$ . Note that steps 1 to 4 can be omitted (and the training set is no longer necessary) if we just keep the dataset  $V$  when training the model.

A third enrichment method can be defined by constructing the bins using distances. This method builds a bin for each prediction by including the nearest estimated values and then averages the squared differences of the *true* values of these elements against the same prediction (instead of averaging the residuals of each pair of prediction and true value). This method is a univariate version of the method *KNC* (see appendix D):

*Definition 5.3.* Given an existing regression model ( $m_y$ ), a train or validation set  $T$ , and a (test) instance  $x$ , the *univariate k-nearest comparison* enrichment method *uKNC* (using parameter  $k$ ) is defined as follows:

- (1) Obtain  $\hat{y}_i = m_y(x_i)$  for each example  $\langle x_i, y_i \rangle \in T$ .
- (2) Construct a dataset  $Q = \{\langle \hat{y}_i, y_i \rangle\}_{i=1..|T|}$ .
- (3) Obtain  $\hat{y} = m_y(x)$  for the example  $x$  to be predicted (in the test set).
- (4) Let  $S = \langle \hat{y}_j, y_j \rangle$  the set of the  $k$  nearest neighbours in  $Q$  (using the distance  $|\hat{y}_i - \hat{y}|$  between each  $\hat{y}_i$  in  $Q$  and the fixed  $\hat{y}$ ).
- (5) Obtain  $\hat{\sigma}^2 = \frac{1}{k} \sum_{\langle \hat{y}_j, y_j \rangle \in S} (\hat{y} - y_j)^2$ .

The estimated conditional mean is  $\hat{\mu}(x) = \hat{y}$  and the estimated conditional variance is  $\hat{\sigma}(x)^2 = \hat{\sigma}^2$ .

Note that this method is different from *RBE* using *kNN*. The method *uKNC* just looks for the closest estimations in the training set to the estimation for example  $x$  and compares their true values with the estimation for  $x$ . This approach resembles the non-monotonic calibration method (similarity-binning) introduced in [Bella et al. 2013]. Note that the rationale behind this method is that we link the variance to the estimations, i.e., given a set of  $k$  examples with similar estimations, we calculate how far (on average) the true values are to the centre estimation. By using the centre estimation and not the estimation for each of the  $k$  examples with most similar estimations, this method can be more robust (since an outlier estimation for one of the  $k$  estimations has no effect on the result).

Again, we apply these two methods (*BIN* and *uKNC*) to the base techniques (*LR*, *kNN* and *Tree*). Table IV shows the results. As we see, the performance is not degraded at all by these extremely straightforward and efficient methods. Much on the contrary, their results are good, especially for *uKNC*.

#### 5.4. Choosing some appropriate *NCDE* methods for cost-sensitive applications

As we have seen, the number of possible methods which can be used to derive a simple (i.e., normal) conditional density function is really large, and some of them could be parameterised and refined. Nonetheless, our goal was to select a small set of simple *NCDE* methods that could produce a reasonably good normal conditional density estimation or, more precisely, a good pair of conditional mean and conditional variance (from which a normal conditional density estimation is built).

In order to make a selection, we have analysed some of the methods seen so far in order to find a small subset of methods with the following criteria: good performance (for any base regression technique), low dependence on the training set and efficiency. Performance results and significance tests are shown in tables XVII, XVIII and XIX in appendix F. These tables also include the results for some of the methods mentioned in section 5.2 (which are explained in full detail in appendices B, C, D and E). According to these results and the previous criteria, we decide to use the following *NCDE* methods:

- *Own*: uses the *own* variance estimation methods from each base regression technique (section 5.1).
- *uKNC*: uses the *univariate k-nearest comparison* enrichment method (definition 5.3, section 5.3).
- *BIN*: uses the residual-based enrichment method using *binning* (definition 5.2, section 5.3).

Table IV. Results (using the datasets in Table XI) for several base methods (*LR*, *kNN* and *Tree*) with the enrichment method *uKNC* and the enrichment method using binning for the residuals (*BIN*) (with  $k = 10$  for both enrichment methods). Method *BIN* uses  $\theta(t) = t^2$ . Results for *mrse* are not shown since they are equal to Table II.

	LR ENR	LR ENR	kNN ENR	kNN ENR	Tree ENR	Tree ENR	LR ENR	LR ENR	kNN ENR	kNN ENR	Tree ENR	Tree ENR
	uKNC msll	uKNC msvr	uKNC msll	uKNC msvr	uKNC msll	uKNC msvr	BIN msll	BIN msvr	BIN msll	BIN msvr	BIN msll	BIN msvr
1	0.79	0.51	0.81	0.50	0.82	0.62	0.79	0.52	0.80	0.49	0.82	0.60
2	0.79	0.52	0.77	0.54	0.70	0.51	0.79	0.55	0.77	0.52	0.67	0.52
3	0.67	0.81	0.72	0.56	0.72	0.52	0.77	0.83	0.69	0.64	0.70	0.47
4	0.78	0.68	0.84	0.58	0.83	0.59	0.73	0.58	0.84	0.52	0.81	0.58
5	0.63	0.54	0.63	0.53	0.64	0.53	0.63	0.54	0.64	0.51	0.63	0.54
6	0.70	0.60	0.66	0.60	0.66	0.57	0.71	0.60	0.68	0.59	0.65	0.60
7	0.85	0.42	0.79	0.60	0.75	0.50	0.97	0.98	0.78	0.59	0.75	0.50
8	0.65	0.52	0.63	0.54	0.60	0.52	0.64	0.54	0.63	0.54	0.59	0.53
9	0.75	0.68	0.99	0.90	0.92	0.72	0.95	0.92	0.99	0.89	0.99	0.93
10	0.69	0.36	0.64	0.49	0.52	0.48	0.71	0.40	0.66	0.52	0.54	0.56
11	0.73	0.52	0.75	0.60	0.73	0.72	0.71	0.49	0.75	0.62	0.67	0.64
12	0.86	0.44	0.77	0.49	0.73	0.61	0.86	0.55	0.77	0.44	0.76	0.65
13	0.78	0.80	0.90	0.66	0.89	0.54	0.69	0.63	0.89	0.60	0.88	0.66
14	0.75	0.46	0.73	0.50	0.67	0.52	0.75	0.46	0.73	0.51	0.66	0.53
15	0.72	0.60	0.74	0.60	0.73	0.55	0.73	0.59	0.75	0.62	0.73	0.56
16	0.75	0.49	0.74	0.53	0.69	0.58	0.77	0.56	0.72	0.58	0.69	0.60
17	0.90	0.45	0.79	0.53	0.78	0.52	0.97	0.88	0.79	0.54	0.78	0.54
18	0.75	0.55	0.75	0.57	0.76	0.54	0.74	0.55	0.75	0.56	0.75	0.53
19	0.69	0.53	0.70	0.52	0.69	0.55	0.69	0.53	0.70	0.52	0.70	0.56
20	0.80	0.58	0.68	0.51	0.69	0.56	0.81	0.61	0.68	0.52	0.70	0.59
Mean	0.75	0.55	0.75	0.57	0.73	0.56	0.77	0.62	0.75	0.57	0.72	0.58

Note that we make a selection for the clarity of exposition. If other *NCDE* methods (either directly or by enrichment) are eventually found to perform better or more efficiently (in general or for a particular problem), this will give further support for (and improve) the probabilistic reframing methods that we will explore in the following sections.

## 6. BID APPLICATIONS

As explained in the introduction, most regression problems require the minimisation of a loss function representing the cost context, rather than an uncontextualised (quadratic) error. Frequently, this loss function is only known at deployment time, so training is usually performed without this information, as shown in Figure 1. Given the selection of *NCDE* methods at the end of previous section, we are ready to apply reframing to several kinds of cost-sensitive decision problems, where particular families of loss functions are used. For instance, in this section we will explore a family of problems which are very common in econometrics, commerce and retailing applications, where we need to estimate the price (or other quantifiable features) for an offer or bid in the context of a sale, deal or auction [Schapire et al. 2002; Dumas et al. 2005; Kitts and Leblanc 2004; Wellman et al. 2004; Ghani 2005]. One of the most relevant features of the loss functions in these applications is that they are highly discontinuous, since an offer which is much too expensive changes loss dramatically: from the maximum attainable benefit to no benefit at all (the offer is not accepted). This is formalised by the following bid loss function:

*Definition 6.1.* The bid loss  $\ell_{\beta}^B$  is a loss function defined as follows:

$$\begin{aligned}\ell_{\beta}^B(\hat{y}, y) &= -\hat{y} + \beta \quad \text{if } \hat{y} \leq y \\ &= 0 \quad \text{otherwise}\end{aligned}$$

where  $\beta$  represents some kind of *base cost*. If  $\hat{y} > \beta$  then we have negative loss and hence positive profits.




Fig. 4. Two loss functions shown as a scatter plot with actual output values on the x-axis and estimated output values on the y-axis. Costs are shown with contour lines and colours (from benefits to high costs represented with the scale green-yellow-reddish-white). Left: Bid loss function with  $\beta = 3$ . Right: Non-losing bid loss function with  $\beta = 3$ .

For positive values of  $y$  we see this as a *sale*, for which we want to increase margin. For negative values of  $y$ , it is a *purchase*, and we want to reduce margin. Figure 4 (left) shows a representation of this loss function with  $\beta = 3$ .

Given the bid loss and a *NCDE* model  $\hat{f}(y|x)$ , we need to determine the *optimal local reframing*, which gets the lowest expected loss (minimum risk). This can be done as follows:

PROPOSITION 6.2. Given  $\ell_\beta^B$ ,

$$r^*(x, \ell_\beta^B, \hat{f}) = \arg \min_t \{(\beta - t)(1 - \hat{F}(t|x))\}$$

PROOF. From eq. (3), we have that  $r^*(x, \ell_\beta^B, \hat{f})$  can be written as follows:

$$\begin{aligned} r^*(x, \ell_\beta^B, \hat{f}) &= \arg \min_t \int_{-\infty}^{\infty} \ell_\beta^B(t, y) \hat{f}(y|x) dy \\ &= \arg \min_t \left\{ \int_{-\infty}^t 0 + \int_t^{\infty} (-t + \beta) \hat{f}(y|x) dy \right\} \\ &= \arg \min_t \{(\beta - t)(1 - \hat{F}(t|x))\} \end{aligned}$$

□

The previous equation has no closed form in general (and it does not reduce either for the normal distribution). For most distributions (exceptions are fat-tailed distributions, such as the Cauchy distribution), the value of the estimated cumulative distribution function  $\hat{F}(t|x)$  goes to 1 faster than  $t$  grows to infinity. So this expression is 0 for  $t \rightarrow \infty$ . Hence, in general, it only has a minimum. We can find the minimum of the previous function numerically or we can calculate the derivative  $(t - \beta)\hat{f}(t|x) + \hat{F}(t|x) - 1$  and try to find (also numerically) the values which make the expression 0 and see which of them are minima. The first option seems the easiest one, especially for a normal distribution. While the loss function is discontinuous, the expression in proposition 6.2 is not, and some efficient numerical methods can be used.

Apart from the *local* reframing using *NCDE* methods and proposition 6.2 seen above, we will compare with a *global* reframing which just uses the expected value (a crisp regression model) and adds a shift which has been optimised for the training set. The methods are then:




Fig. 5. Left: comparing the bid loss using different methods for dataset *rock* with base technique *kNN*. Right: comparing the non-losing bid loss using different methods for dataset *menarche* with base technique *LR*.

- *None*: No reframing. The prediction (conditional estimated mean) is used as it is.
- *Own*, *uKNC*, *BIN*: Probabilistic (local) reframing using a numerical approximation for the expression  $r^*(x, \ell_\beta^B, \hat{f})$  with a normal distribution, as given by proposition 6.2. The conditional normal density is obtained by three different methods: as derived by the base technique (*Own*) and the enrichment methods *uKNC* and *BIN*.
- *CoSh*: Global reframing using a *constant shift*  $s_0$  for all the predictions:  $R^+(x, \ell_\beta^B, \hat{f}) = \mathbb{E}_{\hat{f}}(y|x) + s_0$ . We use this approach as a reference method, even though we do not know of any previous use of global reframing for a bid loss in the literature. In order to calculate a good shift, we look for the best shift for the whole training set<sup>4</sup>. This calculation of  $s_0$  is done numerically on the training set. Since  $\ell_\beta^B$  is discontinuous, we cannot use any common optimisation method, so we use a grid search algorithm (a Monte Carlo method could also be used), assuming that the solution is inside a (wide) interval. The accuracy of the solution depends on the number of points for which the loss on the training set is calculated.

Now we see how the previous methods perform. We will use several values of  $\beta$  using the equation  $\beta = (\max_y - \min_y) \cdot a^2$ , where  $a$  ranges regularly between 0 and 1, and  $\max_y$  and  $\min_y$  are the maximum and minimum values of the output  $y$  for the whole dataset. The equation tries to capture a range of reasonable cases for this family of problems. The rationale is that high values of  $\beta$  imply that benefits can only be obtained with values of  $\hat{y}$  which get close to  $\max_y$ , while low values of  $\beta$  imply that we will almost always get benefits. This is the reason why we have squared  $a$ , in order to make cases with low  $\beta$  more frequent, if we just choose  $a$  regularly. With this, we explore different reasonable possibilities for  $\beta$ .

Figure 5 (left) shows the evolution of this loss for different methods and different values of  $\beta$  (which is a function of  $a$ ) for *one* dataset as an illustration (the figures may vary significantly for other datasets).

The overall results (Table V) show that an appropriate cost-sensitive probabilistic (local) reframing outperforms a constant shift method (global reframing). The results are consistent for the three different base techniques (*LR*, *kNN* and *Tree*).

<sup>4</sup>Note that this method does not require the estimation of a conditional normal distribution (only the expected value is needed), so any crisp regression method can be used directly. However, it requires the complete training set (or at least the

Table V. Results for the bid loss  $\ell_{\beta}^B$  for the datasets in Table XII, using the experimental methodology in section 4. Each row aggregates the ten samples and ten different values for  $\beta$  per fold using the formula  $\beta = (\max_y - \min_y) \cdot a^2$  with  $a \in \{0, 0.111, 0.222, \dots, 1\}$ . For visibility all the losses are multiplied by 10. Each section of five columns shows results for different base techniques (*LR*, *kNN* and *Tree*). The average ranks (AR) are calculated for these three groups separately. The Friedman statistics for the three sections are (65.97, 66.75 and 66.63 respectively), which are greater than the Critical Value (10.92). This means that the null hypothesis is rejected (significance level: 0.05) and the methods do not perform equally. Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.3626) imply that the difference is significant (in bold).

	LR					kNN					Tree				
	None	Own	uKNC	BIN	CoSh	None	Own	uKNC	BIN	CoSh	None	Own	uKNC	BIN	CoSh
1	3.13	-0.45	-0.77	-0.75	0.25	3.44	-0.35	-0.36	-0.36	0.93	3.89	-0.11	-0.18	-0.36	1.04
2	11.00	0.21	-0.27	-0.18	1.79	5.72	-0.26	-0.29	-0.25	0.11	11.36	-0.02	-0.17	-0.01	1.36
3	6.49	-0.50	-0.53	-0.58	-0.02	4.27	-0.65	-0.63	-0.64	-0.30	8.13	-0.51	-0.59	-0.47	0.22
4	22.10	0.90	0.18	4.07	16.19	9.07	3.78	3.93	3.81	6.86	8.06	3.81	3.87	4.12	6.53
5	17.32	-0.02	-0.90	1.35	13.19	10.72	0.35	0.57	0.32	8.45	9.96	3.89	0.31	1.86	8.03
6	10.03	-0.29	-0.29	0.37	3.83	6.79	-0.36	-0.36	-0.37	0.12	7.97	-0.41	-0.39	-0.34	0.66
7	3.52	1.21	-0.22	-0.20	2.07	5.63	-0.26	-0.29	-0.30	2.15	3.41	-0.14	-0.17	-0.12	1.83
8	7.15	-0.02	-0.33	0.27	5.67	10.93	0.37	0.37	0.35	9.95	11.01	1.52	1.76	1.94	10.35
9	27.33	0.65	0.00	0.71	24.87	5.82	-0.19	-0.20	-0.20	0.36	6.23	-0.16	-0.14	-0.07	0.93
10	3.79	0.11	-0.56	0.17	2.05	5.60	-0.37	-0.33	-0.40	0.77	4.35	-0.51	-0.48	-0.34	0.28
11	9.11	5.58	3.67	5.58	8.63	9.97	5.58	5.58	5.58	9.40	9.95	5.58	5.58	5.58	9.55
12	10.00	-0.51	-0.80	-0.80	0.02	8.23	-0.68	-0.68	-0.69	-0.20	11.30	-0.64	-0.64	-0.64	1.00
13	17.65	0.10	-0.40	-0.41	0.08	17.01	-0.37	-0.39	-0.39	-0.05	17.10	-0.36	-0.35	-0.36	0.79
14	7.67	0.52	-0.27	-0.27	0.94	8.01	-0.26	-0.31	-0.29	0.51	8.76	0.34	0.29	0.09	2.86
15	13.02	-0.19	-0.23	-0.12	3.98	8.62	-0.24	-0.28	-0.29	0.43	11.61	-0.22	-0.19	-0.23	1.19
16	15.91	3.20	-0.40	-0.40	-0.17	18.15	-0.39	-0.38	-0.39	-0.21	17.27	-0.33	-0.27	-0.23	-0.22
17	1.29	0.02	-0.34	0.23	0.99	3.32	0.41	0.50	0.46	1.68	3.10	0.78	0.28	1.05	2.39
18	2.21	1.35	-0.30	-0.40	1.04	2.46	-0.31	-0.27	-0.31	1.14	1.58	0.21	-0.41	0.17	1.27
19	6.86	-0.45	-0.27	-0.51	0.79	9.15	-0.54	-0.57	-0.55	0.30	6.79	-0.52	-0.59	-0.51	1.24
20	13.31	-0.53	-0.47	0.16	8.59	6.71	-0.62	-0.65	-0.60	0.35	6.85	-0.48	-0.50	-0.52	0.87
AR	5.00	2.58	<b>1.35</b>	2.23	3.85	5.00	2.30	2.10	<b>1.60</b>	4.00	5.00	<b>1.85</b>	<b>1.75</b>	2.40	4.00

After this first loss function and its results for different methods, we can of course figure out other related loss functions. For instance, a common variant of the bid loss is when the decision rule does not make a bid if we expect no benefit. This is not a rejection rule, which we will see in section 8, but means that there is no offer, no sale and, hence, no profit or loss. In many applications, this is a more realistic loss function, and can be defined as follows.

*Definition 6.3.* The non-losing bid loss  $\ell_{\beta}^{\bar{B}}$  is a loss function defined as follows:

$$\begin{aligned} \ell_{\beta}^{\bar{B}}(\hat{y}, y) &= -\hat{y} + \beta \quad \text{if } (\hat{y} \leq y) \wedge (\beta \leq \hat{y}) \\ &= 0 \quad \text{otherwise} \end{aligned}$$

Figure 4 (right) shows a representation of this loss function with  $\beta = 3$ . We can get its optimal reframing as we did for  $\ell_{\beta}^B$ :

*PROPOSITION 6.4.* Given  $\ell_{\beta}^{\bar{B}}$ ,

$$r^*(x, \ell_{\beta}^{\bar{B}}, \hat{f}) = \arg \min_t \{(\beta - t)(1 - \hat{F}(\max(\beta, t)|x))\}$$

actual output values  $y$ ) for every new context (loss function). This might not be possible in many applications. It also assumes the same parameters for the loss function for the whole dataset.

PROOF. From eq. (3), we have that  $r^*(x, \ell_{\beta}^{\hat{B}}, \hat{f})$  can be written as follows:

$$\begin{aligned} r^*(x, \ell_{\beta}^{\hat{B}}, \hat{f}) &= \arg \min_t \int_{-\infty}^{\infty} \ell_{\beta}^{\hat{B}}(t, y) \hat{f}(y|x) dy \\ &= \arg \min_t \left\{ \int_{-\infty}^{\max(\beta, t)} 0 + \int_{\max(\beta, t)}^{\infty} (-t + \beta) \hat{f}(y|x) dy \right\} \\ &= \arg \min_t \{ (\beta - t)(1 - \hat{F}(\max(\beta, t)|x)) \} \end{aligned} \quad (7)$$

□

Figure 5 (right) shows the evolution of this loss for different values of  $\beta$  (which is a function of  $a$ ) for *one* dataset. The overall results for this variant are shown in Table VI for several base techniques, with the same configuration as Table V. The results are similar to those in Table V, although the

Table VI. Results for the non-losing bid loss  $\ell_{\beta}^{\hat{B}}$  for the datasets in Table XII, using the experimental methodology in section 4. Each row aggregates the ten samples and ten different values for  $\beta$  per fold using the formula  $\beta = (\max_y - \min_y) \cdot a^2$  with  $a \in \{0, 0.111, 0.222, \dots, 1\}$ . For visibility all the losses are multiplied by 10. Each section of five columns shows results for different base techniques (LR, kNN and Tree). The average ranks (AR) are calculated for these three groups separately. The Friedman statistics for the three sections are (32.43, 36.93 and 38.97 respectively), which are greater than the Critical Value (10.92). This means that the null hypothesis is rejected (significance level: 0.05) and the methods do not perform equally. Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.3626) imply that the difference is significant (in bold).

	LR					kNN					Tree				
	None	Own	uKNC	BIN	CoSh	None	Own	uKNC	BIN	CoSh	None	Own	uKNC	BIN	CoSh
1	-0.59	-0.65	-0.66	-0.69	-0.64	-0.25	-0.38	-0.38	-0.36	-0.47	-0.30	-0.37	-0.40	-0.39	-0.48
2	-0.08	-0.23	-0.32	-0.26	-0.21	-0.08	-0.23	-0.28	-0.28	-0.22	-0.12	-0.22	-0.25	-0.24	-0.15
3	-0.25	-0.46	-0.59	-0.59	-0.52	-0.18	-0.46	-0.45	-0.46	-0.39	-0.37	-0.63	-0.67	-0.62	-0.34
4	-0.00	-0.12	-0.34	-0.02	0.00	0.00	-0.02	-0.01	-0.02	0.00	0.00	-0.01	-0.02	-0.01	0.00
5	-0.33	-0.38	-0.85	-0.30	-0.02	0.00	-0.11	-0.10	-0.12	0.00	0.00	-0.02	-0.14	-0.06	0.00
6	-0.27	-0.23	-0.39	-0.35	-0.17	-0.03	-0.30	-0.30	-0.27	-0.11	-0.08	-0.39	-0.39	-0.35	-0.38
7	-0.11	-0.14	-0.27	-0.26	-0.30	-0.15	-0.27	-0.30	-0.30	-0.35	-0.10	-0.20	-0.19	-0.18	-0.17
8	-0.22	-0.20	-0.58	-0.21	-0.15	0.00	-0.01	-0.01	-0.01	0.00	0.00	-0.00	-0.00	-0.00	0.00
9	-0.10	-0.10	-0.18	-0.09	-0.14	-0.05	-0.17	-0.20	-0.18	-0.20	-0.02	-0.18	-0.18	-0.11	-0.27
10	-0.21	-0.21	-0.58	-0.20	-0.35	-0.03	-0.48	-0.47	-0.49	-0.53	-0.03	-0.32	-0.30	-0.32	-0.22
11	0.00	0.00	-0.09	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
12	-0.57	-0.77	-0.78	-0.79	-0.77	-0.58	-0.66	-0.67	-0.68	-0.61	-0.56	-0.64	-0.64	-0.64	-0.56
13	-0.13	-0.23	-0.39	-0.40	-0.33	-0.08	-0.39	-0.39	-0.39	-0.31	-0.13	-0.36	-0.36	-0.37	-0.18
14	-0.03	-0.15	-0.34	-0.32	-0.46	-0.03	-0.34	-0.35	-0.34	-0.42	-0.04	-0.27	-0.27	-0.24	-0.17
15	-0.35	-0.35	-0.39	-0.33	-0.29	-0.19	-0.33	-0.32	-0.40	-0.27	-0.24	-0.31	-0.35	-0.30	-0.24
16	-0.28	-0.26	-0.38	-0.38	-0.37	-0.31	-0.42	-0.40	-0.39	-0.36	-0.31	-0.37	-0.38	-0.39	-0.32
17	-0.17	-0.18	-0.26	-0.18	-0.16	-0.03	-0.06	-0.05	-0.05	-0.08	-0.05	-0.06	-0.07	-0.05	-0.06
18	-0.20	-0.18	-0.26	-0.34	-0.46	-0.21	-0.31	-0.28	-0.29	-0.42	-0.43	-0.37	-0.42	-0.37	-0.37
19	-0.24	-0.51	-0.32	-0.55	-0.20	-0.12	-0.58	-0.62	-0.62	-0.64	-0.33	-0.52	-0.56	-0.51	-0.50
20	-0.40	-0.60	-0.62	-0.45	-0.42	-0.17	-0.54	-0.57	-0.57	-0.53	-0.36	-0.54	-0.49	-0.54	-0.42
AR	4.08	3.38	<b>1.50</b>	2.58	3.48	4.83	<b>2.60</b>	<b>2.35</b>	<b>2.35</b>	<b>2.88</b>	4.47	2.35	<b>1.75</b>	2.80	3.62

differences between the methods are now smaller since there are many more cases where the loss is 0. Again this shows that the use of an appropriate cost-sensitive probabilistic reframing gets better results than a constant shift method (global reframing). The results are again consistent for different base techniques.

While we only show the results for some typical bid functions as for definitions 6.1 and 6.3, the same idea can be used in applications where there can be more bids (see, e.g., [Bella et al. 2011]), or when auctions work in a different way. In that case, a similar result to propositions 6.2 and 6.4 could be obtained by applying the maximisation recursively. In the cases where the expressions cannot be simplified into a closed form, as in this case, we can use a numerical method.






Fig. 6. Two loss functions shown as a scatter plot with actual output values on the  $x$ -axis and estimated output values on the  $y$ -axis. Costs are shown with contour lines and colours (from benefits to high costs represented with the scale green-yellow-reddish-white). Left: Asymmetric absolute (Lin-Lin) loss function (ABS) with  $\alpha = 0.8$ . Right: Asymmetric squared (Quad-Quad) loss function (SQU) with  $\alpha = 0.8$ .

As mentioned above, the global reframing method cannot be applied when the loss function parameters may be different for each example. For instance, the  $\beta$  parameter of the bid function may depend on the instance, such as cases where this represents the production cost and is different for each product. The probabilistic reframing can still be applied in these cases.

## 7. ASYMMETRIC LOSS APPLICATIONS

Many regression problems do not have a symmetric loss. Depending on the application, overestimations might be worse than underestimations (or vice versa). The way in which this asymmetry is modelled has led to the definition of many asymmetric loss functions, such as *Lin-Exp* (approximately linear on one side and exponential on the other side), *Quad-Exp* (approximately quadratic on one side and exponential on the other side), *Lin-Lin* (asymmetric linear) and *Quad-Quad* (asymmetric quadratic). We will focus on the latter two since these are more common and can be seen as generalisations of absolute error and quadratic error respectively.

First, we give a definition for the asymmetric absolute error  $\ell_\alpha^A$ :

*Definition 7.1.* The asymmetric absolute error  $\ell_\alpha^A$  is a loss function defined as follows:

$$\begin{aligned}\ell_\alpha^A(\hat{y}, y) &= \alpha(y - \hat{y}) & \text{if } \hat{y} < y \\ &= (1 - \alpha)(\hat{y} - y) & \text{otherwise}\end{aligned}$$

with  $\alpha$  being the cost proportion (or asymmetry) between 0 and 1, with increasing values meaning higher cost for *low predictions* (underestimation). In other words, when  $\alpha = 0$  we mean that predictions below the actual value have no cost. When  $\alpha = 1$  we mean that predictions above the actual value have no cost. When  $\alpha = 0.5$  we mean that costs above and below are symmetric.

Similarly, we give the definition for the asymmetric squared error:

*Definition 7.2.* The asymmetric squared error  $\ell_\alpha^S$  is a loss function defined as follows:

$$\begin{aligned}\ell_\alpha^S(\hat{y}, y) &= \alpha(y - \hat{y})^2 & \text{if } \hat{y} < y \\ &= (1 - \alpha)(\hat{y} - y)^2 & \text{otherwise}\end{aligned}$$

Figure 6 (left and right) shows a representation of these two loss functions for  $\alpha = 0.8$ .

Now, we look for the optimal choice in both cases. The case for  $\ell_\alpha^A$  is relatively straightforward:

PROPOSITION 7.3. If  $\ell$  is the asymmetric absolute error function  $\ell_\alpha^A(\hat{y}, y)$  given by definition 7.1 and  $\hat{f}(y|x)$  is any conditional distribution (whose mean is denoted by  $\hat{\mu}(x)$ ), then the expected loss for a predicted value  $t$  is given by:

$$\mathcal{L}(x, t, \hat{f}, \ell_\alpha^A) = \alpha \hat{\mu}(x) + t \hat{F}(t|x) - \alpha t - \int_{-\infty}^t y \hat{f}(y|x) dy$$

where  $\hat{F}$  is the cumulative distribution for  $\hat{f}$ .

From now on, we omit all the proofs, which can be found in appendix G.

The previous expression can be used to obtain the optimal prediction easily:

PROPOSITION 7.4. If  $\ell$  is the asymmetric absolute error function  $\ell_\alpha^A(\hat{y}, y)$  given by definition 7.1 and  $\hat{f}(y|x)$  is any conditional distribution, then the optimal prediction is given by the value  $t$  such that the following equality holds:

$$\hat{F}(t|x) = \alpha \quad (8)$$

where  $\hat{F}$  is the cumulative distribution for  $\hat{f}$ .

Clearly, the previous result can be instantiated for any distribution, whose cumulative distribution is invertible, and get:

$$\hat{F}^{-1}(\alpha|x) \quad (9)$$

where  $\hat{F}^{-1}$  is the inverse of the cumulative distribution for  $\hat{f}$ . If  $\hat{f}$  is a normal distribution then we can just use the quantile function (or probit function).

The expression in eq. 9 is easy and intuitive. For a normal distribution, if we have  $\alpha = 0$  predictions below the actual value have no cost, so the best thing to do is to predict  $-\infty$ , since the quantile function returns this for  $p = 0$ . When  $\alpha = 1$  predictions above the actual value have no cost, so the best thing to do is to predict  $\infty$ . If  $\alpha = 0.5$  the best result is given by the result of the quantile function for 0.5, i.e., the median, which for a normal distribution is also the mean.

As in the previous section, we compare the method without reframing (*None*), with probabilistic (local) reframing and two methods with global reframing.

- *None*: No reframing. The prediction (conditional estimated mean) is used as it is.
- *Own, uKNC, BIN*: Probabilistic (local) reframing using  $r^*(x, \ell_\alpha^A, \hat{f})$ , as given by eq. 9. The conditional normal density is obtained by three different methods: as derived from the base technique (*Own*) and enrichment methods (*uKNC* and *BIN*).
- *CoSh*: Global reframing using a constant shift  $s_0$  for all the predictions  $R^+(x, \ell_\alpha^A, \hat{f}) = \mathbb{E}_{\hat{f}}(y|x) + s_0$ . In order to calculate a good shift, we look for the best shift for the whole training set. Interestingly, in this case, the calculation of the optimal constant shift for the training data follows a convex function (from a convex loss function) and can be calculated using efficient numerical methods. For instance, [Bansal et al. 2008] use hill climbing to calculate this optimal  $s_0$ .
- *PoSh*: Global reframing using a polynomial shift  $s(x)$  for all the predictions  $R^p(x, \ell_\alpha^A, \hat{f}) = s(\mathbb{E}_{\hat{f}}(y|x))$  where  $s$  is a polynomial function. Considering that the problem is convex, [Zhao et al. 2011] present a numerical method (also based on hill climbing) to derive this polynomial in a relatively efficient way. We will just show the results for a first-order polynomial because this degree produced the best results.

Using these methods, the evolution of this loss for different values of  $\alpha$  for *one* dataset is shown in Figure 7 (left). The overall results are shown in Table VII for several base techniques. Here we see that some local and global reframing methods perform relatively well.

Now we will derive the minimisation expression for the asymmetric squared error (definition 7.2):

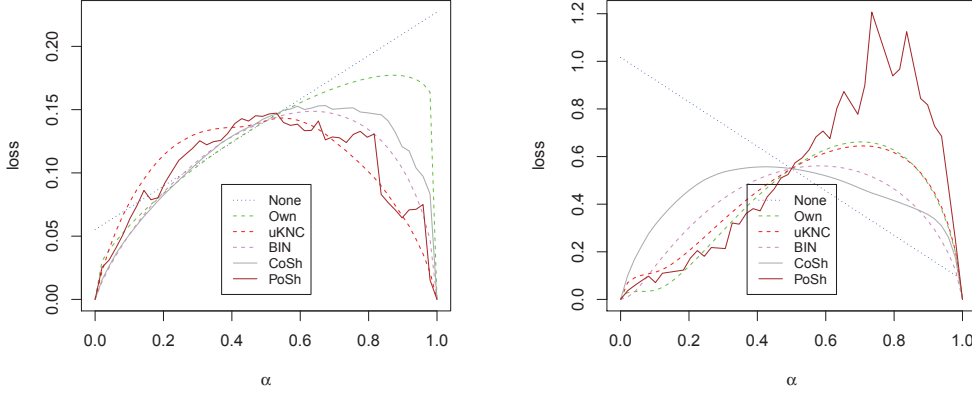


Fig. 7. Left: comparing the asymmetric absolute loss using different methods for dataset *iris3* with base technique *LR*. Right: comparing the asymmetric square loss using different methods for dataset *road* with base technique *Tree*.

Table VII. Results for the absolute loss  $\ell_\alpha^A$  for the datasets in Table XII, using the experimental methodology in section 4. Each row aggregates the ten samples and ten different values for  $\alpha$  per fold with  $\alpha \in \{0, 0.111, 0.222, \dots, 1\}$ . For visibility all the losses are multiplied by 10. Each section of six columns shows results for different base techniques (*LR*, *kNN* and *Tree*). The average ranks (AR) are calculated for these three groups separately. The Friedman statistic for the three sections are (47.46, 44.31 and 49.06 respectively), which are greater than the Critical Value (12.57). This means that the null hypothesis is rejected (significance level: 0.05) and the methods do not perform equally. Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.5217) imply that the difference is significant (in bold).

	LR None	LR Own	LR uKNC	LR BIN	LR CoSh	LR PoSh	kNN None	kNN Own	kNN uKNC	kNN BIN	kNN CoSh	kNN PoSh	Tree None	Tree Own	Tree uKNC	Tree BIN	Tree CoSh	Tree PoSh
1	1.44	0.98	0.99	0.98	0.93	0.91	4.38	3.12	3.10	3.10	3.06	2.90	3.59	2.64	2.57	2.66	2.65	2.51
2	4.66	3.21	2.94	3.18	3.13	3.04	3.54	2.35	2.30	2.36	2.19	2.60	3.68	2.45	2.41	2.43	2.46	2.45
3	2.95	2.00	1.96	1.94	1.97	1.99	3.44	2.18	2.23	2.22	2.24	2.38	3.37	2.15	2.17	2.17	2.17	2.18
4	7.19	4.82	4.88	5.30	5.28	5.11	7.80	5.69	5.56	5.57	5.54	5.53	7.49	5.49	5.44	5.46	5.48	5.38
5	2.92	2.10	2.20	2.19	2.19	2.06	7.12	4.78	4.83	4.73	4.61	4.26	5.03	3.62	3.27	3.59	3.65	3.46
6	6.29	4.27	3.83	4.56	4.57	3.69	3.65	2.41	2.40	2.41	2.39	2.56	3.85	2.58	2.58	2.55	2.50	2.65
7	5.99	4.52	4.23	4.26	4.20	4.00	5.15	3.55	3.51	3.50	3.40	3.30	6.18	4.44	4.45	4.43	4.42	4.20
8	1.86	1.40	1.50	1.46	1.45	1.44	4.34	2.81	2.79	2.89	2.30	2.03	3.99	2.63	2.62	2.70	2.33	2.05
9	9.11	6.35	5.43	7.21	7.19	7.28	3.02	2.21	2.18	2.26	2.27	2.69	3.49	2.37	2.37	2.40	2.39	2.40
10	4.94	3.56	2.85	3.74	3.68	3.55	3.83	2.49	2.48	2.50	2.53	3.02	3.71	2.57	2.56	2.62	2.56	2.72
11	6.57	5.11	4.87	5.08	4.97	4.84	8.29	6.47	6.30	6.28	6.30	6.14	8.16	6.41	6.41	6.41	6.41	6.35
12	1.27	0.94	0.80	0.81	0.80	0.86	1.49	0.98	0.97	0.96	0.97	1.03	2.13	1.39	1.39	1.41	1.39	1.35
13	3.94	2.82	2.54	2.54	2.50	2.50	4.07	2.64	2.61	2.62	2.53	2.47	4.26	2.80	2.82	2.81	2.81	2.71
14	6.19	4.42	4.04	4.16	4.16	4.00	4.77	3.05	3.07	3.07	3.07	3.18	6.23	4.19	4.22	4.24	4.31	4.04
15	2.60	1.76	1.77	1.79	1.81	1.83	2.51	1.71	1.78	1.71	1.61	1.70	2.70	1.72	1.71	1.72	1.69	1.72
16	2.53	1.89	1.60	1.61	1.60	1.68	2.29	1.48	1.49	1.49	1.49	1.56	2.47	1.58	1.58	1.59	1.60	1.61
17	21.14	16.30	11.91	16.67	16.68	15.10	5.09	3.42	3.38	3.42	3.20	2.83	6.53	3.94	3.94	4.20	4.49	3.46
18	7.44	5.64	5.39	5.35	5.41	4.71	5.68	4.10	4.02	3.96	4.14	3.85	6.34	4.84	4.69	4.80	4.88	4.62
19	2.82	1.81	1.86	1.84	1.82	1.89	4.36	2.81	2.79	2.79	2.84	2.75	3.74	2.45	2.44	2.48	2.50	2.45
20	4.36	2.90	2.76	3.12	3.17	3.17	3.62	2.26	2.27	2.29	2.28	2.47	3.33	2.20	2.21	2.21	2.20	2.23
AR	6.00	3.40	<b>2.45</b>	3.40	3.15	<b>2.60</b>	6.00	<b>3.30</b>	<b>2.80</b>	<b>3.25</b>	<b>2.80</b>	<b>2.85</b>	6.00	<b>3.05</b>	<b>2.40</b>	<b>3.60</b>	<b>3.40</b>	<b>2.55</b>

PROPOSITION 7.5. If  $\ell$  is the asymmetric squared error function  $\ell_\alpha^S(\hat{y}, y)$  given by definition 7.2 and  $\hat{f}$  is any distribution with mean  $\hat{\mu}(x)$  and standard deviation  $\hat{\sigma}(x)$ , then the expected loss for a predicted value  $t$  is given by:

$$\mathcal{L}(x, t, \hat{f}, \ell_\alpha^S) = (1 - 2\alpha) \left[ t^2 \hat{F}(t|x) - 2t \int_{-\infty}^t y \hat{f}(y|x) dy + \int_{-\infty}^t y^2 \hat{f}(y|x) dy \right] + \alpha [t^2 - 2t\hat{\mu}(x) + \hat{\mu}_2(x)] \quad (10)$$

where  $\hat{\mu}_2(x)$  is the second raw moment of  $\hat{f}(y|x)$ .

**PROPOSITION 7.6.** *If  $\ell$  is the asymmetric squared error function  $\ell_\alpha^S(\hat{y}, y)$  given by definition 7.2 and  $\hat{f}$  is any distribution with mean  $\hat{\mu}(x)$  and standard deviation  $\hat{\sigma}(x)$ , then the optimal prediction is given by the value  $t$  such that the following equation holds:*

$$(1 - 2\alpha) \left[ 2t\hat{F}(t|x) - 2 \int_{-\infty}^t y\hat{f}(y|x)dy \right] + 2\alpha t - 2\alpha\hat{\mu}(x) = 0 \quad (11)$$

The previous result can be simplified for a normal distribution

**PROPOSITION 7.7.** *If  $\ell$  is the asymmetric squared error function  $\ell_\alpha^S(\hat{y}, y)$  given by definition 7.2 and  $\hat{f}$  is a normal distribution with mean  $\hat{\mu}(x)$  and standard deviation  $\hat{\sigma}(x)$ , then the optimal prediction  $t$  is given by first calculating  $t'$  from the following equation:*

$$t'\Phi(t') + \phi(t') + t' \frac{\alpha}{1 - 2\alpha} = 0 \quad (12)$$

and then getting  $t = \hat{\sigma}(x)t' + \hat{\mu}(x)$ . Note the use of the standardised cumulative normal distribution  $\Phi$  and the standardised normal density function  $\phi$ .

Even though the value of  $t'$  cannot be expressed in a closed form, we only need to calculate this value for each  $\alpha$  once, since it is calculated for the standard normal distribution. Then, we just use the expression  $t = \hat{\sigma}(x)t' + \hat{\mu}(x)$  for each example.

Figure 7 (right) shows the evolution of this loss for different values of  $\alpha$  for *one* dataset, using several reframing methods (again, figures may vary significantly for other datasets). The overall results are shown in Table VIII for several base techniques and the same enrichment methods as in Table VII. The results are good for the probabilistic (local) reframing methods, and especially good for *uKNC*. About the global methods, the results are better for *PoSh* than *CoSh* for two base methods. This indicates that the (more common) asymmetric squared loss, which highly penalises wrong big shifts, may require a more flexible reframing, a polynomial (global) or a probabilistic (local) one.

Apart from  $\ell_\alpha^A$  and  $\ell_\alpha^S$ , there are many other kinds of asymmetric loss. In fact, for instance, we could use a discrete function where loss would be 0 if the error is inside a tolerance band (which can be asymmetric) and, e.g., 1 otherwise. We will discuss this notion of ‘tolerance’ after the following section.

## 8. REJECTION RULE APPLICATIONS

A common situation when working with predictive models appears when there is the possibility of abstention, i.e., to reject the prediction and do nothing (or delegate to an expert or other kind of model). The rationale is to avoid a decision that is likely to have more cost than the abstention itself. In order to do this, we need to know what the cost of an abstention is, which may be constant or may depend on the instance. According to this information, a decision rule which tries to minimise the cost (as an aggregation of the overall prediction loss and the rejection cost) is known as a *rejection rule*. Several works in the literature have been devoted to rejection rules, although most of the work in the area of data mining and machine learning is conceived for classification [Ferri and Hernández-Orallo 2004; Pietraszek 2007].

We can apply a rejection rule on top of any loss function, such as those seen in the previous sections. For instance, for the asymmetric absolute error  $\ell_\alpha^A$  we can derive the corresponding loss with rejection option  $\ell_{\alpha,\rho}^{AR}$  as follows:

Table VIII. Results for the squared loss  $\ell_\alpha^S$  for the datasets in Table XII, using the experimental methodology in section 4. Each row aggregates the ten samples and ten different values for  $\alpha$  per fold with  $\alpha \in \{0, 0.111, 0.222, \dots, 1\}$ . For visibility all the losses are multiplied by 10. Each section of six columns shows results for different base techniques (*LR*, *kNN* and *Tree*). The average ranks (AR) are calculated for these three groups separately. The Friedman statistic for the three sections are (52.2, 44.77 and 51.66 respectively), which are greater than the Critical Value (12.57). This means that the null hypothesis is rejected (significance level: 0.05) and the methods do not perform equally. Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.5217) imply that the difference is significant (in bold).

	LR		LR		LR		kNN		kNN		kNN		Tree		Tree		Tree	
	None	Own	uKNC	BIN	CoSh	PoSh	None	Own	uKNC	BIN	CoSh	PoSh	None	Own	uKNC	BIN	CoSh	PoSh
1	0.60	0.42	0.50	0.42	0.41	0.40	5.79	4.13	4.12	4.12	3.95	3.59	3.45	2.56	2.50	2.56	2.49	2.37
2	9.22	6.55	6.20	6.51	6.54	6.65	6.07	4.37	4.32	4.38	4.28	5.10	6.13	4.39	4.37	4.41	4.37	4.66
3	3.11	2.12	2.09	2.07	2.11	2.20	3.94	2.56	2.61	2.59	2.51	2.67	3.13	2.06	2.05	2.09	2.09	2.14
4	7.33	5.62	5.59	5.63	5.60	5.58	13.46	9.83	9.73	9.73	9.61	9.42	15.08	11.18	11.03	11.01	11.08	10.94
5	2.20	1.60	1.96	1.66	1.65	1.53	12.40	8.69	8.74	8.65	8.87	8.53	7.63	5.80	5.50	5.69	5.72	5.60
6	15.89	10.72	10.01	11.68	11.81	10.55	4.89	3.32	3.32	3.33	3.30	3.63	3.91	2.74	2.72	2.73	2.77	3.13
7	8.83	6.63	6.29	6.29	6.25	6.01	6.57	4.59	4.55	4.54	4.40	4.48	8.66	6.21	6.17	6.27	6.15	5.88
8	5.20	3.81	3.84	4.14	4.14	4.09	10.43	7.99	7.96	8.12	8.46	8.25	9.31	7.02	6.99	7.10	7.54	7.79
9	382.7	322.9	228.1	305.8	305.8	310.2	6.87	5.15	5.14	5.19	5.17	5.35	5.96	4.45	4.43	4.44	4.49	4.71
10	6.48	4.51	4.16	4.84	4.84	4.63	5.82	4.11	4.12	4.10	4.04	4.15	5.28	3.62	3.62	3.68	3.60	3.89
11	10.07	7.87	7.55	7.82	7.73	7.62	14.31	11.13	10.82	10.89	11.17	11.09	14.30	11.23	11.20	11.25	11.21	11.11
12	0.51	0.38	0.34	0.34	0.34	0.37	0.85	0.57	0.57	0.57	0.57	0.59	1.27	0.86	0.84	0.86	0.85	0.84
13	5.41	3.96	3.63	3.63	3.64	3.70	5.76	3.87	3.90	3.89	3.82	3.82	5.95	4.03	4.06	4.05	4.05	3.97
14	12.39	8.95	8.21	8.56	8.57	8.10	6.58	4.51	4.50	4.51	4.46	4.55	10.22	7.03	7.07	7.07	7.11	6.83
15	5.12	3.56	3.49	3.75	3.73	3.70	3.07	2.15	2.18	2.18	2.11	2.03	3.75	2.60	2.52	2.60	2.59	2.61
16	2.43	1.86	1.66	1.68	1.67	1.79	2.76	1.92	1.93	1.94	1.94	2.04	2.92	1.99	2.00	2.02	2.05	1.99
17	148.3	114.2	89.22	117.0	117.2	113.6	10.15	7.16	7.11	7.17	7.49	8.66	16.68	11.05	10.98	11.36	11.94	11.29
18	13.20	9.94	9.49	9.53	9.68	9.14	7.09	5.13	5.05	5.00	5.16	4.64	11.02	8.36	8.16	8.34	8.30	7.85
19	2.45	1.62	1.67	1.61	1.60	1.62	5.74	3.77	3.79	3.77	3.83	3.71	5.51	3.76	3.79	3.80	3.87	3.76
20	14.03	9.66	8.74	10.72	10.71	10.62	4.52	2.96	2.96	2.99	2.96	3.00	3.97	2.71	2.71	2.74	2.76	2.84
AR	6.00	3.70	<b>2.05</b>	3.30	3.25	2.70	6.00	<b>3.05</b>	<b>2.70</b>	<b>3.25</b>	<b>2.70</b>	<b>3.30</b>	6.00	3.15	<b>2.05</b>	3.55	3.50	2.75

*Definition 8.1.* The asymmetric absolute error  $\ell_{\alpha,\rho}^{AR}$  with rejection option is a loss function defined as follows:

$$\begin{aligned} \ell_{\alpha,\rho}^{AR}(\hat{y}, y) &= \rho \quad \text{if } REJECT \\ &= \ell_\alpha^A(\hat{y}, y) \quad \text{otherwise} \end{aligned}$$

A straightforward way of handling this type of loss functions with rejection option is to estimate the expected loss and check whether it is greater than  $\rho$ . If this is the case we should reject. Otherwise, we should use the decision rule as if there were no rejection. For instance, for  $\ell_{\alpha,\rho}^{AR}$ , we would just calculate the expected loss using proposition 7.3, and compare this to  $\rho$ . Only if it is lower than  $\rho$  would we apply the minimisation given by proposition 7.4.

However, we need to derive a more operative expression for the expected loss from proposition 7.3:

*PROPOSITION 8.2.* Consider the asymmetric absolute error function  $\ell_\alpha^A(\hat{y}, y)$  given by definition 7.1 and a normal conditional distribution  $\hat{f}$  with mean  $\hat{\mu}(x)$  and standard deviation  $\hat{\sigma}(x)$ . The expected loss for a prediction value  $t$  can be further simplified to:

$$\mathcal{L}(x, t, \hat{f}, \ell_\alpha^A) = [t' \Phi(t') + \phi(t') - \alpha t'] \hat{\sigma}(x) \quad (13)$$

with  $t' = \frac{t - \hat{\mu}(x)}{\hat{\sigma}(x)}$ . Note the use of the standardised cumulative normal distribution  $\Phi$  and the standardised normal density function  $\phi$ .

The expected loss given by proposition 8.2 is then easy to calculate and can be used to compare it with the cost of rejection. This leads to the decision rule for REJECT:

$$\mathcal{L}(x, t, \hat{f}, \ell_\alpha^A) > \rho \quad (14)$$

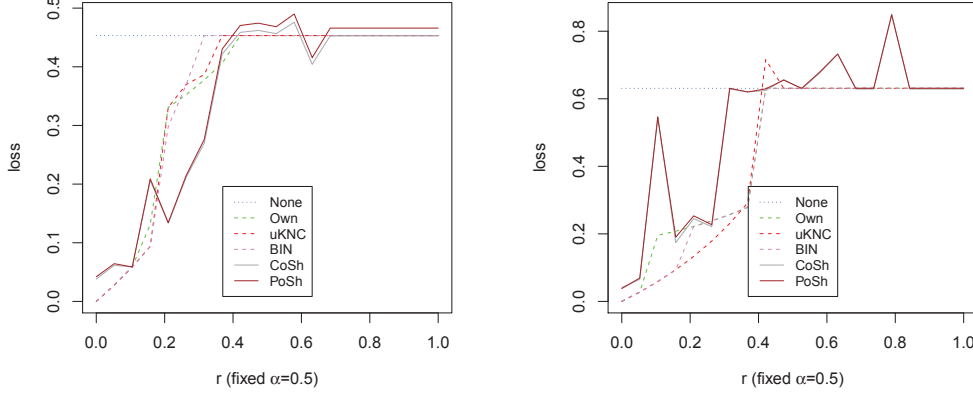


Fig. 8. Left: comparing the absolute loss with reject using different methods for dataset *savings* with base technique *Tree*. Right: comparing the squared loss with reject using different methods for dataset *salinity* with base technique *kNN*.

Now we are ready to compare the methods based on probabilistic (local) reframing using the above rule with methods which do a global reframing, as we did in the previous section. The methods are:

- *None*: No reframing. The prediction (conditional estimated mean) is used as it is.
- *Own*, *uKNC*, *BIN*: We use the rejection rule given by eq. (14) and proposition 8.2. In case of no reject we apply proposition 7.4, eq. (9), as used in the previous section.
- *CoSh*: We decide whether we reject or not using the optimal reject rate calculated on the training dataset. This rate is calculated assuming that a percentage (rate) of examples is just rejected (the examples are necessarily chosen randomly, since there is no information about reliability). Given the optimal rate and the optimal shift, we reject examples using the rate and, if the example is finally not rejected, we use the CoSh method as in the previous section (using the method from [Bansal et al. 2008]).
- *PoSh*: Similar to CoSh but the polynomial approach in [Zhao et al. 2011] is used instead.

For the experiments, we vary both  $\alpha$  and  $\rho$ . For  $\rho$  we apply the following function:

$$\rho = \frac{1}{2} \sigma(D_Y) \frac{r}{1-r}$$

where  $\sigma(D_Y)$  is the standard deviation of the output variable for the dataset  $D$ . We let  $r$  range between 0 and 1. The rationale for the previous function can be explained with the extreme cases. If  $r = 0$  we get  $\rho = 0$  and reject has no cost (so we will always reject). If  $r = 0.5$  we get  $\rho = \frac{1}{2} \sigma(D_Y)$ , which means that with a trivial constant model, residuals will equal the standard deviation (and the expected error), so the expected loss will be  $\frac{1}{2} \sigma(D_Y) = (\alpha \sigma(D_Y) + (1 - \alpha) \sigma(D_Y)) / 2$ . This means that approximately we will reject half of the times (for a trivial model). And finally, for  $r = 1$  we get  $\rho = \infty$  and reject has infinite cost (so we will never reject).

Figure 8 (left) shows the evolution of this loss for different values of  $\rho$ , as derived from  $\rho$  (with  $\alpha$  fixed to 0.5) for *one* dataset. The overall results are shown in Table IX. We see that probabilistic reframing takes advantage of a local decision. In the end, the conditional variance is used to make a *ranking*, which is crucial for rejection rules.

For the squared loss with reject ( $\ell_{\alpha, \rho}^{SR}$ ), we work similarly:

**PROPOSITION 8.3.** Consider the asymmetric squared error function  $\ell_{\alpha}^S(\hat{y}, y)$  given by definition 7.2 and a normal conditional distribution  $\hat{f}$  with mean  $\hat{\mu}(x)$  and standard deviation  $\hat{\sigma}(x)$ . The

Table IX. Results for the absolute reject loss  $\ell_{\alpha,\rho}^{AR}$  for the datasets in Table XII, using the experimental methodology in section 4. Each row aggregates the ten samples and five different values for  $\alpha$  with  $\alpha \in \{0, 0.25, 0.5, 0.75, 1\}$  and ten different values for  $\rho$  with  $r \in \{0, 0.111, 0.222, \dots, 1\}$  and  $\rho = \frac{1}{2}\sigma(D_Y) \frac{r}{1-r}$  (totalling 50 variations per sample). For visibility all the losses are multiplied by 10. Each section of six columns shows results for different base techniques (*LR*, *kNN* and *Tree*). The average ranks (AR) are calculated for these three groups separately. The Friedman statistic for the three sections are (51.8, 47.74 and 53.29 respectively), which are greater than the Critical Value (12.57). This means that the null hypothesis is rejected (significance level: 0.05) and the methods do not perform equally. Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.5217) imply that the difference is significant (in bold).

	LR		LR		LR		kNN		kNN		kNN		Tree		Tree		Tree	
	None	Own	uKNC	BIN	CoSh	PoSh	None	Own	uKNC	BIN	CoSh	PoSh	None	Own	uKNC	BIN	CoSh	PoSh
1	1.24	0.60	0.70	0.59	1.08	1.08	4.27	1.96	1.96	1.93	1.84	1.75	3.68	1.88	1.84	1.89	1.73	1.69
2	4.77	1.93	1.76	1.87	2.07	1.99	3.90	1.57	1.54	1.59	1.58	1.80	3.98	1.63	1.62	1.65	1.82	1.88
3	2.71	1.19	1.14	1.13	1.39	1.42	3.45	1.36	1.39	1.38	1.52	1.50	3.13	1.29	1.27	1.31	1.49	1.50
4	6.05	2.56	2.59	2.84	2.88	2.78	7.55	3.51	3.44	3.39	3.77	3.62	8.04	3.80	3.66	3.72	3.84	3.72
5	2.36	1.12	1.25	1.21	1.44	1.42	7.95	3.20	3.24	3.18	3.39	3.26	4.69	2.33	1.99	2.22	2.29	2.22
6	6.40	1.91	1.70	2.52	2.51	2.08	3.90	1.59	1.59	1.61	1.76	1.81	3.35	1.38	1.38	1.41	1.52	1.55
7	6.14	3.07	2.69	2.74	2.50	2.45	5.06	2.18	2.16	2.15	2.19	2.14	6.18	2.75	2.76	2.79	2.52	2.50
8	1.92	0.91	1.06	1.02	1.09	1.09	4.27	1.97	1.95	2.03	1.83	1.67	4.36	1.89	1.87	1.92	1.82	1.55
9	16.06	2.19	2.07	8.64	7.17	7.15	3.28	1.44	1.42	1.45	1.55	1.69	3.26	1.39	1.38	1.42	1.59	1.58
10	4.85	2.06	1.58	2.37	2.35	2.23	3.70	1.57	1.57	1.58	1.77	1.92	3.96	1.62	1.64	1.67	1.90	2.04
11	6.67	3.54	3.21	3.46	2.98	2.95	8.41	4.35	4.14	4.11	3.97	3.93	8.22	4.40	4.40	4.40	4.04	4.02
12	1.30	0.67	0.58	0.58	0.91	0.93	1.51	0.71	0.69	0.69	0.92	0.92	2.12	0.94	0.92	0.95	1.12	1.11
13	4.00	1.85	1.58	1.57	1.55	1.55	4.03	1.60	1.59	1.58	1.55	1.53	4.55	1.80	1.81	1.80	1.76	1.71
14	6.33	2.72	2.29	2.47	2.38	2.28	4.85	1.89	1.91	1.90	1.95	1.99	5.68	2.30	2.28	2.31	2.23	2.15
15	3.64	1.50	1.31	1.71	1.74	1.72	3.04	1.30	1.31	1.32	1.40	1.41	3.25	1.38	1.30	1.40	1.57	1.52
16	2.43	1.24	0.99	1.01	1.11	1.14	2.33	0.97	0.98	0.98	1.10	1.12	2.37	0.99	0.99	1.00	1.13	1.13
17	21.12	8.81	4.06	10.61	5.56	5.53	5.55	2.23	2.21	2.22	2.20	1.89	7.27	2.55	2.54	2.72	2.80	2.47
18	7.22	3.52	3.13	3.14	2.67	2.46	5.37	2.42	2.36	2.34	2.20	2.13	6.31	3.21	3.09	3.20	2.46	2.43
19	2.73	1.14	1.18	1.16	1.40	1.41	4.06	1.58	1.61	1.59	1.70	1.67	3.92	1.61	1.58	1.63	1.69	1.65
20	5.46	2.00	1.79	2.54	2.44	2.42	3.41	1.36	1.36	1.36	1.54	1.67	3.56	1.45	1.43	1.46	1.63	1.63
AR	6.00	2.95	<b>2.00</b>	3.25	3.70	3.10	6.00	<b>2.65</b>	<b>2.70</b>	<b>2.65</b>	3.60	3.40	6.00	2.85	<b>2.00</b>	3.55	3.70	2.90

expected loss can be expressed as:

$$\mathcal{L}(x, t, \hat{f}, \ell_{\alpha}^S) = \Phi(t') (1 - 2\alpha) [(t'\sigma)^2 + 3t'\sigma^2 q(t') - 2\mu^2 + 4\mu\sigma q(t') - \sigma^2] + \alpha\sigma^2 (t' + 1) \quad (15)$$

with  $t' = \frac{t - \hat{\mu}(x)}{\hat{\sigma}(x)}$ ,  $q(t') = \frac{\phi(t')}{\Phi(t')}$  and notation  $\mu$  for  $\hat{\mu}(x)$  and  $\sigma$  for  $\hat{\sigma}(x)$ .

Although the previous expression is long, it can be computed easily with the standard normal distribution. Figure 8 (right) shows the evolution of this loss for a fixed  $\alpha = 0.5$  and different values of  $r$  (from which  $\rho$  is derived) for *one* dataset. The overall results with the same configuration as the absolute loss with reject are shown in Table X. The results are even more clear-cut in this case.

## 9. DISCUSSION

After this jaunt through several families of loss functions using different kinds of reframing methods (local or global) we are ready to make a comprehensive analysis of the results, see other (many) applications and close the paper with the overall contributions and some future work.

### 9.1. Overview of results and contributions

As a short recapitulation of results, we can just summarise tables V, VI, VII, VIII, IX and X by counting the number of cases where each reframing is in the group of the (statistically significant) best results. This is 18 (out of 18) for probabilistic local reframing in front of 6 (out of 18) for global reframing. If we focus on particular methods, the probabilistic local reframing based on the enrichment method *uKNC* is in the group of best results 17 times (out of 18) in front of only 6 (out of 18) for the best global reframing (*PoSh*, and *CoSh* for the first two tables). In fact, if we compare *uKNC* against *PoSh* (*CoSh* for the first two tables) using the Nemenyi post-hoc test difference in each case, we have 15 *wins*, 3 *ties* and 0 *losses*. As a possible explanation of why and when one approach is better than the other, we can see that global reframing seems to be competitive in asymmetric loss

Table X. Results for the squared reject loss  $\ell_{\alpha,\rho}^{SR}$  for the datasets in Table XII, using the experimental methodology in section 4. Each row aggregates the ten samples and five different values for  $\alpha$  with  $\alpha \in \{0, 0.25, 0.5, 0.75, 1\}$  and ten different values for  $\rho$  with  $r \in \{0, 0.111, 0.222, \dots, 1\}$  and  $\rho = \frac{1}{2} \sigma(D_Y) \frac{r}{1-r}$  (totalling 50 variations per sample). For visibility all the losses are multiplied by 10. Each section of six columns shows results for different base techniques (LR, kNN and Tree). The average ranks (AR) are calculated for these three groups separately. The Friedman statistic for the three sections are (48.49, 56.11 and 47.94 respectively), which are greater than the Critical Value (12.57). This means that the null hypothesis is rejected (significance level: 0.05) and the methods do not perform equally. Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.5217) imply that the difference is significant (in bold).

	LR		LR		LR		kNN		kNN		kNN		Tree		Tree		Tree	
	None	Own	uKNC	BIN	CoSh	PoSh	None	Own	uKNC	BIN	CoSh	PoSh	None	Own	uKNC	BIN	CoSh	PoSh
1	0.50	0.25	0.51	0.25	1.54	1.54	5.39	2.65	2.59	2.65	2.39	2.27	3.85	2.00	1.98	2.00	1.90	1.86
2	8.29	3.50	3.00	3.45	3.51	3.47	5.71	2.51	2.42	2.48	2.57	2.77	5.98	2.68	2.53	2.69	2.73	2.79
3	3.89	1.78	1.63	1.59	2.01	2.04	3.55	1.38	1.46	1.44	1.78	1.88	3.37	1.45	1.42	1.48	1.71	1.72
4	9.65	4.01	4.59	4.69	4.87	4.75	13.17	6.16	5.88	5.93	6.12	6.00	15.23	7.52	7.24	7.21	7.21	7.23
5	2.25	1.13	1.75	1.17	1.95	1.90	12.54	5.50	5.53	5.42	6.02	5.85	8.81	4.61	4.37	4.59	4.66	4.61
6	4.77	1.97	1.96	2.21	2.46	2.62	6.13	2.58	2.56	2.56	2.94	2.99	4.75	2.21	2.19	2.11	2.37	2.42
7	8.63	4.49	4.15	4.21	3.94	3.85	6.43	3.05	2.98	2.95	3.08	3.06	8.96	4.26	4.46	4.50	3.92	3.87
8	3.47	1.83	1.16	1.87	1.96	1.94	5.62	2.86	2.83	2.98	3.14	3.15	8.01	4.22	4.17	4.28	4.43	4.50
9	204.64	13.78	11.53	110.38	94.69	95.73	7.00	3.68	3.66	3.63	3.77	4.04	6.00	3.14	3.15	3.12	3.19	3.30
10	14.34	6.27	2.64	7.59	5.83	5.78	5.60	2.42	2.51	2.47	2.79	2.97	5.70	2.56	2.57	2.62	2.96	3.19
11	10.18	5.42	5.31	5.40	5.01	4.98	14.51	7.62	7.47	7.46	7.28	7.25	14.30	7.64	7.63	7.64	7.26	7.22
12	0.55	0.29	0.27	0.27	1.05	1.05	0.82	0.41	0.40	0.40	1.12	1.13	1.48	0.70	0.69	0.71	1.37	1.37
13	5.20	2.64	2.03	2.03	2.05	2.07	5.88	2.29	2.30	2.31	2.25	2.24	6.06	2.42	2.43	2.43	2.28	2.27
14	13.53	5.98	4.26	5.17	4.59	4.53	7.20	3.05	3.03	3.07	3.21	3.27	11.98	4.78	4.80	5.05	4.24	4.15
15	5.69	1.95	1.71	2.63	2.64	2.62	2.71	1.27	1.24	1.31	1.42	1.43	3.21	1.50	1.37	1.50	1.68	1.68
16	2.60	1.36	1.05	1.11	1.45	1.47	2.40	0.93	1.01	1.05	1.36	1.37	2.53	1.01	1.03	1.08	1.44	1.44
17	161.53	63.71	11.99	84.64	23.92	23.57	7.34	3.70	3.54	3.66	3.70	3.90	15.95	4.33	4.33	4.74	5.07	5.09
18	13.61	6.85	6.17	6.34	4.55	4.44	7.23	3.64	3.58	3.42	3.18	3.03	9.01	4.74	4.66	4.73	3.63	3.57
19	2.82	1.20	1.31	1.30	1.98	1.96	5.38	2.11	2.10	2.11	2.26	2.24	4.92	2.12	2.17	2.28	2.31	2.27
20	12.62	3.21	2.86	6.45	5.17	5.16	4.12	1.66	1.63	1.65	1.88	1.82	3.32	1.45	1.49	1.43	1.75	1.75
AR	5.75	3.05	<b>1.80</b>	3.10	3.90	3.40	5.90	<b>2.85</b>	<b>2.05</b>	<b>2.40</b>	3.85	3.95	6.00	<b>2.65</b>	<b>2.40</b>	<b>3.05</b>	<b>3.45</b>	<b>3.45</b>

applications, as the notion of shift seems to work better here because of the continuous and convexity conditions. In fact, the only three ties take place for the asymmetric absolute loss. This suggests that as the reframing problem becomes more complex because of the non-linearity or non-convexity of the loss function, local reframing seems to be more flexible and powerful. Nonetheless, better specific global (and local) reframing methods could be developed in the future that may yield different results for these or other cost functions. Overall, the experimental results just support the claim that an appropriate probabilistic reframing with the use of a lightweight normal conditional distribution using appropriate estimations (directly or through enrichment methods) is a good approach for a wide variety of cost-sensitive problems.

This claim is accompanied by a series of major contributions:

- We push forward an *appropriate mapping between classification and regression*, and develop the right parallelism between crisp and soft models in classification and regression.
- We vindicate *reframing* as a flexible and powerful approach for context-sensitive applications (as an alternative to re-training), and characterise the distinction between *global* reframing (typically based on crisp regression models) and *local* reframing (typically based on soft regression models).
- We uphold a lightweight view of soft regression models as *normal conditional density estimators*, only requiring two parameters. This entails several benefits: easier, more robust estimation and simpler optimisation formulae resulting from expected loss.
- We introduce *new metrics* for the evaluation of soft regression models.
- We present *enrichment* as a way to convert any traditional one-parameter crisp regression model into a two-parameter soft regression model, by just working with the actual and predicted values. This has several advantages: enrichment methods are easily applicable to any regression method by any data mining practitioner, they only require the actual output values of a training (or vali-



dation) dataset, and the two-parameter normal density estimator does not need to be recalculated whenever the loss function changes.

- We develop new *straightforward enrichment methods which show good performance* as conditional density estimators.
- We show that local reframing has a broad range of applicability over many different problems. We illustrate its *effectiveness on three families of problems* (bid applications, asymmetric loss applications and rejection rule applications) by the theoretical derivation of the expression that leads to optimal reframing in each case and a thorough empirical validation against other approaches, such as global reframing.

These contributions are important for data mining and machine learning because regression, unlike classification, has lacked a comprehensive and effective approach to deal with cost-sensitive problems by the reuse (and not a re-training) of general regression models.

Naturally, there are also some limitations of the approach presented here. For instance, equation 3 is not easily solvable for some loss functions, so we need to rely on the results and expressions previously derived for each loss function (as done here in sections 6, 7 and 8) or use some general numerical (e.g., Monte Carlo) methods. Similarly, decision rules used by reframing may be less understandable than those embedded in a loss-specific model. This is yet another factor to be considered in the re-training/reframing dilemma. Also, even if we choose reframing instead of re-training, we have several options. Global reframing is a more traditional option, which requires a re-evaluation of the shift whenever the loss function changes. Also, global reframing needs to keep the training data in order to perform the process. The advantages of local reframing are that more information is available for each example, so a specific reframing can be done for each example, especially if the cost function changes between examples. The disadvantages come from the need of having a soft regression model and the risk of overfitting, a problem that is minimised by some results of this work, as we have seen very simple methods to enrich a crisp model and convert it into a soft one. Also, note that the conditional mean is not modified, only the variance, so this overfitting can only affect the latter. In fact, overfitting is not specific to local reframing. Some global reframing methods can overfit, especially if a high-degree polynomial is used [Zhao et al. 2011]. Finally, local reframing (using a soft model or a model that has been enriched, once and for all) does not need to keep the training data, while global reframing does.

Finally, there is a case where global reframing may be more flexible than local reframing. On occasions we may have a regression model that was trained using some knowledge about the operating condition or the loss function. This means that the model is specialised for that context. If the deployment context is the same, reframing should not be applied. In fact, it is relatively easy to see that reframing with an asymmetric (or neutral) loss function may give a different (undesirable) result. The question is what to do if the loss function changes. In this case, some global reframing methods may be able to recalculate a new shift (on top of the old shift, for instance) without problems. However, local reframing may be more problematic, as the estimated conditional density is expected to be unbiased and not including any information about the loss. Trying to ‘discount’ or compensate two different loss functions may be possible to do analytically for some loss functions, but it may be tricky in general.

A good way to look at this is in terms of calibration of the soft regression model, as discussed in section 4.1. If we are given a model (or we perform an enrichment) with very bad conditional density estimation (especially in terms of its conditional variance), then it may be better to use a global reframing than a local reframing. In this paper we have focussed on showing that if we are able to get reasonably good conditional density estimations, local reframing is a very good option in general. An experimental analysis of what to do when the conditional density estimations are not good, or to determine when to use one kind of reframing or the other would be a contribution on its own. In the end, we do not claim that the enrichment and reframing approaches are the best solution for any possible cost-sensitive situation, but we claim that it is a feasible alternative that deserves to be considered.

## 9.2. Alternative approaches and other applications

The experimental results shown in previous sections could even be better for probabilistic reframing were we able to get better enrichment methods (or other normal conditional density estimation methods) or other soft regression methods (or adaptations of the ones seen in appendices B, D, B and E), which directly or indirectly can produce good normal conditional density estimations. In fact, we envisage an intensive work in this line, similar to what was done in the last decade for probability estimation in classification, where many classification methods were rethought and re-designed to get good probabilities or good rankings (e.g., probability estimation trees [Provost and Domingos 2003; Ferri et al. 2003; Ferri et al. 2002]). Similarly, an important progress was made in calibration methods [Zadrozny and Elkan 2002; Bella et al. 2009]. This parallel between enrichment and calibration could be exploited to find theoretical foundations and connections with some of the methods inspired or closely related to calibration methods in classification, such as *BIN* or *uKNC*, or to refine them. Other possibilities for soft regression models could be conceived, leading to possibly simpler minimisation solutions (e.g., triangular or uniform distributions). Also, distributions with more parameters (e.g., asymmetric normal, truncated normal or Lévy distributions) could be explored.

The applicability of the enrichment and reframing methods for a diversity of base regression techniques (from non-parametric regression trees and *kNN* to parametric *LR*) with very different mathematical and statistical properties has led us to validate the approaches experimentally. However, the definition of specific combinations of a base technique with a particular enrichment method (e.g., *kNN* with *uKNC*, or *LR* with *ENR-LR*) could be analysed theoretically to derive statistical properties that may characterise their general behaviour better.

One important element in the development and improvement of soft and probabilistic regression models is the use of appropriate evaluation metrics and graphical representations, prior to any specific context-sensitive application, as has been done here in section 5 (before sections 6, 7, 8). We have used *msll* and *msvr* here, but other (new) metrics could be used, inspired by classification plots such as ROC curves, cost curves, calibration plots, etc. The recently introduced ROC curves for regression [Hernández-Orallo 2013], is a possible pathway, also because its area is equivalent to the error variance.

This paper has only included some representative applications, by choosing some common loss functions. There are, of course, many other domains and possible loss functions. For instance, tolerance is a concept that has been frequently used to bring ideas from classification to regression, since a tolerance level can be used to classify estimations as ‘correct’ or ‘incorrect’. An example of a general tolerance loss can be defined as follows, by considering asymmetric losses and asymmetric tolerance levels for overestimations and underestimations:

*Definition 9.1.* The tolerance loss  $\ell_{T,\alpha,\tau^-, \tau^+}$  is a loss function defined as follows:

$$\begin{aligned} \ell_{T,\alpha,\tau^-, \tau^+}(\hat{y}, y) &= \alpha && \text{if } \hat{y} + \tau^- < y \\ &= (1 - \alpha) && \text{if } \hat{y} - \tau^+ > y \\ &= 0 && \text{otherwise} \end{aligned}$$

A related loss could originate from ordinal prediction if we define a loss function in such a way that it is 0 if the prediction is inside the bin of the discretisation (e.g., low (0..3), mid (3..7), high(7..10)), and, say, the number of bins it has to cross to go to the right bin otherwise. It would be interesting to see how probabilistic reframing could work in these two cases.

Apart from the loss functions which relate the true value  $y$  and the estimated value  $\hat{y}$ , there are many other kinds of costs and contexts [Turney 2000]. For instance, loss functions can be instance-dependent, such as those that are a function of the input values, represented as  $\ell(\hat{y}, y, x)$ . It is important to note again that this invalidates global reframing methods (but not local reframing methods). More generally, we can even have a relevance or prior distribution  $U(x)$ . This can be addressed by giving more relevance to some examples than others, in the integration (or sum) of the over-

all estimated cost ( $\sum U(x)\ell(\hat{y}, y)$ ) or in graphical representations, such as ROCIV (instance-varying ROC curves, [Fawcett 2006b]). In other cases, this relevance function can be more complex, as in the so-called utility-based regression [Torgo and Ribeiro 2007; Ribeiro 2011]. One way or another, it is important to realise that the methods in this paper are applicable when there is a change of the prior distribution (the experiments in this paper used five samples over 2-fold cross-validation without shuffling to simulate this situation), since the minimisation of the loss is local to each example. Changing the output distribution (or relevance) does not change the methods, since each optimisation is independent from the rest.

Finally, a soft regression model (issuing probabilities, reliabilities or confidence intervals) can be useful for other data mining tasks, such as quantification [Forman 2008] for regression [Bella et al. 2013], in the same way it has been shown beneficial for quantification for classification [Bella et al. 2010]. Also, screening applications can also take advantage of enrichment methods, since some elements in a rank could be considered a tie if their conditional distributions overlap a certain degree (possibly determining this with a test over the two normals, such as a KS-statistic). This can be applied to preference learning, where we can answer not only whether for two given examples  $x_1$  and  $x_2$  we have that  $\hat{y}_1 > \hat{y}_2$ , but we can also calculate the probability  $Prob(\hat{y}_1 > \hat{y}_2)$  (if the regression model is probabilistic). This, for instance, suggests an evaluation metric related to the Wilcoxon-Mann-Whitney statistic interpretation of the *AUC* (area under the ROC curve), simply as  $Pr(\hat{y}_1 > \hat{y}_2 | y_1 > y_2)$ .

### 9.3. Concluding remarks

The goal of the paper was to show that cost-sensitive applications in regression can be successfully handled by a probabilistic reframing using enriched regression models in the form of a two-parameter normal conditional distribution. In order to accomplish this goal we needed to compare enrichment methods to other approaches for conditional density estimation in terms of estimation quality and efficiency. Another important issue that we needed to consider is the simplicity of the expressions leading to the optimal reframing with minimum expected loss (minimum risk). The choice of a normal distribution consummates all this, and consolidates a view of regression as a two-parameter estimation problem: conditional mean and variance. Also, we have seen that we can *enrich* any existing regression technique with reasonable good variance estimations, using some existing techniques and, most especially, some novel *enrichment* methods that are extremely lightweight. Enrichment methods in regression are somewhat similar to calibration methods in classification (the connections were made explicit in sections 5.2 and 5.3). However, the key difference is that the original prediction is kept and complemented by a second parameter, the variance.

Other approaches for context-sensitive applications build a model which is specialised for a very specific context, embedding the context into the model. In reframing, we reuse a general model for a wide range of contexts and operating contexts. The philosophy is completely different: models can be reused and validated across different operating contexts, improving robustness and efficiency.

Local reframing uses information about each prediction (reliability, confidence or probability) to adapt each local prediction. When we have probabilities (from the use of a conditional density function, e.g., a normal distribution), we can solve the decision rules analytically and, in cases where a closed form cannot be derived, use simple numerical approximations. In fact, probabilistic reframing only needs to derive the conditional variance once and for all, either from the method itself (e.g., regression trees derive this variance as the variance in each leaf of the tree) or by the use of enrichment methods, which only require the comparison to the output value. Once a regression model is equipped with a good conditional variance estimation we can apply the model to a variety of problems. Moreover, we can even use a different loss function (or different loss parameters) for each individual example.

Global reframing, on the other hand, tries to infer one global function from the training set which is applied to all the examples. This implies an optimisation procedure over the whole training set whenever the loss function (or any of its parameters) changes. Also, except for some convex loss functions where some efficient numerical methods can be used, this procedure may be time-

consuming. These differences and the fact that the experimental results are, in general, favourable to probabilistic reframing, suggest that an instance-based (local) approach is a good option for reframing. It also links much better to the areas of risk minimisation in decision theory.

Overall, this paper contains a number of contributions and integrates a wide range of techniques that should trigger further research on conditional variance estimation, enrichment methods, calibration techniques for regression, evaluation metrics for regression, and better reframing techniques on these and other context-sensitive applications of regression models.

## A. DATASETS

Datasets, shown in tables XI and XII, are obtained from eight packages of the CRAN distribution of R-project [R Team et al. 2012; Torgo 2010], namely: ‘class’, ‘boot’, ‘MASS’, ‘nlme’, ‘lattice’, ‘np’, ‘survival’ and ‘faraway’. Some of them have been processed to eliminate redundant or null attributes. On other occasions, the original dataset was not a regression problem or the output attribute was not determined, so we (arbitrarily) chose one of the attributes as the dependent variable. All the datasets and the scripts in R (for the methods and tests) are available at <http://www.dsic.upv.es/~jorallo/reframe-reg/>.

Table XI. Dataset battery used in section 5 and related appendices. We show the size, the number of attributes, the relative difference in means (of the output value) between train and test (TrTeMD) and the Kolmogorov-Smirnov statistic between train and test (TrTeKS).

	name	size	attr	TrTeMD	TrTeKS
1	seatbelts	192	7	1.22	0.53
2	theoph	132	4	0.10	0.11
3	USjudgeratings	44	12	0.24	0.25
4	cars	50	2	1.72	0.74
5	faithful	272	2	0.02	0.07
6	boston	506	14	0.40	0.24
7	UScrime	48	16	0.32	0.19
8	gilgais	364	9	0.20	0.10
9	wtloss	52	2	3.57	1.00
10	cefamandole	84	3	0.37	0.14
11	dialyzer	140	4	0.42	0.47
12	earthquake	182	5	0.31	0.18
13	gasoline	32	6	2.09	0.75
14	glucose	376	4	0.07	0.10
15	IGF	236	3	0.03	0.11
16	nitrendipene	88	4	0.09	0.14
17	wheat	48	4	0.74	0.42
18	environmental	112	4	0.28	0.20
19	wage1	526	21	0.23	0.12
20	ozone	330	10	0.28	0.20

## B. CONDITIONAL DENSITY ESTIMATION METHODS

It can be argued that if we want to obtain a conditional density estimation, we should use conditional density estimation techniques, instead of crisp regression methods. Conditional density estimation techniques [Hyndman et al. 1996] are methods which directly<sup>5</sup> obtain  $\hat{f}(y|x)$ . While this is the most

<sup>5</sup>Some other methods calculate the joint distribution  $\hat{f}(y,x)$  or the likelihood  $\hat{f}(x|y)$ , from which the conditional density is just derived by dividing by  $\hat{f}(x)$  or applying Bayes theorem respectively. However, this is usually more complex than the original problem.

Table XII. Dataset battery used in sections 6, 7 and 8. We show the size, the number of attributes, the relative difference in means (of the output value) between train and test (TrTeMD) and the Kolmogorov-Smirnov statistic between train and test (TrTeKS).

	name	size	attr	TrTeMD	TrTeKS
1	iris3	150	4	2.48	0.70
2	savings	50	5	0.22	0.20
3	USArrests	50	4	0.38	0.20
4	rock	48	4	7.73	0.83
5	trees	32	3	3.37	0.94
6	salinity	28	4	0.16	0.43
7	birthwt	188	10	0.67	0.63
8	menarche	24	3	4.97	1.00
9	road	26	6	0.09	0.46
10	stormer	24	3	0.32	0.25
11	bodyweight	176	4	11.09	1.00
12	oxboys	234	3	0.10	0.09
13	occdpanel	616	7	0.58	0.26
14	lungcancer	168	10	0.67	0.33
15	chicago	48	7	0.34	0.17
16	diabetes	402	3	0.11	0.07
17	divusa	76	7	0.54	0.55
18	exa	256	3	0.40	0.47
19	prostate	96	9	1.33	0.50
20	seatpos	38	9	0.28	0.26

general and informative way for the regression problem (since conditional means, variances, confidence intervals and other measures can be obtained from it), the techniques are usually slower and suffer from a number of restrictions. A general way to tackle this estimation is through non-parametric methods (see, e.g., [Hwang et al. 1994]). For instance, many approaches are restricted to only one (or two input variables, such as R's `hdrcde` package [Hyndman et al. 1996]), or just calculate multivariate densities, which have to be normalised for each input value  $x$  to get a univariate density.

It is not the goal of this paper to evaluate several of the approaches for density estimation methods, but it is important to see whether these methods are better, in general, than 'augmented' or 'enriched' methods for which just a conditional mean and variance are obtained (from which a simple Gaussian density function is estimated). In addition, we are interested in the result of 'reducing' a local and detailed density function into a Gaussian. In order to do all this, we illustrate this approach with a kernel-based (non parametric) conditional density method. We use the function `npcondens` in the R's `np` (non parametric) package. This function computes estimates for the density function (i.e.,  $\hat{f}(y|x)$ ), for a bandwidth specification using the method in [Hall et al. 2004]. From the density we calculate the conditional mean and the conditional variance as a pointwise average which approximates the integral of the expected value and the moment respectively. We used the median point and the four points at  $\pm\sigma$  and  $\pm 2\sigma$ , given by a Gaussian.

Table XIII shows the results of this method. Non-parametric conditional density estimation methods can get good estimations for large datasets with complex densities (e.g., bimodal) but here we see that the results are, in general, worse than those of simple regression methods such as a  $kNN$  or *Tree* for the conditional mean, as shown in Table II. In fact, while the conditional variances seem better (*msvr* value of 0.55 in front of 0.66, 0.56 and 0.57 in Table II), the conditional densities (as measured by *msll*) are not better (except for *LR*) and the squared error (*mrse*) is also worse. A possible idea is then to combine the good conditional means from the base techniques with the conditional variance from the conditional density estimation method. This is what the six last columns show. However, as expected, this does not increase the quality of the conditional densities (as measured by

Table XIII. Results (using the datasets in Table XI) for the kernel-based (non parametric) conditional density method given by the function `npcondens` in the R's `np` (non parametric) package. The first three columns show the results for the parametric density estimation (using the Gaussian approximation). The column "CDE orig msl" shows the *msll* result by using the original non-parametric density function. The six rightmost columns show the results for the mean given by the base method with the variance estimation given by the conditional density method (using the Gaussian approximation). Results for *mrse* are not shown for the six most right columns since they are equal to Table II.

	CDE gaus mrse	CDE gaus msll	CDE gaus msvr	CDE orig msll	LR CDE msll	LR CDE msvr	kNN CDE msll	kNN CDE msvr	Tree CDE msll	Tree CDE msvr
1	0.53	0.84	0.55	0.78	0.83	0.49	0.85	0.55	0.84	0.54
2	0.36	0.73	0.42	0.68	0.74	0.51	0.73	0.51	0.72	0.52
3	0.15	0.61	0.57	0.59	0.61	0.57	0.63	0.63	0.61	0.54
4	0.45	0.86	0.60	0.86	0.85	0.55	0.84	0.57	0.86	0.54
5	0.18	0.68	0.49	0.63	0.68	0.53	0.68	0.51	0.68	0.53
6	0.22	0.68	0.55	0.66	0.67	0.61	0.66	0.59	0.66	0.58
7	0.57	0.77	0.58	0.75	0.79	0.83	0.77	0.57	0.77	0.56
8	0.18	0.64	0.51	0.61	0.63	0.54	0.64	0.55	0.63	0.59
9	0.43	0.98	0.84	0.99	0.97	0.58	0.97	0.81	0.97	0.64
10	0.21	0.58	0.56	0.56	0.58	0.60	0.59	0.64	0.58	0.65
11	0.26	0.69	0.64	0.78	0.70	0.65	0.67	0.67	0.72	0.64
12	0.57	0.75	0.56	0.72	0.75	0.74	0.75	0.59	0.77	0.51
13	0.30	0.82	0.57	0.80	0.83	0.60	0.89	0.74	0.86	0.72
14	0.41	0.74	0.47	0.68	0.74	0.49	0.74	0.51	0.74	0.60
15	0.80	0.80	0.45	0.73	0.81	0.69	0.82	0.67	0.81	0.64
16	0.61	0.80	0.60	0.78	0.81	0.54	0.81	0.55	0.77	0.57
17	0.46	0.80	0.50	0.87	0.81	0.62	0.81	0.59	0.83	0.52
18	0.66	0.82	0.44	0.75	0.82	0.58	0.83	0.56	0.82	0.55
19	0.53	0.74	0.46	0.68	0.76	0.59	0.75	0.56	0.76	0.61
20	0.37	0.73	0.58	0.71	0.77	0.66	0.73	0.54	0.73	0.55
Mean	0.41	0.75	0.55	0.73	0.76	0.60	0.76	0.60	0.76	0.58

*msll*), because conditional variance estimates must be *linked* to a mean estimation. Consequently, neither as a standalone method nor combined with the base classifier can we get better performance. Also, this method is about two orders of magnitude slower than the direct methods in subsection 5.1 (and many other methods that we see in the rest of section 5).

### C. CONDITIONAL VARIANCE ESTIMATION METHODS

Instead of deriving a full conditional density, we can just (re-)use the conditional mean of a classical (crisp) regression method and derive a conditional variance. A usual, but not generally well-known, way of estimating the conditional variance is as follows (steps 1 and 2 can be omitted if we already have a regression model):

*Definition C.1.* Given a training or validation set  $T$ , and a (test) instance  $x$ , the *two-step conditional variance estimation* method 2SCVE is defined as follows:

- (1) Train a regression model  $m_y$  using  $T$ .
- (2) Obtain  $\hat{y}_i \leftarrow m_y(x_i)$  for each example  $\{x_i, y_i\} \in T$ .
- (3) Calculate the residuals:  $u_i \leftarrow (y_i - \hat{y}_i)$ .
- (4) Apply a transformation function  $\theta$ :  $v_i \leftarrow \theta(u_i)$ .
- (5) Train a regression model  $m_v$  for the dataset  $H = \{x_i, v_i\}_{i=1..|T|}$ .
- (6) Obtain  $\hat{y} = m_y(x)$  and  $\hat{v} = m_v(x)$  for the example  $x$  to be predicted (in the test set).

This estimates the conditional mean as  $\hat{\mu}(x) = \hat{y}$  and the conditional standard deviation as  $\hat{\sigma}(x_j) = \theta^{-1}(\hat{v})$ .

Usual choices for  $\theta(t)$  are  $\theta(t) = t^2$  and  $\theta(t) = \ln(t^2)$ , i.e., we model the (logarithm of) the squared residuals [Yu and Jones 2004; Wasserman 2006]. The square is usually included since these methods are aimed at estimating the variance and also because otherwise we would need to remove the sign.

The estimation will depend on the quality of the regression model  $m_y$  and most especially on the second regression model  $m_v$ . In fact, the previous algorithm (from steps 1 to 5) is usually iterated by retraining the regression model  $m_y(x)$  using the heteroscedasticity information about the recently estimated conditional variance. This information can only be used by some regression techniques, e.g., a weighted least squares with inverse variance weights. This is usually called *iteratively re-weighted least squares*.

At this point it is important to notice that we really estimate the variance of the residuals of our model conditional to  $x$ , not the variance of  $y$  conditional to  $x$ .

It is usual to apply a non-parametric model in step 5. For instance, if we use nearest neighbours, the previous algorithm boils down to estimating  $\hat{\sigma}^2(x_i)$  as the mean of the squared residuals of the  $k$ -closest examples, which is similar to what we did for  $kNN$  in section 5.1.

Here we will explore the  $kNN$  and *Tree* techniques for the residual model  $m_v(x)$ , jointly with the three base techniques *LR*,  $kNN$  and *Tree* as usual. Table XIV shows some of these methods for  $\theta(t) = t^2$  (we ran the same experiments with other configurations of  $\theta$  with equal or worse results). The results only show a slight improvement for *LR* compared to the results in Table II. For  $kNN$  or *Tree*, the results are similar to the results in Table II.

Table XIV. Results (using the datasets in Table XI) for several base methods (*LR*,  $kNN$  and *Tree*) with conditional variance estimation using  $kNN$  and *Tree* as models for the residuals. All the methods use  $\theta(t) = t^2$ . The results for *mrse* are not shown since they are equal to Table II.

	LR		LR		kNN		kNN		Tree		Tree	
	CVE	CVE	CVE	CVE	CVE	CVE	CVE	CVE	CVE	CVE	CVE	CVE
	kNN	kNN	Tree	Tree	kNN	kNN	Tree	Tree	kNN	kNN	Tree	Tree
	mssl	msvr	mssl	msvr	mssl	msvr	mssl	msvr	mssl	msvr	mssl	msvr
1	0.79	0.51	0.79	0.54	0.81	0.51	0.82	0.54	0.81	0.62	0.81	0.58
2	0.76	0.52	0.77	0.50	0.76	0.48	0.76	0.46	0.69	0.49	0.72	0.54
3	0.75	0.82	0.79	0.84	0.66	0.55	0.69	0.63	0.70	0.49	0.71	0.48
4	0.72	0.53	0.70	0.55	0.83	0.51	0.82	0.47	0.82	0.63	0.80	0.53
5	0.64	0.54	0.63	0.53	0.64	0.52	0.63	0.53	0.63	0.52	0.63	0.54
6	0.69	0.58	0.70	0.60	0.66	0.58	0.65	0.56	0.64	0.57	0.65	0.56
7	0.98	0.99	0.96	0.97	0.76	0.57	0.76	0.58	0.75	0.54	0.75	0.54
8	0.64	0.52	0.64	0.53	0.61	0.54	0.62	0.55	0.61	0.54	0.61	0.54
9	0.94	0.91	0.95	0.91	1.00	0.93	0.98	0.87	0.99	0.93	1.00	0.94
10	0.71	0.48	0.72	0.45	0.65	0.54	0.67	0.60	0.57	0.55	0.54	0.56
11	0.70	0.47	0.73	0.51	0.74	0.61	0.76	0.64	0.68	0.65	0.68	0.65
12	0.88	0.67	0.88	0.72	0.78	0.44	0.82	0.56	0.74	0.57	0.76	0.63
13	0.66	0.69	0.71	0.64	0.88	0.64	0.91	0.67	0.91	0.68	0.90	0.68
14	0.74	0.49	0.74	0.45	0.73	0.51	0.72	0.48	0.69	0.57	0.67	0.53
15	0.73	0.60	0.73	0.58	0.73	0.60	0.74	0.61	0.74	0.57	0.73	0.56
16	0.78	0.53	0.78	0.60	0.73	0.56	0.73	0.60	0.70	0.60	0.69	0.58
17	0.97	0.87	0.96	0.86	0.78	0.53	0.79	0.50	0.78	0.56	0.77	0.55
18	0.74	0.52	0.74	0.57	0.75	0.55	0.75	0.56	0.75	0.52	0.76	0.55
19	0.70	0.55	0.70	0.54	0.70	0.53	0.71	0.54	0.71	0.57	0.70	0.58
20	0.81	0.64	0.81	0.65	0.69	0.52	0.70	0.55	0.69	0.58	0.70	0.60
Mean	0.77	0.62	0.77	0.63	0.74	0.56	0.75	0.58	0.73	0.59	0.73	0.59

#### D. CONDITIONAL VARIANCE ESTIMATION BASED ON RELIABILITY

A reliability measure for regression is any numerical value which is directly related to the degree of certainty about an accurate prediction being produced or, more precisely, inversely related to the expected (absolute) residual. However, the magnitude of this value can follow any scale. Bosnic & Kononenko [Bosnić and Kononenko 2008] compare several reliability estimators for regression.

Among them, *CNK* is a simple method which shows good performance (as a reliability estimator). This method works as follows. For each example  $\langle x, y \rangle$ , this method just calculates the  $k$ -closest elements in the training set to  $x$  and calculates the mean of their output values, denoted by  $C$  or, in other words, calculates the  $kNN$  prediction for  $x$ . Then it calculates the absolute difference between  $C$  and the prediction (of presumably a regression technique which is not  $kNN$ ). This is the estimated standard deviation.

The previous approach makes an average of the true values, and then compares this to a single estimation. While this is good as a reliability measure, the magnitude of the estimation will typically be low (for a standard deviation), since it just compares the prediction of two methods. Consequently, we suggest a correction, which goes as follows:

- (1) Given an example  $\langle x, y \rangle$ , we estimate  $\hat{y}$  by any base regression technique.
- (2) Let  $S = \langle x_i, y_i \rangle$  the set of the  $k$  nearest neighbours of  $x$  in a training or validation dataset.
- (3) Calculate  $\frac{1}{k} \sum_{\langle x_i, y_i \rangle \in S} (\hat{y} - y_i)^2$  as the output estimated variance.

Since it can be seen as a symmetric version to method *CNK*, we call it *KNC*. The results are shown in Table XV. As we can see, *CNK* does not produce good variance estimations. In addition, it cannot work, by definition, for  $kNN$ , since this method uses  $kNN$  as the true value, and the estimated ‘residuals’ will be 0. This can be clearly seen on the two columns “ $kNN$  *CNK*”. On the contrary, *KNC* works well. In fact, it improves the results for the *LR* base technique shown in Table II, but gets equal or slightly worse results for the rest.

Table XV. Results (using the datasets in Table XI) for several base techniques (*LR*,  $kNN$  and *Tree*) with Bosnic & Kononenko’s *CNK* and a refined variation that we dub *KNC*.

	LR msll	LR msvr	kNN msll	kNN msvr	Tree msll	Tree msvr	LR msll	LR msvr	kNN msll	kNN msvr	Tree msll	Tree msvr
1	0.85	0.70	0.99	0.99	0.84	0.70	0.80	0.52	0.81	0.50	0.82	0.54
2	0.82	0.73	1.00	1.00	0.76	0.63	0.77	0.49	0.76	0.49	0.75	0.61
3	0.61	0.72	1.00	1.00	0.74	0.54	0.68	0.78	0.67	0.60	0.71	0.53
4	0.76	0.65	1.00	1.00	0.88	0.81	0.78	0.72	0.84	0.58	0.81	0.58
5	0.74	0.69	0.99	0.99	0.80	0.78	0.64	0.53	0.63	0.52	0.64	0.53
6	0.72	0.55	1.00	1.00	0.71	0.64	0.69	0.53	0.66	0.59	0.66	0.59
7	0.86	0.35	1.00	1.00	0.77	0.54	0.87	0.37	0.77	0.57	0.79	0.48
8	0.71	0.64	1.00	1.00	0.68	0.63	0.65	0.54	0.63	0.54	0.63	0.60
9	0.74	0.69	1.00	1.00	0.99	0.94	0.75	0.72	0.99	0.90	0.93	0.77
10	0.76	0.61	1.00	1.00	0.63	0.65	0.71	0.40	0.63	0.53	0.64	0.67
11	0.77	0.66	1.00	1.00	0.75	0.75	0.74	0.55	0.74	0.59	0.77	0.76
12	0.86	0.42	1.00	1.00	0.79	0.69	0.87	0.39	0.78	0.46	0.78	0.58
13	0.77	0.84	1.00	1.00	0.91	0.75	0.78	0.87	0.90	0.65	0.87	0.52
14	0.83	0.72	1.00	1.00	0.76	0.69	0.75	0.49	0.73	0.50	0.74	0.64
15	0.77	0.70	1.00	1.00	0.80	0.68	0.74	0.62	0.72	0.60	0.74	0.59
16	0.79	0.65	1.00	1.00	0.76	0.63	0.74	0.47	0.73	0.55	0.72	0.65
17	0.90	0.42	1.00	1.00	0.77	0.62	0.88	0.40	0.79	0.51	0.79	0.55
18	0.81	0.73	1.00	1.00	0.79	0.63	0.75	0.54	0.75	0.54	0.77	0.54
19	0.75	0.68	1.00	1.00	0.74	0.63	0.72	0.58	0.70	0.52	0.72	0.58
20	0.80	0.52	1.00	1.00	0.73	0.62	0.78	0.43	0.69	0.51	0.71	0.54
Mean	0.78	0.63	1.00	1.00	0.78	0.68	0.75	0.55	0.75	0.56	0.75	0.59

## E. CONDITIONAL VARIANCE ESTIMATION USING CONFORMAL PREDICTION

Confidence intervals are an alternative (and statistically convenient) way of measuring the reliability of a prediction. Conformal prediction [Shafer and Vovk 2008; Papadopoulos et al. 2002] is a general technique for deriving confidence intervals. It can be applied to any predictive task, such as classification and regression. While originally introduced for a transductive scenario, it has also been



extended to a more classical inductive setting [Papadopoulos 2008]. Conformal prediction works as follows. Given an error probability  $\varepsilon$  and any regression method that makes a prediction  $\hat{y}$ , it produces a region  $\Gamma^\varepsilon$  such that it contains the true value  $y$  in at least a proportion  $1 - \varepsilon$  of the cases (the confidence level). Logically, by making the region infinitely large, we can get any confidence level. The key issue is that the tightness and therefore usefulness of the prediction region depends on the *nonconformity measure* used. A nonconformity measure is any measure that evaluates how unusual an example is (with respect to the others). In regression, for instance, a typical nonconformity measure is the absolute difference:  $|y - \hat{y}|$ .

The idea of outputting a confidence region is richer and more informative than only a prediction point (which in regression is just the conditional mean). One advantage of confidence regions is that there is no assumption about the conditional distribution. However, this is one of its drawbacks for cost-sensitive learning, because we cannot quantify the probability of error in the prediction<sup>6</sup>.

One way of deriving a density function from an interval is by assuming a distribution. Since we advocate for the normal distribution for context-sensitive applications, we can devise a simple method by assuming this distribution. In particular, for a normal distribution (with cumulative distribution  $\Phi_{\mu, \sigma^2}$ ) we know that a proportion  $p$  of the values inside  $\mu \pm a\sigma$  is given by  $p = \Phi_{\mu, \sigma^2}(\mu + a\sigma) - \Phi_{\mu, \sigma^2}(\mu - a\sigma)$ . For a conformal region  $\Gamma^\varepsilon$  we know that a proportion  $1 - \varepsilon$  of the values fall inside the region. By taking different values for  $a$  we can get different points where we can derive the correspondence. For instance, for  $a = 1$ , we get  $p = 0.6827$ . Setting  $\varepsilon = 1 - 0.6827$  we then calculate the conformal region  $\Gamma^{0.3173}$ . The width of this region, denoted by  $width(\Gamma^{0.3173})$ , has to be  $2a\sigma$ . Since we chose  $a = 1$ , we have that:

$$\sigma = \frac{1}{2} width(\Gamma^{0.3173}) \quad (16)$$

We have implemented the inductive conformal regression presented in [Papadopoulos et al. 2011]. It presents seven nonconformity measures. The first one, that we denote by  $A$ , is just  $|y - \hat{y}|$ . The other six are just modifications which take some metrics of the  $k$ -nearest neighbours into account, such as the mean (or median) (input domain) distance in relation to the average distance of the training dataset, or the mean (or median) standard deviation (of the output domain), corresponding to formulas (24), (25), (29), (30), (31), (32) in [Papadopoulos et al. 2011]. Some of them have parameters ( $\gamma$  and  $\rho$ ), which we set to 0.5 (as in [Papadopoulos et al. 2011]).

We analysed the results for all the nonconformity measures, but we just show the results for the nonconformity measure  $A$  in Table XVI, since this measure gives the best results (although the results are relatively similar for all of them). Comparing to the results in Table II, it seems that there is no improvement in the variance estimation, except for linear regression.

## F. COMPARISON BETWEEN NCDE METHODS

At the end of section 5 we perform a selection of some of the NCDE methods seen in the section. In this appendix, we include the results for a selection of the most relevant methods: the own estimation from the base techniques (section 5.1), conformal prediction (appendix E), a conditional density estimation (CDE) method (appendix B), a conditional variance estimation (CVE) method using Tree for residual regression (appendix C), and three enrichment methods described in section 5.3: RBE (using Tree as residual regression), *uKNC* and *BIN*. For all these methods, Tables XVII, XVIII and XIX show the comparison for the base techniques *LR*, *KNN* and *Tree* respectively. We include the results for the metric *msll*, which considers both the quality of the conditional mean and the

<sup>6</sup>For instance, if we have two regions  $\Gamma_1^\varepsilon = [3.2, 5.4]$  and  $\Gamma_2^\varepsilon = [5.3, 15.9]$  for two different examples, we see that the first interval is much tighter. However, we cannot directly see whether an actual value of 5.2 has higher probability for the first example than a value of 14.2 for the second, because we cannot directly derive probabilities. A possible way of answering this specific question is by adjusting the error probability  $\varepsilon$ . If we increase our tolerance, we may get tighter intervals and we may see that some of the values fall out of the interval. However, we cannot derive the probability for each point either. In order to do this we need a conditional probability density function.

Table XVI. Results (using the datasets in Table XI) for three regression methods for which the variance estimation has been replaced by the variance given by conformal prediction using the nonconformity measure  $A$ , denoted by  $LRc$ ,  $kNNc$  and  $Trec$ . The method  $Conf$  keeps the mean which is estimated by conformal prediction.

	LRc mrse	LRc msll	LRc msvr	kNNc mrse	kNNc msll	kNNc msvr	Trec mrse	Trec msll	Trec msvr	Conf mrse	Conf msll	Conf msvr
1	0.32	0.80	0.54	0.36	0.81	0.50	0.43	0.84	0.68	0.38	0.82	0.56
2	0.58	0.79	0.56	0.46	0.78	0.56	0.25	0.70	0.55	0.45	0.78	0.56
3	0.12	0.83	0.85	0.28	0.75	0.71	0.22	0.71	0.54	0.35	0.94	0.95
4	0.15	0.73	0.59	0.39	0.87	0.65	0.32	0.82	0.60	0.37	0.85	0.64
5	0.10	0.64	0.53	0.11	0.63	0.52	0.12	0.64	0.52	0.10	0.63	0.54
6	0.55	0.70	0.61	0.20	0.65	0.58	0.18	0.65	0.56	0.21	0.67	0.58
7	0.97	0.99	1.00	0.43	0.78	0.63	0.46	0.76	0.58	0.46	1.00	1.00
8	0.14	0.64	0.53	0.14	0.63	0.56	0.15	0.61	0.55	0.14	0.64	0.56
9	0.07	0.93	0.89	0.32	0.93	0.72	0.29	0.98	0.90	0.35	1.00	0.99
10	0.31	0.74	0.54	0.32	0.71	0.67	0.19	0.60	0.68	0.34	0.71	0.67
11	0.19	0.70	0.47	0.18	0.79	0.69	0.14	0.70	0.67	0.20	0.71	0.50
12	0.93	0.88	0.67	0.54	0.77	0.49	0.48	0.75	0.70	0.54	0.79	0.53
13	0.03	0.68	0.68	0.50	0.91	0.65	0.44	0.88	0.63	0.51	0.95	0.91
14	0.46	0.76	0.55	0.36	0.75	0.58	0.23	0.68	0.59	0.37	0.75	0.59
15	0.50	0.72	0.62	0.54	0.71	0.57	0.60	0.73	0.58	0.52	0.71	0.55
16	0.49	0.77	0.56	0.32	0.75	0.59	0.24	0.68	0.61	0.33	0.75	0.56
17	0.96	0.96	0.85	0.43	0.82	0.63	0.37	0.78	0.56	0.40	0.84	0.74
18	0.33	0.76	0.57	0.36	0.77	0.58	0.43	0.77	0.58	0.34	0.75	0.55
19	0.35	0.70	0.55	0.37	0.72	0.56	0.45	0.71	0.61	0.36	0.71	0.54
20	0.60	0.82	0.69	0.20	0.70	0.57	0.27	0.72	0.60	0.20	0.70	0.56
Mean	0.41	0.78	0.64	0.34	0.76	0.60	0.31	0.74	0.61	0.35	0.78	0.66

conditional variance estimation, and the metric  $msvr$ , which considers the quality of the conditional variance estimation. The conditional mean is the same for many of these methods. Note that in those cases where the differences are significant,  $uKNC$  and the own estimation are the best (or among the best) methods.  $BIN$  is also a good method. While performance is a very important reason to make a selection, some other criteria (such as simplicity) are used to finally select the methods in section 5.

## G. PROOFS

Here we include the proofs for several results in the paper.

PROOF. (for proposition 3.3) We have that  $r^*(x, \ell, \hat{f})$  can be written as follows:

$$\begin{aligned}
 r^*(x, \ell, \hat{f}) &= \arg \min_t \int_{-\infty}^{\infty} \ell(t, y) \hat{f}(y|x) dy \\
 &= \arg \min_t \int_{-\infty}^{\infty} \ell(t, \hat{\mu}_{\hat{f}}(x) + s) \hat{f}((\hat{\mu}_{\hat{f}}(x) + s)|x) ds \\
 &= \arg \min_t \left\{ \int_{-\infty}^0 \ell(t, \hat{\mu}_{\hat{f}}(x) + s) \hat{f}((\hat{\mu}_{\hat{f}}(x) + s)|x) ds + \int_0^{\infty} \ell(t, \hat{\mu}_{\hat{f}}(x) + s) \hat{f}((\hat{\mu}_{\hat{f}}(x) + s)|x) ds \right\} \\
 &= \arg \min_t \left\{ \int_{-\infty}^0 \ell(t, \hat{\mu}_{\hat{f}}(x) + s) \hat{f}((\hat{\mu}_{\hat{f}}(x) + s)|x) ds - \int_{-\infty}^0 \ell(t, \hat{\mu}_{\hat{f}}(x) - s) \hat{f}((\hat{\mu}_{\hat{f}}(x) - s)|x) ds \right\} \\
 &= \arg \min_t \int_{-\infty}^0 \left\{ \ell(t, \hat{\mu}_{\hat{f}}(x) + s) \hat{f}((\hat{\mu}_{\hat{f}}(x) + s)|x) - \ell(t, \hat{\mu}_{\hat{f}}(x) - s) \hat{f}((\hat{\mu}_{\hat{f}}(x) - s)|x) \right\} ds
 \end{aligned}$$

Since  $\hat{f}$  is symmetric relative to the mean:

$$r^*(x, \ell, \hat{f}) = \arg \min_t \int_{-\infty}^0 \left\{ \ell(t, \hat{\mu}_{\hat{f}}(x) + s) - \ell(t, \hat{\mu}_{\hat{f}}(x) - s) \right\} \hat{f}((\hat{\mu}_{\hat{f}}(x) + s)|x) ds$$

Table XVII. Results (using the datasets in Table XI) for base technique *LR* using a selection of the methods seen in section 5 as described in appendix F. Each section of seven columns shows results for a measure (*msll* or *msvr*). The average ranks (AR) are calculated for these two groups separately. The Friedman statistic for the two sections are (12.41 and 19.89 respectively). Comparing to the Critical Value (14.16), the null hypothesis is not rejected for *msll* but it is rejected for *msvr* (significance level: 0.05). This means that for *msvr* the methods do not perform equally, but we cannot affirm the same for *msll*. Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.6922) imply that the difference is significant (in bold).

	LR msll	LRc msll	LR CDE msll	LR CVE Tree msll	LR ENR Tree msll	LR ENR uKNC msll	LR ENR BIN msll	LR msvr	LRc msvr	LR CDE msvr	LR CVE Tree msvr	LR ENR Tree msvr	LR ENR uKNC msvr	LR ENR BIN msvr
1	0.81	0.80	0.83	0.79	0.79	0.79	0.79	0.63	0.54	0.49	0.54	0.53	0.51	0.52
2	0.79	0.79	0.74	0.77	0.79	0.79	0.79	0.59	0.56	0.51	0.50	0.54	0.52	0.55
3	0.54	0.83	0.61	0.79	0.81	0.67	0.77	0.52	0.85	0.57	0.84	0.84	0.81	0.83
4	0.72	0.73	0.85	0.70	0.73	0.78	0.73	0.55	0.59	0.55	0.55	0.56	0.68	0.58
5	0.81	0.64	0.68	0.63	0.64	0.63	0.63	0.82	0.53	0.53	0.53	0.55	0.54	0.54
6	0.74	0.70	0.67	0.70	0.69	0.70	0.71	0.63	0.61	0.61	0.60	0.59	0.60	0.60
7	0.90	0.99	0.79	0.96	0.96	0.85	0.97	0.54	1.00	0.83	0.97	0.98	0.42	0.98
8	0.73	0.64	0.63	0.64	0.64	0.65	0.64	0.70	0.53	0.54	0.53	0.54	0.52	0.54
9	0.98	0.93	0.97	0.95	0.96	0.75	0.95	0.95	0.89	0.58	0.91	0.92	0.68	0.92
10	0.74	0.74	0.58	0.72	0.71	0.69	0.71	0.55	0.54	0.60	0.45	0.44	0.36	0.40
11	0.74	0.70	0.70	0.73	0.74	0.73	0.71	0.57	0.47	0.65	0.51	0.57	0.52	0.49
12	0.89	0.88	0.75	0.88	0.86	0.86	0.86	0.66	0.67	0.74	0.72	0.57	0.44	0.55
13	0.63	0.68	0.83	0.71	0.66	0.78	0.69	0.60	0.68	0.60	0.64	0.65	0.80	0.63
14	0.83	0.76	0.74	0.74	0.74	0.75	0.75	0.77	0.55	0.49	0.45	0.47	0.46	0.46
15	0.73	0.72	0.81	0.73	0.71	0.72	0.73	0.68	0.62	0.69	0.58	0.58	0.60	0.59
16	0.77	0.77	0.81	0.78	0.79	0.75	0.77	0.58	0.56	0.54	0.60	0.58	0.49	0.56
17	0.97	0.96	0.81	0.96	0.99	0.90	0.97	0.84	0.85	0.62	0.86	0.93	0.45	0.88
18	0.75	0.76	0.82	0.74	0.74	0.75	0.74	0.67	0.57	0.58	0.57	0.58	0.55	0.55
19	0.71	0.70	0.76	0.70	0.70	0.69	0.69	0.60	0.55	0.59	0.54	0.54	0.53	0.53
20	0.83	0.82	0.77	0.81	0.83	0.80	0.81	0.68	0.69	0.66	0.65	0.61	0.58	0.61
Mean	0.78	0.78	0.76	0.77	0.77	0.75	0.77	0.66	0.64	0.60	0.63	0.63	0.55	0.62
AR	5.30	4.40	3.70	3.60	4.10	3.15	3.75	5.25	4.75	3.80	3.70	4.35	<b>2.55</b>	3.60

But  $\ell$  is symmetric, so we have that for every  $y$  and  $r$  we have that  $\ell(y+r, y) = \ell(y-r, y)$  which, jointly with its commutativity, implies  $\ell(y, y-r) = \ell(y, y+r)$ , so a minimum of the above expression can be found when  $t = \hat{\mu}_{\hat{f}}(x)$ , leading to the expression  $\ell(\hat{\mu}_{\hat{f}}(x), \hat{\mu}_{\hat{f}}(x) + s) - \ell(\hat{\mu}_{\hat{f}}(x), \hat{\mu}_{\hat{f}}(x) - s) = 0$ . So,  $r^*(x, \ell, \hat{f}) = \hat{\mu}_{\hat{f}}(x)$ .  $\square$

Table XVIII. Results (using the datasets in Table XI) for base technique  $kNN$  using a selection of the methods seen in section 5 as described in appendix F. Each section of seven columns shows results for a measure ( $msll$  or  $msvr$ ). The average ranks (AR) are calculated for these two groups separately. The Friedman statistic for the two sections are (9.9 and 24.88 respectively). Comparing to the Critical Value (14.16), the null hypothesis is not rejected for  $msll$  but it is rejected for  $msvr$  (significance level: 0.05). This means that for  $msvr$  the methods do not perform equally, but we cannot affirm the same for  $msll$ . Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.6922) imply that the difference is significant (in bold).

	msll							msvr						
	kNN msll	kNNc msll	kNN CDE msll	kNN CVE Tree msll	kNN ENR Tree msll	kNN ENR uKNC msll	kNN ENR BIN msll	kNN msvr	kNNc msvr	kNN CDE msvr	kNN CVE Tree msvr	kNN ENR Tree msvr	kNN ENR uKNC msvr	kNN ENR BIN msvr
1	0.81	0.81	0.85	0.82	0.82	0.81	0.80	0.48	0.50	0.55	0.54	0.52	0.50	0.49
2	0.76	0.78	0.73	0.76	0.78	0.77	0.77	0.49	0.56	0.51	0.46	0.52	0.54	0.52
3	0.68	0.75	0.63	0.69	0.71	0.72	0.69	0.60	0.71	0.63	0.63	0.61	0.56	0.64
4	0.84	0.87	0.84	0.82	0.80	0.84	0.84	0.56	0.65	0.57	0.47	0.53	0.58	0.52
5	0.63	0.63	0.68	0.63	0.63	0.63	0.64	0.52	0.52	0.51	0.53	0.53	0.53	0.51
6	0.65	0.65	0.66	0.65	0.67	0.66	0.68	0.57	0.58	0.59	0.56	0.59	0.60	0.59
7	0.77	0.78	0.77	0.76	0.76	0.79	0.78	0.57	0.63	0.57	0.58	0.60	0.60	0.59
8	0.63	0.63	0.64	0.62	0.63	0.63	0.63	0.55	0.56	0.55	0.55	0.54	0.54	0.54
9	0.98	0.93	0.97	0.98	0.99	0.99	0.99	0.89	0.72	0.81	0.87	0.93	0.90	0.89
10	0.63	0.71	0.59	0.67	0.68	0.64	0.66	0.52	0.67	0.64	0.60	0.58	0.49	0.52
11	0.76	0.79	0.67	0.76	0.76	0.75	0.75	0.61	0.69	0.67	0.64	0.63	0.60	0.62
12	0.78	0.77	0.75	0.82	0.79	0.77	0.77	0.47	0.49	0.59	0.56	0.50	0.49	0.44
13	0.90	0.91	0.89	0.91	0.89	0.90	0.89	0.61	0.65	0.74	0.67	0.59	0.66	0.60
14	0.73	0.75	0.74	0.72	0.73	0.73	0.73	0.51	0.58	0.51	0.48	0.52	0.50	0.51
15	0.73	0.71	0.82	0.74	0.73	0.74	0.75	0.60	0.57	0.67	0.61	0.60	0.60	0.62
16	0.73	0.75	0.81	0.73	0.73	0.74	0.72	0.55	0.59	0.55	0.60	0.55	0.53	0.58
17	0.80	0.82	0.81	0.79	0.79	0.79	0.79	0.50	0.63	0.59	0.50	0.55	0.53	0.54
18	0.76	0.77	0.83	0.75	0.75	0.75	0.75	0.54	0.58	0.56	0.56	0.56	0.57	0.56
19	0.70	0.72	0.75	0.71	0.71	0.70	0.70	0.52	0.56	0.56	0.54	0.54	0.52	0.52
20	0.68	0.70	0.73	0.70	0.69	0.68	0.68	0.50	0.57	0.54	0.55	0.52	0.51	0.52
Mean	0.75	0.76	0.76	0.75	0.75	0.75	0.75	0.56	0.60	0.60	0.58	0.58	0.57	0.57
AR	3.60	5.15	4.40	3.45	4.15	3.80	3.45	<b>2.50</b>	5.55	4.70	4.20	4.10	3.65	3.30

PROOF. (for proposition 7.3) We use the expression for  $\ell_\alpha^A(\hat{y}, y)$  and decompose it depending on whether  $t < y$  or not.

$$\begin{aligned}
 \mathcal{L}(x, t, \hat{f}, \ell_\alpha^A) &= \int_{-\infty}^{\infty} \ell_\alpha^A(t, y) \hat{f}(y|x) dy \\
 &= \int_{-\infty}^t (1 - \alpha)(t - y) \hat{f}(y|x) dy + \int_t^{\infty} \alpha(y - t) \hat{f}(y|x) dy \\
 &= \int_{-\infty}^t (1 - \alpha)t \hat{f}(y|x) dy - \int_{-\infty}^t (1 - \alpha)y \hat{f}(y|x) dy + \int_t^{\infty} \alpha(y) \hat{f}(y|x) dy - \alpha \int_t^{\infty} t \hat{f}(y|x) dy \\
 &= (1 - \alpha)t \hat{F}(t|x) - \int_{-\infty}^t (1 - \alpha)y \hat{f}(y|x) dy + \int_t^{\infty} \alpha y \hat{f}(y|x) dy - \alpha t (1 - \hat{F}(t|x)) \\
 &= (1 - \alpha)t \hat{F}(t|x) - \int_{-\infty}^t y \hat{f}(y|x) dy + \int_{-\infty}^{\infty} \alpha y \hat{f}(y|x) dy - \alpha t (1 - \hat{F}(t|x)) \\
 &= (1 - \alpha)t \hat{F}(t|x) - \int_{-\infty}^t y \hat{f}(y|x) dy + \alpha \hat{\mu}(x) - \alpha t (1 - \hat{F}(t|x)) \\
 &= \alpha \hat{\mu}(x) + t \hat{F}(t|x) - \alpha t - \int_{-\infty}^t y \hat{f}(y|x) dy
 \end{aligned}$$

□

Table XIX. Results (using the datasets in Table XI) for base technique *Tree* using a selection of the methods seen in section 5 as described in appendix F. Each section of seven columns shows results for a measure (*msll* or *msvr*). The average ranks (AR) are calculated for these two groups separately. The Friedman statistic for the two sections are (34.22 and 28.61 respectively). Comparing to the Critical Value (14.16), the null hypothesis is rejected for both *msll* and *msvr* (significance level: 0.05). This means that for both mterics the methods do not perform equally. Differences in average ranks higher than the critical difference for the Nemenyi post-hoc test (0.6922) imply that the difference is significant (in bold).

	Tree msll	Treec msll	Tree CDE msll	Tree CVE Tree msll	Tree ENR Tree msll	Tree ENR uKNC msll	Tree ENR BIN msll	Tree msvr	Treec msvr	Tree CDE msvr	Tree CVE Tree msvr	Tree ENR Tree msvr	Tree ENR uKNC msvr	Tree ENR BIN msvr
1	0.82	0.84	0.84	0.81	0.83	0.82	0.82	0.57	0.68	0.54	0.58	0.63	0.62	0.60
2	0.69	0.70	0.72	0.72	0.68	0.70	0.67	0.53	0.55	0.52	0.54	0.52	0.51	0.52
3	0.70	0.71	0.61	0.71	0.70	0.72	0.70	0.50	0.54	0.54	0.48	0.49	0.52	0.47
4	0.81	0.82	0.86	0.80	0.83	0.83	0.81	0.55	0.60	0.54	0.53	0.60	0.59	0.58
5	0.63	0.64	0.68	0.63	0.64	0.64	0.63	0.51	0.52	0.53	0.54	0.51	0.53	0.54
6	0.65	0.65	0.66	0.65	0.64	0.66	0.65	0.55	0.56	0.58	0.56	0.56	0.57	0.60
7	0.76	0.76	0.77	0.75	0.75	0.75	0.75	0.52	0.58	0.56	0.54	0.53	0.50	0.50
8	0.61	0.61	0.63	0.61	0.60	0.60	0.59	0.51	0.55	0.59	0.54	0.53	0.52	0.53
9	1.00	0.98	0.97	1.00	0.99	0.92	0.99	0.94	0.90	0.64	0.94	0.92	0.72	0.93
10	0.53	0.60	0.58	0.54	0.53	0.52	0.54	0.50	0.68	0.65	0.56	0.56	0.48	0.56
11	0.69	0.70	0.72	0.68	0.68	0.73	0.67	0.65	0.67	0.64	0.65	0.65	0.72	0.64
12	0.71	0.75	0.77	0.76	0.73	0.73	0.76	0.56	0.70	0.51	0.63	0.61	0.61	0.65
13	0.89	0.88	0.86	0.90	0.90	0.89	0.88	0.64	0.63	0.72	0.68	0.62	0.54	0.66
14	0.65	0.68	0.74	0.67	0.66	0.67	0.66	0.51	0.59	0.60	0.53	0.52	0.52	0.53
15	0.73	0.73	0.81	0.73	0.73	0.73	0.73	0.56	0.58	0.64	0.56	0.58	0.55	0.56
16	0.68	0.68	0.77	0.69	0.70	0.69	0.69	0.58	0.61	0.57	0.58	0.57	0.58	0.60
17	0.77	0.78	0.83	0.77	0.78	0.78	0.78	0.51	0.56	0.52	0.55	0.56	0.52	0.54
18	0.77	0.77	0.82	0.76	0.76	0.76	0.75	0.54	0.58	0.55	0.55	0.56	0.54	0.53
19	0.69	0.71	0.76	0.70	0.69	0.69	0.70	0.56	0.61	0.61	0.58	0.55	0.55	0.56
20	0.70	0.72	0.73	0.70	0.69	0.69	0.70	0.56	0.60	0.55	0.60	0.57	0.56	0.59
Mean	0.72	0.74	0.76	0.73	0.73	0.73	0.72	0.57	0.61	0.58	0.59	0.58	0.56	0.58
AR	<b>2.80</b>	5.10	6.00	<b>4.00</b>	<b>3.15</b>	<b>3.65</b>	<b>3.30</b>	<b>2.80</b>	5.90	4.15	4.55	3.70	<b>2.90</b>	4.00

PROOF. (for proposition 7.4) From proposition 7.3 we just derive the expression for minimising the expected loss:

$$r^*(x, \ell_{A, \alpha}, \hat{f}) = \arg \min_t \left\{ \alpha \hat{\mu}(x) + t \hat{F}(t|x) - \alpha t - \int_{-\infty}^t y \hat{f}(y|x) dy \right\}$$

In order to find the minimum, we calculate the first derivative and equal it to 0:

$$\begin{aligned} \hat{F}(t|x) + t \hat{f}(t|x) - \alpha - t \hat{f}(t|x) &= 0 \\ \hat{F}(t|x) &= \alpha \end{aligned}$$

Since the second derivate is positive this is a minimum.  $\square$

PROOF. (for proposition 7.5) We follow the same initial steps as in the absolute case. We derive the expected loss (eq. 2) and decompose the expression for  $\ell_{\alpha}^S(\hat{y}, y)$  depending on whether  $t < y$  or

not.

$$\begin{aligned}
 \mathcal{L}(x, t, \hat{f}, \ell_\alpha^S) &= \int_{-\infty}^{\infty} \ell_\alpha^S(t, y) \hat{f}(y|x) dy \\
 &= \int_{-\infty}^t (1 - \alpha)(t - y)^2 \hat{f}(y|x) dy + \int_t^{\infty} \alpha(y - t)^2 \hat{f}(y|x) dy \\
 &= \int_{-\infty}^t (1 - \alpha)t^2 \hat{f}(y|x) dy + \int_{-\infty}^t (1 - \alpha)(-2ty + y^2) \hat{f}(y|x) dy \\
 &\quad + \int_t^{\infty} \alpha t^2 \hat{f}(y|x) dy + \int_t^{\infty} \alpha(-2ty + y^2) \hat{f}(y|x) dy \\
 &= (1 - \alpha)t^2 \hat{F}(t|x) + \alpha t^2 (1 - \hat{F}(t|x)) \\
 &\quad + \int_{-\infty}^t (1 - \alpha)(-2ty + y^2) \hat{f}(y|x) dy + \int_t^{\infty} \alpha(-2ty + y^2) \hat{f}(y|x) dy \\
 &= (1 - 2\alpha)t^2 \hat{F}(t|x) + \alpha t^2 + \int_{-\infty}^t (1 - 2\alpha)(-2ty + y^2) \hat{f}(y|x) dy - 2\alpha t \hat{\mu}(x) + \alpha \hat{\mu}_2(x) \\
 &= (1 - 2\alpha) \left[ t^2 \hat{F}(t|x) - 2t \int_{-\infty}^t y \hat{f}(y|x) dy + \int_{-\infty}^t y^2 \hat{f}(y|x) dy \right] + \alpha [t^2 - 2t \hat{\mu}(x) + \hat{\mu}_2(x)]
 \end{aligned}$$

where  $\hat{\mu}_2(x)$  is the second raw moment of  $\hat{f}(y|x)$ .  $\square$

PROOF. (for proposition 7.6) From proposition 7.5, we have:

$$\begin{aligned}
 r^*(x, \ell_{S, \alpha}, \hat{f}) &= \arg \min_t \mathcal{L}(x, t, \hat{f}, \ell_\alpha^S) \\
 &= \arg \min_t \left\{ (1 - 2\alpha) \left[ t^2 \hat{F}(t|x) - 2t \int_{-\infty}^t y \hat{f}(y|x) dy + \int_{-\infty}^t y^2 \hat{f}(y|x) dy \right] + \alpha [t^2 - 2t \hat{\mu}(x) + \hat{\mu}_2(x)] \right\}
 \end{aligned}$$

Again, in order to find the minimum, we calculate the first derivative and equal it to 0:

$$\begin{aligned}
 (1 - 2\alpha) \left[ t^2 \hat{f}(t|x) + 2t \hat{F}(t|x) - 2 \int_{-\infty}^t y \hat{f}(y|x) dy - 2t \cdot t \hat{f}(t|x) + t^2 \hat{f}(t|x) \right] + 2\alpha t - 2\alpha \hat{\mu}(x) + 0 &= 0 \\
 (1 - 2\alpha) \left[ 2t \hat{F}(t|x) - 2 \int_{-\infty}^t y \hat{f}(y|x) dy \right] + 2\alpha t - 2\alpha \hat{\mu}(x) &= 0
 \end{aligned}$$

The second derivative is:

$$(1 - 2\alpha) [2\hat{F}(y|x) + 2t\hat{f}(y|x) - 2t\hat{f}(y|x)] + 2\alpha = (1 - 2\alpha)2\hat{F}(y|x) + 2\alpha$$

which is always positive since both  $\hat{F}(y|x)$  and  $\alpha$  are between 0 and 1. Consequently, we have a minimum.  $\square$

PROOF. (for proposition 7.7) Assuming  $\hat{f}(t|x)$  is a normal distribution, we can standardise  $\hat{f}(t|x)$  as  $\phi(t')$  with  $t' = \frac{t - \hat{\mu}(x)}{\hat{\sigma}(x)}$ . Then, proposition 7.6 reduces to:

$$(1 - 2\alpha) \left[ 2t' \Phi(t') - 2 \int_{-\infty}^{t'} y \phi(y) dy \right] + 2\alpha t' - 0 = 0$$

The partial (from  $-\infty$  to  $t$ ) first moment of the standard normal distribution is just  $-\phi(t)$ . This can also be seen as a truncated standard normal distribution whose expected value is:  $\mathbb{E}(u|u \leq t) = -\frac{\phi(t)}{\Phi(t)}$ . Since the truncated standard normal distribution is normalised by  $\Phi(t)$  we get  $-\phi(t)$ . This

can also be obtained by just solving the integral. From here,

$$\begin{aligned} (1 - 2\alpha) [2t'\Phi(t') + 2\phi(t')] + 2\alpha t' &= 0 \\ t'\Phi(t') + \phi(t') + t' \frac{\alpha}{1 - 2\alpha} &= 0 \end{aligned}$$

where  $t$  is obtained using  $t = \hat{\sigma}(x)t' + \hat{\mu}(x)$ .  $\square$

PROOF. (for proposition 8.2) We start from the expression of the expected loss (proposition 7.3):

$$\mathcal{L}(x, t, \hat{f}, \ell_\alpha^A) = - \int_{-\infty}^t y \hat{f}(y|x) dy + \alpha \hat{\mu}(x) + t \hat{F}(t|x) - \alpha t$$

In order to reduce  $\int_{-\infty}^t y \hat{f}(y|x) dy$ , we see that it is a partial moment of the normal distribution. This is equal to an unnormalised version of the expected value of a truncated distribution, which is  $\mathbb{E}(X|X \leq T) = \mu - \sigma \frac{\phi(\tau)}{\Phi(\tau)}$  with  $\tau = \frac{T-\mu}{\sigma}$ . Consequently, this term reduces to  $\Phi(t')\hat{\mu}(x) - \hat{\sigma}(x)\phi(t')$  with  $t' = \frac{t-\hat{\mu}(x)}{\hat{\sigma}(x)}$  ( $t = \hat{\sigma}(x)t' + \hat{\mu}(x)$ ). So, we have:

$$\begin{aligned} \mathcal{L}(x, t, \hat{f}, \ell_\alpha^A) &= \alpha \hat{\mu}(x) + t \Phi(t') - \alpha t - \Phi(t')\hat{\mu}(x) + \hat{\sigma}(x)\phi(t') \\ &= (t - \hat{\mu}(x))\Phi(t') + \hat{\sigma}(x)\phi(t') - \alpha(t - \hat{\mu}(x)) \\ &= \left[ \frac{t - \hat{\mu}(x)}{\hat{\sigma}(x)} \Phi(t') + \frac{\hat{\sigma}(x)}{\hat{\sigma}(x)} \phi(t') - \alpha \frac{t - \hat{\mu}(x)}{\hat{\sigma}(x)} \right] \hat{\sigma}(x) \\ &= [t'\Phi(t') + \phi(t') - \alpha t'] \hat{\sigma}(x) \end{aligned}$$

$\square$

PROOF. (for proposition 8.3) We start from the expression of the expected loss (proposition 7.5):

$$\begin{aligned} \mathcal{L}(x, t, \hat{f}, \ell_\alpha^S) &= (1 - 2\alpha) \left[ t^2 \hat{F}(t|x) - 2t \int_{-\infty}^t y \hat{f}(y|x) dy - \int_{-\infty}^t y^2 \hat{f}(y|x) dy \right] \\ &\quad + \alpha [t^2 - 2t\hat{\mu}(x) + \hat{\mu}_2(x)] \end{aligned} \tag{17}$$

We reduce  $\int_{-\infty}^t y \hat{f}(y|x) dy$  as we did in the proof of proposition 8.2, as a partial moment of the normal distribution, to  $\Phi(t')\hat{\mu}(x) - \hat{\sigma}(x)\phi(t')$  with  $t' = \frac{t-\hat{\mu}(x)}{\hat{\sigma}(x)}$  ( $t = \hat{\sigma}(x)t' + \hat{\mu}(x)$ ).

We reduce  $\int_{-\infty}^t y^2 \hat{f}(y|x) dy$  as a partial second order moment of the normal distribution (or a full second order moment of the truncated normal distribution), to  $\Phi(t')(\hat{\mu}(x)^2 - 2\hat{\mu}(x)\hat{\sigma}(x)\frac{\phi(t')}{\Phi(t')} + \hat{\sigma}(x)^2(1 - t'\frac{\phi(t')}{\Phi(t')}))$ . The last term  $\hat{\mu}_2(x)$  is the second order moment of the normal distribution, which is just  $\hat{\mu}(x)^2 + \hat{\sigma}(x)^2$ . Plugging all this into (17), and using the short notation  $\mu$  for  $\hat{\mu}(x)$  and  $\sigma$  for  $\hat{\sigma}(x)$  we have:

$$\begin{aligned}
\mathcal{L}(x, t, \hat{f}, \ell_\alpha^S) &= (1 - 2\alpha) \left[ t^2 \Phi(t') - 2t(\Phi(t')\mu - \sigma\phi(t')) - \Phi(t')(\mu^2 - 2\mu\sigma\frac{\phi(t')}{\Phi(t')} + \sigma^2(1 - t'\frac{\phi(t')}{\Phi(t')})) \right] \\
&\quad + \alpha[t^2 - 2t\mu + (\mu^2 + \sigma^2)] \\
&= (1 - 2\alpha)\Phi(t') \left[ t^2 + -2t(\mu - \sigma\frac{\phi(t')}{\Phi(t')}) - (\mu^2 - 2\mu\sigma\frac{\phi(t')}{\Phi(t')} + \sigma^2(1 - t'\frac{\phi(t')}{\Phi(t')})) \right] \\
&\quad + \alpha((t - \mu)^2 + \sigma^2) \\
&= (1 - 2\alpha)\Phi(t') [t^2 + -2t(\mu - \sigma q(t')) - (\mu^2 - 2\mu\sigma q(t') + \sigma^2(1 - t'q(t')))] \\
&\quad + \alpha((t - \mu)^2 + \sigma^2) \\
&= \Phi(t')(1 - 2\alpha) [t^2 - 2t(\mu - \sigma q(t')) - (\mu^2 - 2\mu\sigma q(t') + \sigma^2(1 - t'q(t')))] + \alpha((t - \mu)^2 + \sigma^2) \\
&= \Phi(t')(1 - 2\alpha) [(t'\sigma + \mu)^2 - 2(t'\sigma + \mu)(\mu - \sigma q(t')) - (\mu^2 - 2\mu\sigma q(t') + \sigma^2(1 - t'q(t')))] \\
&\quad + \alpha((t'\sigma)^2 + \sigma^2) \\
&= \Phi(t')(1 - 2\alpha) [(t'\sigma)^2 + \mu^2 + 2t'\sigma^2 q(t') - 2\mu^2 + 2\mu\sigma q(t') - \mu^2 + 2\mu\sigma q(t') - \sigma^2 + \sigma^2 t'q(t')] \\
&\quad + \alpha\sigma^2(t' + 1) \\
&= \Phi(t')(1 - 2\alpha) [(t'\sigma)^2 + 3t'\sigma^2 q(t') - 2\mu^2 + 4\mu\sigma q(t') - \sigma^2] + \alpha\sigma^2(t' + 1)
\end{aligned}$$

with  $q(t') = \frac{\phi(t')}{\Phi(t')}$ .  $\square$

## ACKNOWLEDGMENTS

Some ideas in this paper have benefited from early discussions with Peter Flach, around the notions of generalising operating conditions and reframing, as for the view shown in Table I and the relation to conformal prediction. I'm also grateful to Nicolas Lachiche for his very useful comments and suggestions on a preliminary version of this paper. The anonymous reviewers also contributed with many helpful comments. This work was supported by the MEC/MINECO projects CONSOLIDER-INGENIO CSD2007-00022 and TIN 2010-21062-C02-02, and GVA project PROMETEO/2008/051. Finally, part of this work was motivated by the *REFRAME* project (<http://www.reframe-d2k.org>) granted by the European Coordinated Research on Long-term Challenges in Information and Communication Sciences & Technologies ERA-Net (CHIST-ERA), and funded by Fundación Española para la Ciencia y la Tecnología - Ministerio de Economía y Competitividad.

## REFERENCES

- G. Bansal, A. Sinha, and H. Zhao. 2008. Tuning Data Mining Methods for Cost-Sensitive Regression: A Study in Loan Charge-Off Forecasting. *J. Management Information System* 25 (December 2008), 315–336. Issue 3. [1](#), [2](#), [5](#), [8](#), [25](#), [29](#)
- A. P. Basu and N. Ebrahimi. 1992. Bayesian approach to life testing and reliability estimation using asymmetric loss function. *Journal of statistical planning and inference* 29, 1-2 (1992), 21–31. [5](#)
- A. Bella, C. Ferri, J. Hernández-Orallo, and M.J. Ramírez-Quintana. 2010. Quantification via Probability Estimators. In *2010 IEEE International Conference on Data Mining*. IEEE, 737–742. [34](#)
- A. Bella, C. Ferri, J. Hernández-Orallo, and M.J. Ramírez-Quintana. 2013. Aggregative quantification for regression. *Data Mining and Knowledge Discovery* (2013), 1–44. [18](#), [34](#)
- A. Bella, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. 2009. Calibration of Machine Learning Models. In *Handbook of Research on Machine Learning Applications*. IGI Global, 128–146. [2](#), [11](#), [17](#), [33](#)
- A. Bella, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. 2011. Using negotiable features for prescription problems. *Computing* 91, 2 (2011), 135–168. [6](#), [23](#)
- J. Bi and K. P. Bennett. 2003. Regression error characteristic curves. In *Twentieth International Conference on Machine Learning (ICML-2003)*. Washington, DC. [2](#)
- Z. Bosnić and I. Kononenko. 2008. Comparison of approaches for estimating reliability of individual regression predictions. *Data & Knowledge Engineering* 67, 3 (2008), 504–516. [2](#), [15](#), [38](#)
- Z. Bosnić and I. Kononenko. 2009. An overview of advances in reliability estimation of individual predictions in machine learning. *Intelligent Data Analysis* 13, 2 (2009), 385–401. [2](#)



- L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. 1984. *Classification and Regression Trees*. Wadsworth Publishing Company. 14
- P. F. Christoffersen and F. X. Diebold. 1996. Further results on forecasting and model selection under asymmetric loss. *Journal of applied econometrics* 11, 5 (1996), 561–571. 5
- P. F. Christoffersen and F. X. Diebold. 1997. Optimal prediction under asymmetric loss. *Econometric Theory* 13 (1997), 808–817. 5
- I. Cohen and M. Goldszmidt. 2004. Properties and benefits of calibrated classifiers. *Knowledge Discovery in Databases: PKDD 2004* (2004), 125–136. 2, 11
- S. Crone. 2002. Training artificial neural networks for time series prediction using asymmetric cost functions. In *Proceedings of the 9th International Conference on Neural Information Processing*. 5
- J. Demšar. 2006. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research* 7 (2006), 1–30. 13
- M. Dumas, L. Aldred, G. Governatori, and A.H.M. Ter Hofstede. 2005. Probabilistic automated bidding in multiple auctions. *Electronic Commerce Research* 5, 1 (2005), 25–49. 19
- C. Elkan. 2001. The Foundations of Cost-Sensitive Learning. In *Proceedings of the seventeenth International Conference on Artificial Intelligence (IJCAI-01)*, Bernhard Nebel (Ed.). San Francisco, CA, 973–978. 2
- G. Elliott and A. Timmermann. 2004. Optimal forecast combinations under general loss functions and forecast error distributions. *Journal of Econometrics* 122, 1 (2004), 47 – 79. 5
- T. Fawcett. 2006a. An introduction to ROC analysis. *Pattern Recognition Letters* 27, 8 (2006), 861–874. 2
- T. Fawcett. 2006b. ROC graphs with instance-varying costs. *Pattern Recognition Letters* 27, 8 (2006), 882–891. 34
- C. Ferri, P. Flach, and J. Hernández-Orallo. 2002. Learning decision trees using the area under the ROC curve. In *International Conference on Machine Learning*. 139–146. 33
- C. Ferri, P. Flach, and J. Hernández-Orallo. 2003. Improving the AUC of Probabilistic Estimation Trees.. In *Machine Learning: ECML 2003, 14th European Conference on Machine Learning, Proceedings (Lecture Notes in Computer Science)*. Springer, 121–132. 33
- C. Ferri and J. Hernández-Orallo. 2004. Cautious Classifiers. In *ROC Analysis in Artificial Intelligence, 1st International Workshop, ROCAI-2004, Valencia, Spain, August 22, 2004*, J. Hernández-Orallo, C. Ferri, N. Lachiche, and P. A. Flach (Eds.). 27–36. 6, 27
- P. Flach. 2012. *Machine Learning: The Art and Science of Algorithms that Make Sense of Data*. Cambridge University Press. 10, 13
- G. Forman. 2008. Quantifying counts and costs via classification. *Data Min. Knowl. Discov.* 17, 2 (2008), 164–206. 34
- S. García and F. Herrera. 2008. An extension on statistical comparisons of classifiers over multiple data sets for all pairwise comparisons. *The Journal of Machine Learning Research* 9, 2677–2694 (2008), 66. 13
- R. Ghani. 2005. Price prediction and insurance for online auctions. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining (KDD '05)*. ACM, New York, NY, USA, 411–418. 19
- C. W. J. Granger. 1969. Prediction with a generalized cost of error function. *Operational Research* (1969), 199–207. 5
- C. W. J. Granger. 1999. Outline of forecast theory using generalized cost functions. *Spanish Economic Review* 1, 2 (1999), 161–173. 5, 6
- P. Hall, J. Racine, and Q. Li. 2004. Cross-validation and the estimation of conditional probability densities. *J. Amer. Statist. Assoc.* 99, 468 (2004), 1015–1026. 36
- P. Hall, R. C. L. Wolff, and Q. Yao. 1999. Methods for estimating a conditional distribution function. *J. Amer. Statist. Assoc.* (1999), 154–163. 3
- T. J. Hastie, R.J. Tibshirani, and J.H. Friedman. 2009. *The elements of statistical learning: data mining, inference, and prediction*. Springer. 10
- J. Hernández-Orallo. 2013. ROC curves for regression. *Pattern Recognition* 46, 12 (2013), 3395–3411. 33
- J. Hernández-Orallo, P. Flach, and C. Ferri. 2012. A Unified View of Performance Metrics: Translating Threshold Choice into Expected Classification Loss. *Journal of Machine Learning Research* 13 (2012), 2813–2869. 2
- J. Hernández-Orallo, P. Flach, and C. Ferri. 2013. ROC curves in cost space. *Machine Learning* 93, 1 (2013), 71–91. 2
- J. N. Hwang, S. R. Lay, and A. Lippman. 1994. Nonparametric multivariate density estimation: a comparative study. *Signal Processing, IEEE Transactions on* 42, 10 (1994), 2795–2810. 36
- R. J. Hyndman, D. M. Bashtannyk, and G. K. Grunwald. 1996. Estimating and visualizing conditional densities. *Journal of Computational and Graphical Statistics* (1996), 315–336. 3, 35, 36
- N. Japkowicz and M. Shah. 2011. *Evaluating Learning Algorithms: A Classification Perspective*. Cambridge Univ Pr. 13
- M. Jino, B. T. de Abreu, and others. 2010. Machine learning methods and asymmetric cost function to estimate execution effort of software testing. In *Software Testing, Verification and Validation (ICST), 2010 Third International Conference on*. IEEE, 275–284. 5

- B. Kitts and B. Leblanc. 2004. Optimal bidding on keyword auctions. *Electronic Markets* 14, 3 (2004), 186–201. [19](#)
- N. Lachiche and P. Flach. 2003. Improving accuracy and cost of two-class and multi-class probabilistic classifiers using ROC curves. In *International Conference on Machine Learning*, Vol. 20-1. 416. [7](#)
- H. Papadopoulos. 2008. Inductive conformal prediction: Theory and application to neural networks. *Tools in Artificial Intelligence* 18 (2008), 315–330. [2](#), [40](#)
- H. Papadopoulos, K. Proedrou, V. Vovk, and A. Gammerman. 2002. Inductive Confidence Machines for Regression. In *Machine Learning: ECML 2002*, Tapio Elomaa, Heikki Mannila, and Hannu Toivonen (Eds.). Lecture Notes in Computer Science, Vol. 2430. Springer Berlin / Heidelberg, 185–194. [2](#), [39](#)
- H. Papadopoulos, V. Vovk, and A. Gammerman. 2011. Regression conformal prediction with nearest neighbours. *Journal of Artificial Intelligence Research* 40, 1 (2011), 815–840. [2](#), [40](#)
- T. Pietraszek. 2007. On the use of ROC analysis for the optimization of abstaining classifiers. *Machine Learning* 68, 2 (2007), 137–169. [6](#), [27](#)
- J. C. Platt. 1999. Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods. In *Advances in Large Margin Classifiers*. MIT Press, Boston, 61–74. [16](#)
- F. Provost and P. Domingos. 2003. Tree induction for probability-based ranking. *Machine Learning* 52, 3 (2003), 199–215. [33](#)
- R Team and others. 2012. R: A language and environment for statistical computing. *R Foundation for Statistical Computing, Vienna, Austria* (2012). [14](#), [35](#)
- R. Ribeiro. 2011. Utility-based regression. *PhD thesis, Dep. computer Science, Faculty of Sciences - University of Porto* (2011). [34](#)
- M. Rosenblatt. 1969. Conditional probability density and regression estimators. *Multivariate analysis ii* 25 (1969), 31. [3](#)
- S. Rosset, C. Perlich, and B. Zadrozny. 2007. Ranking-based evaluation of regression models. *Knowledge and Information Systems* 12, 3 (2007), 331–353. [2](#)
- R. E. Schapire, P. Stone, D. McAllester, M. L. Littman, and J. A. Csirik. 2002. Modeling auction price uncertainty using boosting-based conditional density estimation. In *International Conference on Machine Learning*. 546–553. [19](#)
- G. Shafer and V. Vovk. 2008. A tutorial on conformal prediction. *The Journal of Machine Learning Research* 9 (2008), 371–421. [2](#), [39](#)
- J. A. Swets, R. M. Dawes, and J. Monahan. 2000. Better Decisions through Science. *Scientific American* 283, 4 (Oct. 2000), 82–87. [2](#)
- R. D. Thompson and A. P. Basu. 1996. Asymmetric loss functions for estimating system reliability. *Bayesian Analysis in Statistics and Econometrics, John Wiley & Sons* (1996), 471–482. [5](#)
- L. Torgo. 2005. Regression error characteristic surfaces. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. ACM, 697–702. [2](#)
- L. Torgo. 2010. *Data mining with R*. Chapman and Hall / CRC Press. [14](#), [35](#)
- L. Torgo and R. Ribeiro. 2007. Utility-based regression. *Knowledge Discovery in Databases: PKDD 2007* (2007), 597–604. [34](#)
- L. Torgo and R. Ribeiro. 2009. Precision and Recall for Regression. In *Discovery Science*. Springer, 332–346. [2](#)
- P. Turney. 2000. Types of cost in inductive concept learning. *Canada National Research Council Publications Archive* (2000). [33](#)
- L. Wasserman. 2006. *All of nonparametric statistics*. Springer-Verlag New York Inc. [16](#), [38](#)
- M. P. Wellman, D. M. Reeves, K. M. Lochner, and Y. Vorobeychik. 2004. Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research* 21 (2004), 19–36. [19](#)
- K. Yu and M. C. Jones. 2004. Likelihood-based local linear estimation of the conditional variance function. *J. Amer. Statist. Assoc.* 99, 465 (2004), 139–144. [16](#), [38](#)
- B. Zadrozny and C. Elkan. 2002. Transforming classifier scores into accurate multiclass probability estimates. In *The 8th ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining*. 694–699. [33](#)
- A. Zellner. 1986. Bayesian estimation and prediction using asymmetric loss functions. *J. Amer. Statist. Assoc.* (1986), 446–451. [5](#)
- H. Zhao, A. P. Sinha, and G. Bansal. 2011. An extended tuning method for cost-sensitive regression and forecasting. *Decision Support Systems* (2011). [1](#), [2](#), [5](#), [8](#), [25](#), [29](#), [32](#)