
Nonparametric Distribution Regression Re-calibration

Anonymous Authors¹

Abstract

A key challenge in probabilistic regression is ensuring that predictive distributions accurately reflect true empirical uncertainty. Minimizing overall prediction error often encourages models to prioritize informativeness over calibration, producing narrow but overconfident predictions. However, in safety-critical settings, trustworthy uncertainty estimates are often more valuable than narrow intervals. Realizing the problem, several recent works have focused on post-hoc corrections; however, existing methods either rely on weak notions of calibration (such as PIT uniformity) or impose restrictive parametric assumptions on the nature of the error. To address these limitations, we propose a novel nonparametric re-calibration algorithm based on conditional kernel mean embeddings, capable of correcting calibration error without restrictive modeling assumptions. For efficient inference with real-valued targets, we introduce a novel characteristic kernel over distributions that can be evaluated in $\mathcal{O}(n \log n)$ time for empirical distributions of size n . We demonstrate that our method consistently outperforms prior re-calibration approaches across a diverse set of regression benchmarks and model classes.

1. Introduction

In safety-critical applications, such as autonomous systems or medical domains, the value of a predictive model hinges not only on its raw predictive power but also on its reliability. Models that predict the outcomes of medical procedures or vehicle trajectories must also provide trustworthy estimates of their own certainty. In probabilistic regression, this requires balancing two complementary objectives: concentrating probability mass around the ground truth (*sharpness*) and accurately reflecting the true empirical error distribu-

tion (*calibration*). However, common training objectives, such as the negative log-likelihood, tend to reward improvements in the former even when achieved at the expense of the latter. As a result, modern neural networks often fail to achieve calibration, manifesting as narrow predictive intervals that fail to capture the true range of outcomes (Minderer et al., 2021). In high-stakes environments, such miscalibration is dangerous; a system that cannot accurately quantify its own uncertainty provides a false sense of security, potentially leading to severe consequences (Kompa et al., 2021).

While calibration has long been a central problem in statistical forecasting (Dawid, 1984), it has recently gained renewed attention in machine learning through post-hoc re-calibration (Song et al., 2019; Gruber & Buettner, 2024). The primary tool in this domain is the *Probability Integral Transform* (PIT) (Dawid, 1984; Diebold et al., 1997; Dawid, 1998; Mitchell & Wallis, 2011; Kuleshov et al., 2018), which serves as both a design objective and evaluation criterion by testing if the cumulative distribution function (CDF) values of observed targets are uniformly distributed. However, as also noted by Gneiting & Resin (2023), PIT uniformity is a necessary but insufficient condition: it is a *marginal* property that allows for “error cancellation”. Consider an autonomous vehicle: a model that is dangerously overconfident in difficult conditions (e.g., heavy fog) can mask this behavior by being underconfident in easy conditions (e.g., clear weather). As long as errors average out globally, the PIT statistic will appear uniform, concealing the model’s unreliability.

The inadequacy of PIT has motivated stronger notions of calibration (Tsyplakov, 2013; Gneiting & Resin, 2023; Widmann et al., 2021; Glaser et al., 2023; Moskvichev & Sejdinovic, 2025); however, existing methods either focus solely on quantifying calibration error or impose restrictive parametric assumptions when attempting to mitigate it. Such constraints limit the use of existing recalibration approaches on complex, real-world data, forcing a choice between flexible but weak methods (PIT-based) or rigorous but unrealistic ones (parametric).

To address this gap, we propose a novel recalibration framework. By leveraging *Conditional Kernel Mean Embeddings* (CKME) to map model representations directly to a calibrated distribution, our approach avoids both the failure

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

modes of marginal PIT and the restrictiveness of parametric assumptions. Our specific contributions are as follows: First, we introduce a nonparametric recalibration algorithm that enforces the strict property of auto-calibration. Second, to achieve scalability, we propose novel characteristic kernel over distributions, based on the notion of energy distance (Szekely & Rizzo, 2004). For real-valued targets, this kernel can be evaluated in $\mathcal{O}(n \log n)$ time, overcoming the quadratic bottlenecks of standard nonparametric distribution kernels. Third, to better understand the theoretical concepts of this approach, we derive a novel calibration–sharpness decomposition of the population-level error, leveraging a notion of generalized conditional mutual information between the target, feature, and prediction.

In our experiments, we show that models such as Distributional Random Forest (Cevid et al., 2022), Mixture Density Networks (Bishop, 1994), Bayesian Neural Networks (Blundell et al., 2015) are often miscalibrated. We show that our algorithm provides better calibration compared to the state-of-the-art PIT calibration (Kuleshov et al., 2018). We evaluate our method by Squared Kernelized Calibration Error (SKCE) (Widmann et al., 2021), a PIT calibration test with Kolmogorov-Smirnov test, and average test set CRPS relative to the same model without recalibration. Our experiments using the UCI Regression Benchmark (Hernandez-Lobato & Adams, 2015) and other data confirm the practical usability of our algorithm and the supporting theory.

2. Related Work

Our approach to calibration is grounded in the statistical properties of proper scoring rules. We build on the theories of proper scoring rules (Gneiting & Raftery, 2007), formal notions of calibration (Gneiting & Resin, 2023), and generalized definitions of entropy, divergence, and mutual information induced by scoring rules (Dawid & Musio, 2014).

Statistical tests of calibration have appeared early in statistical literature, with perhaps (Dawid, 1984) the first to introduce the Probability Integral Transform (PIT) with a goodness-of-fit test. The relation of PIT and calibration is explored among others in (Strähl & Ziegel, 2015; Modeste et al., 2024). Perhaps the first auto-calibration test overcoming the heuristic limitation of the PIT test is introduced in (Tsyplakov, 2013), however it is not consistent in general. In our work, we employ the Squared Kernelized Calibration Error (SKCE) (Widmann et al., 2021) as the most reliable test of calibration.

Towards understanding and decomposing the sources of prediction error, (Gruber & Buettner, 2024) introduce the notion of the Proper Calibration Error, and provide a similar decomposition to ours, but without separating alleatoric

uncertainty explicitly. They introduce a somewhat limited variance regression recalibration, and provide an algorithm to recalibrate with respect to their notion only.

Most similar to our solution are the recalibration algorithms, out of which we use the best performing ones as baseline. Most important and widely used is PIT recalibration, for which we use the state-of-the-art method of (Kuleshov et al., 2018). (Marx et al., 2023) introduce the trainable calibration metric. During training, they add the kernelized auto-calibration error as regularization to balance between sharpness and calibration. However, with their technique they require a post-hoc recalibration algorithm, for which they use PIT recalibration. While an empirically well performing method, it does not provide any guarantees for auto-calibration.

Closest to our work is (Song et al., 2019), with the key difference that they rely on strong parametric assumptions about the nature of the calibration error. They model the parameters describing the calibration error as a Gaussian process, dependent on the first two moments of the original prediction. Also, they do not report the hypothesis test result of (Widmann et al., 2021), as their result is earlier. As future work, the authors propose to solve the same problem with a nonparametric approach, which we address in our paper.

An other strongly related result is (Moskvichev & Sejdinovic, 2025), which evaluates calibration with conditional kernel mean embeddings. Just as in their paper, we propose nonparametric calibration, but with the following differences. First, we consider regression, not just classification. Second, we provide a recalibration algorithm and not just quantify the calibration error. Finally, our evaluation includes not only calibration but the overall score as well.

3. Background

Let \mathcal{X} and \mathcal{Y} be the feature and response spaces, respectively. Let $\mathcal{M}(\mathcal{Y})$ denote the space of probability measures over \mathcal{Y} . Given random variables ξ and η , let \mathbb{P}_ξ , $\mathbb{P}_{\xi|\eta}$ and $\mathbb{P}_{\xi,\eta}$ denote the marginal distribution of ξ , the conditional distribution of ξ given η and the joint distribution of (ξ, η) , respectively. Suppose that we have an i.i.d. sample $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$ from the joint distribution of the feature and the target $(X, Y) \sim \mathbb{P}_{X,Y}$.

3.1. Proper Scoring Rules

Distribution fitting can be performed under various notions of alignment of the observations and predicted distributions. We will consider the concept of *Scoring Rules* (Gneiting & Raftery, 2007), as it provides a unifying framework.

Let $S : \mathcal{M}(\mathcal{Y}) \times \mathcal{Y} \rightarrow \mathbb{R}$ be a function that assigns a score $S(q, y)$ to a prediction $q \in \mathcal{M}(\mathcal{Y})$ and an observation

110 $y \in \mathcal{Y}$. We will use negatively oriented scores, i.e., smaller
 111 scores are better. S is said to be *strictly proper* if

$$112 \quad \mathbb{E}[S(q, Y)] > \mathbb{E}[S(\mathbb{P}_Y, Y)] , \quad (1)$$

114 for all $q \in \mathcal{M}(\mathcal{Y}) \setminus \{\mathbb{P}_Y\}$. That is, in expectation the
 115 minimum score is uniquely obtained at the true distribution
 116 of the target.

117 The well known negative log-likelihood metric corresponds
 118 to the *logarithmic score* $S(q, y) = -\log p_q(y)$ (with p_q
 119 denoting the probability density function of q), and is a
 120 strictly proper scoring rule.

121 For real-valued distributions which cannot be represented as
 122 a density, the *Continuous Ranked Probability Score* (CRPS)
 123 is a widely used (strictly proper) scoring rule:

$$124 \quad S(q, y) = \mathbb{E}|M - y| - \frac{1}{2}\mathbb{E}|M - M'| , \quad (2)$$

125 where $M \sim q$ and M' is an i.i.d. copy of M .

126 The excess score stemming from making an imperfect pre-
 127 diction is called the *divergence* and is denoted with

$$128 \quad d(q^*, q) = \mathbb{E}_{M \sim q^*}[S(q, M)] - \mathbb{E}_{M \sim q^*}[S(q^*, M)] . \quad (3)$$

129 The divergence is always non-negative, and for strictly
 130 proper scoring rules $d(q^*, q) = 0$ implies $q = q^*$. The
 131 second term in (3) is the expected score of a perfect predic-
 132 tion, which is called the *generalized entropy* and denoted
 133 with

$$134 \quad H(q) = \mathbb{E}_{M \sim q}[S(q, M)] . \quad (4)$$

135 For our demonstration purposes, we will introduce a *con-*
 136 *ditional* version of the generalized notion of mutual infor-
 137 mation (Dawid & Musio, 2014). It quantifies conditional
 138 dependence via the expected reduction of entropy, as fol-
 139 lows.

140 **Definition 3.1.** Let Y, X, Z be jointly distributed random
 141 variables. The *generalized conditional mutual information*
 142 of Y and X given Z , induced by the entropy function H is
 143 defined as

$$144 \quad I(Y; X|Z) = \mathbb{E}_{Z, X}[H(\mathbb{P}_{Y|Z}) - H(\mathbb{P}_{Y|Z, X})] . \quad (5)$$

145 If H is induced by a strictly proper scoring rule, then
 146 $I(Y; X|Z)$ is nonnegative and $I(Y; X|Z) = 0$ iff Y is
 147 conditionally independent of X given Z . See Appendix A.3
 148 for details. For the logarithmic score, H and d coincides
 149 with the differential entropy, and Kullback-Leibler diver-
 150 gence, respectively. Whereas in the case of the CRPS score,
 151 H is half the mean absolute error

$$152 \quad H(q) = \frac{1}{2}\mathbb{E}|M - M'| , \quad (6)$$

153 with $M, M' \sim q$ i.i.d., and $d(q^*, q)$ equals to

$$154 \quad \mathbb{E}|M - M^*| - \frac{1}{2}\mathbb{E}|M - M'| - \frac{1}{2}\mathbb{E}|M^* - M'^*| , \quad (7)$$

155 where $M^*, M'^* \sim q^*$ i.i.d., independent of M and M' .
 156 Equation (7) is a well known metric in the statistical liter-
 157 ature, called the *energy distance* (Szekely & Rizzo, 2004;
 158 Baringhaus & Franz, 2004).

3.2. Kernel Mean Embedding of distributions

Kernel based algorithms are a powerful and well developed
 branch of statistical machine learning. We will present
 only the most important concepts needed to introduce ker-
 nel mean embeddings, which is a nonparametric technique
 capable of estimating distances between distributions effi-
 ciently, and estimate conditional distributions, applicable
 over very general feature and target spaces. See (Muandet
 et al., 2017) for a gentle introduction.

Consider a set \mathcal{X} and a positive definite kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. It is a well known result of (Aronszajn, 1950) that every positive definite kernel uniquely defines a *Reproducing Kernel Hilbert Space* $(\mathcal{H}, \langle \cdot, \cdot \rangle_{\mathcal{H}})$ and vice versa. Note that \mathcal{H} is a Hilbert space of functions $\mathcal{X} \rightarrow \mathbb{R}$. For every $x \in \mathcal{X}$ the *canonical feature map* $\phi(x) := k(x, \cdot)$ is an element of \mathcal{H} .

Consider a random variable $X \sim \mathbb{P}_X$, taking values in \mathcal{X} . The expected value

$$159 \quad \mu_X = \mathbb{E}[\phi(X)] \quad (8)$$

is called the *kernel mean embedding* (KME) of \mathbb{P}_X . If the kernel k is so called *characteristic* then the the feature map ϕ is rich enough, so that the expected value (8) encodes the whole distribution \mathbb{P}_X , i.e. the mapping $\mathbb{P}_X \mapsto \mu_X$ is injective.

3.2.1. DISTANCE OF MEAN EMBEDDINGS

The norm of \mathcal{H} (induced by $\langle \cdot, \cdot \rangle_{\mathcal{H}}$) is a powerful tool to define a distance on distributions, via their kernel mean embeddings. Given another random variable $Y \sim \mathbb{P}_Y$ independent of X , it is a well known fact that

$$160 \quad \|\mu_X - \mu_Y\|_{\mathcal{H}}^2 = \mathbb{E}[k(X, X')] + \mathbb{E}[k(Y, Y')] - 2\mathbb{E}[k(X, Y)] , \quad (9)$$

where X' and Y' are i.i.d. copies of X and Y , respectively.

For empirical distributions $\hat{\mathbb{P}}_X = \sum_{i=1}^n \delta_{x_i}$ and $\hat{\mathbb{P}}_Y = \sum_{i=1}^m \delta_{y_i}$, it is easy to see that the distance of their empirical kernel mean embeddings $\hat{\mu}_X = \frac{1}{n} \sum_{i=1}^n \phi(x_i)$ and $\hat{\mu}_Y = \frac{1}{m} \sum_{i=1}^m \phi(y_i)$ can be evaluated in $\mathcal{O}((n+m)^2)$ time, via Equation (9).

An interesting fact is that in the special case of $\mathcal{X} = \mathbb{R}$, it is possible to reduce this time complexity to $\mathcal{O}(n \log n + m \log m)$, when the kernel k is either the Laplace kernel $k(u, v) = e^{-|u-v|/\sigma}$ or the Brownian motion covariance kernel $k(u, v) = |u| + |v| - |u - v|$. See (Bodenham &

Kawahara, 2023; Baringhaus & Franz, 2004; Sejdinovic et al., 2013) for details.

3.2.2. CONDITIONAL KERNEL MEAN EMBEDDING

Let $l : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a kernel on the target space, with canonical feature map $\psi(y) := l(y, \cdot)$. Since we are interested in estimating certain conditional distributions of the target, a central object of this work will be the *Conditional Kernel Mean Embedding* (CKME), defined as

$$\mu_{Y|X} = \mathbb{E}[\psi(Y) | X] . \quad (10)$$

The estimation of $\mu_{Y|X=x}$ based on an i.i.d. sample $\{(y_i, x_i)\}_{i=1}^n$ from $\mathbb{P}_{Y,X}$ and a query point $x \in \mathcal{X}$ can be done as

$$\hat{\mu}_{Y|X=x} = \sum_{i=1}^n \psi(y_i) \beta_i(x) , \quad (11)$$

where $\beta(x) = (K + \lambda I_n)^{-1} (k(x, x_1), \dots, k(x, x_n))^T$ with $\lambda > 0$ being a regularization parameter, K the kernel matrix $[k(x_i, x_j)]_{i,j=1}^n$ and $I_n \in \mathbb{R}^{n \times n}$ the identity matrix. See (Song et al., 2009; Park & Muandet, 2020) for details.

3.3. Notions of calibration

3.3.1. PIT CALIBRATION

A basic notion of calibration in regression is defined via the *Probability Integral Transform* (PIT) of the predictions:

$$Z := F_Q(Y) , \quad (12)$$

where F_Q is the cumulative distribution function (CDF) of the predicted distribution Q .

If $Z \sim U[0, 1]$, then we say that the model is *PIT calibrated*, which implies that the predicted quantiles match the empirical frequencies of observing the target below the the given quantile.

The reason why PIT calibration is a weak notion of calibration is that it can easily happen that model errors cancel out on average (e.g. systematic over- and under-estimation of the target), leading to $Z \sim U[0, 1]$. Therefore, it is possible to satisfy this notion of calibration with unreliable uncertainty estimates. Another problem is that the reliance on CDFs restrict the applicability of this notion to the real-valued target setting. See (Gneiting & Resin, 2023) for interesting negative examples.

3.3.2. AUTO-CALIBRATION

A much stronger notion of calibration can be motivated by the idea of enforcing PIT calibration conditionally on the predictions, i.e. requiring $Z|Q \sim U[0, 1]$. This notion is called *auto-calibration* or *calibration in the strong sense*,

and is defined more generally (Tsyplakov, 2013) via the condition

$$Q = \mathbb{P}_{Y|Q} . \quad (13)$$

Auto-calibration implies that the predicted distribution Q matches the true conditional distribution of the target given the prediction itself. See (Gneiting & Resin, 2023) for a detailed discussion of the hierarchies between different notions of calibration. In particular PIT and other weak notions of calibration follows from auto-calibration under mild technical assumptions.

3.3.3. HYPOTHESIS TESTING CALIBRATION

Testing PIT calibration can be straightforwardly done via goodness-of-fit tests (e.g. Kolmogorov-Smirnov test), applied to the PIT values computed on a test data split.

Testing auto-calibration needs more sophisticated approaches. An early attempt was made by (Tsyplakov, 2013), who proposed a test based on checking the correlation of certain real-valued transformations (such as mean or a given quantile) of the predictions and the PIT values. However, this test is not consistent in general, and can only be applied on real-valued targets.

A consistent and very generally applicable hypothesis test was introduced by (Widmann et al., 2021), based on the *Squared Kernelized Calibration Error* (SKCE), which is defined as the (squared) distance of the mean embedding of the joint distribution of (Q, Y) and the joint distribution of (Q, M) , where $M \sim Q$, i.e.

$$\text{SKCE} = \|\mu_{Q,M} - \mu_{Q,Y}\|_{\mathcal{H}}^2 . \quad (14)$$

Here \mathcal{H} corresponds to a kernel k , which is defined on the product space $\mathcal{M}(\mathcal{Y}) \times \mathcal{Y}$. Usually k is constructed as a product kernel $k((q, y), (q', y')) = k_1(q, q')k_2(y, y')$, where k_1 and k_2 are kernels over $\mathcal{M}(\mathcal{Y})$ and \mathcal{Y} respectively.

Note that if k is characteristic on $\mathcal{M}(\mathcal{Y}) \times \mathcal{Y}$, then $\text{SKCE} = 0$ implies

$$\mathbb{P}_{Q,M} = \mathbb{P}_{Q,Y} , \quad (15)$$

which further implies auto-calibration. See (Widmann et al., 2021) for technical details and (Glaser et al., 2023) for a more efficient variant, applicable to unnormalized densities.

4. Calibration vs. Sharpness principle

There are usually two distinguished sources of uncertainty in probabilistic modeling. The first is called *aleatoric* uncertainty, which stems from the inherent randomness of the target given the features. It cannot be reduced, unless one introduces new features that describe more information about the target. The other source of uncertainty is referred to as *epistemic* uncertainty, which is the result of insuffi-




Figure 1. Fraction of random train-test splits where the hypothesis of auto-calibration was accepted by SKCE at $\alpha = 5\%$. The numbers after the dataset name indicate the size of the test set $|\mathcal{D}_{test}|$, allowing the power of the hypothesis test to be assessed. See Appendix C.2 for detailed results.

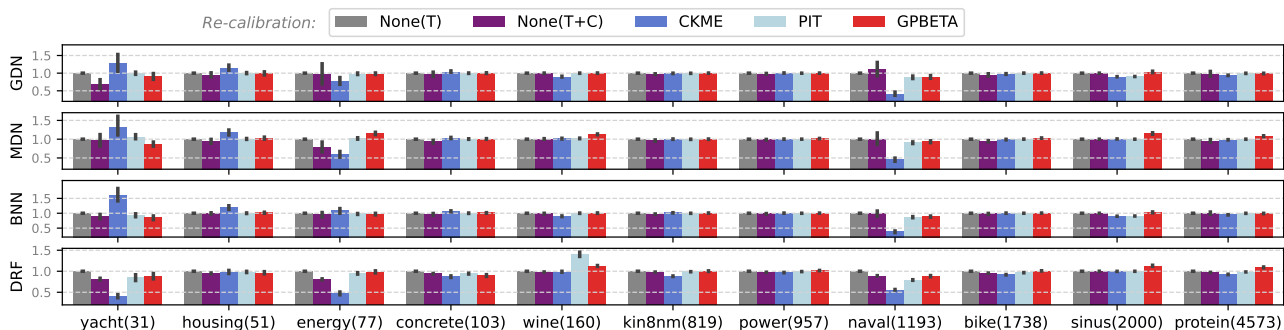


Figure 2. CRPS loss relative to the base model trained only on the test set (None(T)). See Appendix C.2 for detailed results.

cient training data, and is a lack of knowledge which can be entirely eliminated in the limit of $n \rightarrow \infty$.

The paradigm of maximizing *sharpness* subject to *calibration* was first introduced by (Gneiting et al., 2007). Calibration corresponds to how accurately the model represents both aleatoric and epistemic uncertainty, whereas sharpness measures the informativeness of the predictions, i.e., the extent to which the predictions capture the information provided by the features about the target.

The performance of a model Q is usually quantified via the expected error score $\mathbb{E}[S(Q, Y)]$ (such as the negative log-likelihood) it achieves on the whole data distribution. This however conflates the two fundamentally different sources of error: *i)* calibration error, and *ii)* lack of sharpness. Consequently a model with low overall score may not be calibrated, i.e., reliable.

To present this argument more formally, we state the following lemma.

Lemma 4.1. *The sum of calibration error and lack of sharp-*

ness is equal to the divergence from perfect predictions, i.e., the expected error score $\mathbb{E}[S(Q, Y)]$ is equal to

$$\underbrace{\mathbb{E}[d(\mathbb{P}_{Y|Q}, Q)]}_{\text{calibration error}} + \underbrace{I(Y; X|Q)}_{\text{lack of sharpness}} + \underbrace{\mathbb{E}[H(\mathbb{P}_{Y|X})]}_{\text{aleatoric uncertainty}}. \quad (16)$$

Proof. See Appendix A.1. \square

Here the *aleatoric uncertainty* term has nothing to do with the model, it just captures the irreducible inherent randomness of the target, given the features.

We define *lack of sharpness* as the conditional mutual information (Definition 3.1) between the target and the feature given the prediction. That is, the amount of information the feature carries about the target, beyond what is already captured by the prediction. For a perfectly sharp model $I(Y; X|Q) = 0$, i.e., the feature and the target are conditionally independent given the prediction.

Lack of sharpness quantifies the excess entropy of the *target* that the model does not even attempt to capture, even though

it could, in principle, be modeled from the features. A sharp model is often associated with low entropy predictions (e.g. narrow confidence intervals in the real-valued setting). In this formalism, however, this is only a consequence, not the primary definition of sharpness. Narrower confidence intervals arise only from the combination of increased sharpness and accurate uncertainty representation, i.e., low calibration error. Consequently, sharpness only *enables* predictions to be more certain.

By *calibration error* quantified by the first term of (16), we refer to auto-calibration, which is calibration in the strong sense of Section 3.3.2). In other words, when the features are assumed to be *hidden*, the realization of the target corresponding to a *given prediction* is distributed identically to a synthetic sample drawn from that prediction, reflecting our general expectation on the behavior of a reliable uncertainty estimate.

The first two terms of Equation (16) are nonnegative by definition, and are both 0 for the perfect model $Q = \mathbb{P}_{Y|X}$. The divergence from perfect predictions, i.e., $\mathbb{E}[d(\mathbb{P}_{Y|X}, Q)]$ can be manifested in arbitrary combinations of the first two terms. It follows that the same expected score may result from either a well-calibrated model or a sharp yet unreliable one.

5. Re-calibration

Motivated by the calibration–sharpness principle presented in Section 4, in this section, we give a nonparametric, kernel based algorithm to recalibrate a given model while preserving its sharpness, thereby obtaining reliable and useful predictions. To formalize the goal of correcting calibration error, we introduce the following definition.

Definition 5.1. Assume we have a prediction, target, feature tuple $(Q, Y, X) \sim \mathbb{P}_{Q,Y,X}$. Let us define the *recalibrated prediction* \tilde{Q} as

$$\tilde{Q} = \mathbb{P}_{Y|Q}. \quad (17)$$

The following proposition shows that \tilde{Q} indeed achieves our objective. See (Bröcker, 2009, Appendix A) and (Widmann, 2021) for a similar proposition stated for the classification setting.

Proposition 5.2. *The recalibrated prediction \tilde{Q} is Auto-calibrated, and has the same sharpness as the original prediction Q , i.e., we have*

$$\tilde{Q} = \mathbb{P}_{Y|\tilde{Q}} \quad \text{and} \quad I(Y; X|\tilde{Q}) = I(Y; X|Q). \quad (18)$$

Proof. See Appendix A.2. □

5.1. Non-parametric calibration map estimation

Having established the desired properties of \tilde{Q} , the remaining challenge is to estimate the *calibration map* $Q \mapsto \tilde{Q}$. In the pioneering work of (Song et al., 2019), the authors assumed that Q and \tilde{Q} are close enough that their difference can be described by a low-dimensional parametric transformation, where the parameters are dependent on the first two moments of Q .

We relax the heuristic and restrictive assumption on Q and \tilde{Q} , and estimate the calibration map in a fully non-parametric manner using conditional kernel mean embeddings (Song et al., 2009; Park & Muandet, 2020).

In the first step, we embed the predictions of the original model into an RKHS over distributions, i.e., one induced by a kernel $k : \mathcal{M}(\mathcal{Y}) \times \mathcal{M}(\mathcal{Y}) \rightarrow \mathbb{R}$. A general recipe for this is given by the so called *Gaussian-type kernels* introduced by (Christmann & Steinwart, 2010), namely

$$k(q, q') = \exp\left(-\sigma^2 \|\mu_M - \mu_{M'}\|_{\mathcal{H}_r}^2\right). \quad (19)$$

Here μ_M and $\mu_{M'}$ are the kernel mean embeddings of q and q' , respectively, in a RKHS \mathcal{H}_r over \mathcal{Y} , and σ is a bandwidth parameter. The induced kernel k is characteristic provided that r is characteristic.

Given the Gaussian-type kernel k and a kernel $l : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ that encodes the targets via its canonical feature map $\psi : \mathcal{Y} \rightarrow \mathcal{H}_l$, the CKME framework yields an estimator of Q given Q of the form

$$\hat{\mu}_{\tilde{Q}} = \sum_{i=1}^n \psi(y_i) \beta_i(Q). \quad (20)$$

The coefficient vector $\beta(Q) \in \mathbb{R}^n$ is computed as

$$\beta(Q) = (K + \lambda I_n)^{-1} (k(Q, q_1), \dots, k(Q, q_n))^T, \quad (21)$$

where $K \in \mathbb{R}^{n \times n}$ is the Gram matrix $[k(q_i, q_j)]_{i,j=1}^n$, $\lambda > 0$ is a regularization parameter, and $\{(q_i, y_i)\}_{i=1}^n$ is a size n calibration data-set, containing predictions q_i of the original model when we observed $(X, Y) = (x_i, y_i)$.

Although $\hat{\mu}_{\tilde{Q}}$ is a consistent estimator of $\mu_{\tilde{Q}}$ (see (Park & Muandet, 2020, Theorem 4.4.)), note that we only obtain the kernel mean embedding of the recalibrated prediction, rather than an explicit representation of the distribution itself.

Recovering an estimate \hat{Q} corresponds to the distributional pre-image problem (Muandet et al., 2017, sec. 3.8.1). In our approach, this is solved by projecting the weight vector $\beta(Q)$ onto the probability simplex Δ_n to obtain an empirical distribution of the form

$$\hat{Q} = \sum_{i=1}^n w_i \delta_{y_i}, \quad w \in \Delta_n. \quad (22)$$

See Appendix B for details.

5.2. The Energy Distance Kernel (EDK)

Up to this point, we have not made any assumptions on the target space \mathcal{Y} ; in particular, we have not restricted ourselves to recalibrating real-valued distributions, as is done in CDF-based approaches such as (Kuleshov et al., 2018; Song et al., 2019). However, in order to obtain a more efficient algorithm in the special case $\mathcal{Y} = \mathbb{R}$, we propose the *Energy Distance Kernel* (EDK): a specific instantiation of the Gaussian-type kernel (19) by the choice

$$r(u, v) = |u| + |v| - |u - v|. \quad (23)$$

In this case, $\|\mu_M - \mu_{M'}\|_{\mathcal{H}_r}^2$ coincides with the so-called energy distance (cf. Eq. (7); (Szekely & Rizzo, 2004; Sejdinovic et al., 2013)), which admits closed-form expressions for many well-known parametric distribution families and, more importantly, can be evaluated for empirical distributions of size m in $\mathcal{O}(m \log m)$ time. This contrasts with the quadratic complexity (cf. Eq. (9)) incurred when using an arbitrary kernel r on \mathcal{Y} . Efficiency is crucial in practice, since constructing the kernel matrix K requires $\mathcal{O}(n^2)$ evaluations of k .

6. Experiments

We perform a comprehensive benchmark of our proposed CKME based recalibration algorithm described in Section 5. We compare it against the recalibration methods of (Kuleshov et al., 2018) (PIT) and (Song et al., 2019) (GPBETA), as well as against uncalibrated original models trained on \mathcal{D}_{train} (None(T)) or on $\mathcal{D}_{train} \cup \mathcal{D}_{cal}$ (None(T + C)), to ensure fair comparisons. We report p -values from auto-calibration and PIT-calibration hypothesis tests (cf. Section 3.3.3; see aggregated results on Figure 1 and 3 respectively), as well as the mean CRPS score (2) achieved on the test set, across several real-world datasets and a range of machine learning models to be recalibrated. Our experiment code is publicly available at https://anonymous.4open.science/r/recalibration_icml2026-231B.

6.1. Datasets

We use the UCI regression benchmark datasets (Hernandez-Lobato & Adams, 2015), which consist of nine real-word datasets with 20 predefined train-validation-test splits (10% test size, with 20% of the training set held out for validation). The exact splits were taken from the repository of (Gal & Ghahramani, 2016).

In addition, we evaluate our algorithm on the Bike Sharing dataset introduced by (Fanaee-T & Gama, 2014), as well as on a synthetic data set with bimodal target distribution (see Appendix C.1 for details).

6.2. Base Models

We evaluate our recalibration method on a diverse set of probabilistic regression models. Specifically, we consider Distributional Random Forests (DRF; Cevid et al., 2022), Mixture Density Networks (MDN; Bishop, 1994) using the implementation of (Kelen et al., 2025), and Bayesian Neural Network-based MDNs (BNN; Blundell et al., 2015), also following the implementation of (Kelen et al., 2025). In addition, we include a single-component Mixture Density Network, corresponding to a heteroscedastic Gaussian density network (GDN), as a simpler baseline model.

6.3. Experiment setup

We treat the validation split of the original dataset as a calibration set, denoted by \mathcal{D}_{cal} . For every combination of dataset, base model, and recalibration method, we evaluate performance on the test set \mathcal{D}_{test} , collecting results over 20 repetitions of model training, recalibration (when applicable), and testing. When necessary, 20% of the data available for training is held out as a validation set.

Note that in the case of (GPBETA; (Song et al., 2019)), the official implementation¹ can only operate if the output of the base model is a single Gaussian. Consequently, we benchmarked base models with GPBETA recalibration only after approximating the base model’s predicted distribution using a Gaussian fitted to its first two moments.

The hypothesis test of (Widmann et al., 2021) was performed using the authors’ implementation,² but the kernel over distributions was the Energy Distance Kernel (c.f. Section 5.2). The kernel l over the target was the Laplace kernel, and all kernel bandwidth parameters were chosen using the median heuristic.

6.4. Discussion

The aggregated results in Figure 1 demonstrate that, with the exception of our proposed nonparametric recalibration approach, there was generally sufficient evidence to reject the hypothesis of auto-calibration across most datasets by the SKCE test (excluding those with extremely small sample sizes). This highlights the effectiveness of our approach to correct calibration error superior to previous attempts.

Regarding PIT calibration, Figure 3 indicates that the method of (Kuleshov et al., 2018) remains the most effective. Still, our approach achieves performance comparable to the original models and remains more effective than GPBETA, while simultaneously addressing the stronger notions of calibration discussed previously.

¹<https://github.com/Srceh/DistCal>

²<https://github.com/devmotion/CalibrationTests.jl>




Figure 3. Ratio of spits when the hypothesis of PIT-calibration was accepted at $\alpha = 5\%$. The numbers after the dataset name indicate the size of the test set $|\mathcal{D}_{test}|$, allowing the power of the hypothesis test to be assessed. See Appendix C.2 for detailed results.

While sharpness cannot be assessed directly, Figure 2 and Table 1 show that our recalibration method usually was able to marginally improve on the overall score compared to the base model it modified. This suggests that even if there might be some loss of sharpness, it is less important given the improvement on calibration.

By examining Table 1 carefully, we observe that in addition to our algorithm, the base model trained on $\mathcal{D}_{train} \cup \mathcal{D}_{cal}$ provides comparable best scores. The outcomes are consistent with the calibration–sharpness principle discussed in Section 4: data can be utilized to either improve calibration or enhance the overall predictive score. Importantly, we demonstrate the motivating negative example for our work: a model can improve its overall score while remaining significantly miscalibrated. This paradox emphasizes the need for careful testing and correction of calibration alongside standard performance metrics.

7. Conclusions

In this work, we examined the limitations of commonly used calibration techniques in safety-critical regression settings and demonstrated that predictive accuracy alone is insufficient to guarantee reliable uncertainty estimates. In particular, we highlighted the limitations of PIT-based calibration. To address this gap, we introduced a novel recalibration framework based on Conditional Kernel Mean Embeddings (CKME), which directly maps model representations to calibrated predictive distributions without relying on restrictive parametric assumptions.

Empirical results on the UCI Regression Benchmark and additional datasets show that widely used models are frequently miscalibrated, even when they achieve strong predictive scores. Across these experiments, our method consistently improves calibration relative to state-of-the-art recalibration methods, validating both the theoretical foundations

and the practical utility of the proposed framework. The results also reinforce the calibration–sharpness trade-off: available data can be used either to improve calibration or to enhance predictive performance, but gains in one do not necessarily imply gains in the other. Crucially, we demonstrated a negative example in which a model achieves a better overall score while remaining significantly miscalibrated, underscoring the danger of relying solely on standard performance metrics. Together, these findings emphasize the necessity of explicitly testing and correcting calibration, particularly in high-stakes applications where reliable uncertainty quantification is as important as pointwise predictive accuracy.

Impact Statement

This paper presents work whose goal is to advance the field of machine learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

Aronszajn, N. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950. URL <http://dx.doi.org/10.2307/1990404>.

Baringhaus, L. and Franz, C. On a new multivariate two-sample test. *Journal of Multivariate Analysis*, 88:190–206, 02 2004. doi: 10.1016/S0047-259X(03)00079-4.

Bishop, C. M. Mixture density networks. 1994.

Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *International conference on machine learning*, pp. 1613–1622. PMLR, 2015.

- 440 Bodenham, D. and Kawahara, Y. eummd: efficiently computing the mmd two-sample test statistic for univariate data. *Statistics and Computing*, 33, 07 2023. doi: 10.1007/s11222-023-10271-x.
- 441
442
443
444
- 445 Bröcker, J. Reliability, sufficiency, and the decomposition of proper scores. *Quarterly Journal of the Royal Meteorological Society*, 135(643):1512–1519, July 2009. ISSN 1477-870X. doi: 10.1002/qj.456. URL <http://dx.doi.org/10.1002/qj.456>.
- 446
447
448
449
- 450 Cevid, D., Michel, L., Näf, J., Bühlmann, P., and Meinhäusen, N. Distributional random forests: Heterogeneity adjustment and multivariate distributional regression. *Journal of Machine Learning Research*, 23(333):1–79, 2022.
- 451
452
453
454
455
- 456 Chen, Y., Welling, M., and Smola, A. Super-samples from kernel herding, 2012. URL <https://arxiv.org/abs/1203.3472>.
- 457
458
459
- 460 Christmann, A. and Steinwart, I. Universal kernels on non-standard input spaces. In Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., and Culotta, A. (eds.), *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc., 2010. URL https://proceedings.neurips.cc/paper_files/paper/2010/file/4e0cb6fb5fb446d1c92ede2ed8780188-Paper.pdf.
- 461
462
463
464
465
466
467
468
- 469 Dawid, A. and Musio, M. Theory and applications of proper scoring rules. *METRON*, 72, 01 2014. doi: 10.1007/s40300-014-0039-y.
- 470
471
472
- 473 Dawid, A. P. Present position and potential developments: Some personal views statistical theory the prequential approach. *Journal of the Royal Statistical Society: Series A (General)*, 147(2):278–290, 1984.
- 474
475
476
- 477 Dawid, A. P. Coherent measures of discrepancy, uncertainty and dependence, with applications to bayesian predictive experimental design. Technical Report 139, Department of Statistical Science, University College London, 1998. URL <http://www.ucl.ac.uk/Stats/research/pdfs/139b.zip>.
- 478
479
480
481
482
483
- 484 Diebold, F. X., Gunther, T. A., and Tay, A. Evaluating density forecasts, 1997.
- 485
- 486 Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T. Efficient projections onto the 11-ball for learning in high dimensions. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pp. 272–279, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605582054. doi: 10.1145/1390156.1390191. URL <https://doi.org/10.1145/1390156.1390191>.
- 487
488
489
490
491
492
493
494
- Fanaee-T, H. and Gama, J. Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence*, 2(2):113–127, 2014.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/gal16.html>.
- Glaser, P., Widmann, D., Lindsten, F., and Gretton, A. Fast and scalable score-based kernel calibration tests. In Evans, R. J. and Shpitser, I. (eds.), *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pp. 691–700. PMLR, 31 Jul–04 Aug 2023. URL <https://proceedings.mlr.press/v216/glaser23a.html>.
- Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437. URL <https://doi.org/10.1198/016214506000001437>.
- Gneiting, T. and Resin, J. Regression diagnostics meets forecast evaluation: conditional calibration, reliability diagrams, and coefficient of determination. *Electronic Journal of Statistics*, 17(2), January 2023. ISSN 1935-7524. doi: 10.1214/23-ejs2180. URL <http://dx.doi.org/10.1214/23-EJS2180>.
- Gneiting, T., Balabdaoui, F., and Raftery, A. E. Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2):243–268, 2007. doi: <https://doi.org/10.1111/j.1467-9868.2007.00587.x>. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9868.2007.00587.x>.
- Gruber, S. G. and Buettner, F. Better uncertainty calibration via proper scores for classification and beyond, 2024. URL <https://arxiv.org/abs/2203.07835>.
- Hernandez-Lobato, J. M. and Adams, R. Probabilistic back-propagation for scalable learning of bayesian neural networks. In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1861–1869, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/hernandez-lobatoc15.html>.

- 495 Kelen, D. M., Jung, Á., Kersch, P., and Benczur, A. A.
 496 Distribution-free data uncertainty for neural network re-
 497 gression. In *The Thirteenth International Conference on*
 498 *Learning Representations*, 2025.
- 499 Kompa, B., Snoek, J., and Beam, A. L. Second opinion
 500 needed: communicating uncertainty in medical machine
 501 learning. *NPJ Digital Medicine*, 4(1):4, 2021.
- 503 Kuleshov, V., Fenner, N., and Ermon, S. Accurate uncertain-
 504 ties for deep learning using calibrated regression, 2018.
 505 URL <https://arxiv.org/abs/1807.00263>.
- 507 Marx, C., Zalouk, S., and Ermon, S. Calibration by distribu-
 508 tion matching: Trainable kernel calibration metrics, 2023.
 509 URL <https://arxiv.org/abs/2310.20211>.
- 511 Minderer, M., Djolonga, J., Romijnders, R., Hubis, F., Zhai,
 512 X., Houlsby, N., Tran, D., and Lucic, M. Revisiting
 513 the calibration of modern neural networks. *Advances in*
 514 *neural information processing systems*, 34:15682–15694,
 515 2021.
- 516 Mitchell, J. and Wallis, K. F. Evaluating density forecasts:
 517 Forecast combinations, model mixtures, calibration and
 518 sharpness. *Journal of Applied Econometrics*, 26(6):1023–
 519 1040, 2011.
- 521 Modeste, T., Dombry, C., and Fougères, A.-L. Testing
 522 ideal calibration for sequential predictions. working pa-
 523 per or preprint, August 2024. URL [https://hal.](https://hal.science/hal-04679804)
 524 [science/hal-04679804](https://hal.science/hal-04679804).
- 526 Moskvichev, P. and Sejdinovic, D. All models are mis-
 527 calibrated, but some less so: Comparing calibration
 528 with conditional mean operators. In *AI 2025: Ad-*
 529 *vances in Artificial Intelligence: 38th Australasian Joint*
 530 *Conference on Artificial Intelligence, AI 2025, Can-*
 531 *berra, ACT, Australia, December 1–5, 2025, Proceed-*
 532 *ings, Part I*, pp. 274–287, Berlin, Heidelberg, 2025.
 533 Springer-Verlag. ISBN 978-981-95-4968-9. doi: 10.
 534 1007/978-981-95-4969-6_21. URL [https://doi.](https://doi.org/10.1007/978-981-95-4969-6_21)
 535 [org/10.1007/978-981-95-4969-6_21](https://doi.org/10.1007/978-981-95-4969-6_21).
- 536 Muandet, K., Fukumizu, K., Sriperumbudur, B., and
 537 Schölkopf, B. Kernel mean embedding of distributions: A
 538 review and beyond. *Foundations and Trends® in Machine*
 539 *Learning*, 10(1–2):1–141, 2017. ISSN 1935-8245. doi:
 540 10.1561/22000000060. URL [http://dx.doi.org/](http://dx.doi.org/10.1561/22000000060)
 541 [10.1561/22000000060](http://dx.doi.org/10.1561/22000000060).
- 543 Park, J. and Muandet, K. A measure-theoretic approach to
 544 kernel conditional mean embeddings. In Larochelle, H.,
 545 Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.),
 546 *Advances in Neural Information Processing Systems*,
 547 volume 33, pp. 21247–21259. Curran Associates, Inc.,
 548 2020. URL [https://proceedings.neurips.](https://proceedings.neurips.cc/paper_files/paper/2020/file/f340f1b1f65b6df5b5e3f94d95b11daf-Paper.pdf)
 549 [cc/paper_files/paper/2020/file/](https://proceedings.neurips.cc/paper_files/paper/2020/file/f340f1b1f65b6df5b5e3f94d95b11daf-Paper.pdf)
[f340f1b1f65b6df5b5e3f94d95b11daf-Paper.](https://proceedings.neurips.cc/paper_files/paper/2020/file/f340f1b1f65b6df5b5e3f94d95b11daf-Paper.pdf)
[pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/f340f1b1f65b6df5b5e3f94d95b11daf-Paper.pdf).
- Schuster, I., Mollenhauer, M., Klus, S., and Muandet,
 K. Kernel conditional density operators. In Chiappa,
 S. and Calandra, R. (eds.), *Proceedings of the Twenty*
Third International Conference on Artificial Intelligence
and Statistics, volume 108 of *Proceedings of Machine*
Learning Research, pp. 993–1004. PMLR, 26–28 Aug
 2020. URL [https://proceedings.mlr.press/](https://proceedings.mlr.press/v108/schuster20a.html)
[v108/schuster20a.html](https://proceedings.mlr.press/v108/schuster20a.html).
- Sejdinovic, D., Sriperumbudur, B., Gretton, A., and Fuku-
 mizu, K. Equivalence of distance-based and rkhs-based
 statistics in hypothesis testing. *The Annals of Statis-*
tics, 41(5), October 2013. ISSN 0090-5364. doi:
 10.1214/13-aos1140. URL [http://dx.doi.org/](http://dx.doi.org/10.1214/13-AOS1140)
[10.1214/13-AOS1140](http://dx.doi.org/10.1214/13-AOS1140).
- Song, H., Diethe, T., Kull, M., and Flach, P. Distribution
 calibration for regression. In *International Conference*
on Machine Learning, pp. 5897–5906. PMLR, 2019.
- Song, L., Huang, J., Smola, A., and Fukumizu, K.
 Hilbert space embeddings of conditional distributions
 with applications to dynamical systems. In *Proceed-*
ings of the 26th Annual International Conference on
Machine Learning, ICML '09, pp. 961–968, New
 York, NY, USA, 2009. Association for Computing
 Machinery. ISBN 9781605585161. doi: 10.1145/
 1553374.1553497. URL [https://doi.org/10.](https://doi.org/10.1145/1553374.1553497)
[1145/1553374.1553497](https://doi.org/10.1145/1553374.1553497).
- Strähl, C. and Ziegel, J. F. Cross-calibration of probabilistic
 forecasts, 2015. URL [https://arxiv.org/abs/](https://arxiv.org/abs/1505.05314)
[1505.05314](https://arxiv.org/abs/1505.05314).
- Szekely, G. and Rizzo, M. Testing for equal distributions in
 high dimension. *InterStat*, 5, 11 2004.
- Tsyplakov, A. Evaluation of probabilistic forecasts: Proper
 scoring rules and moments. *SSRN Electronic Journal*, 03
 2013. doi: 10.2139/ssrn.2236605.
- Widmann, D. Calibration analysis of probabilistic
 models in julia. YouTube video, 2021. URL
[https://www.youtube.com/watch?v=](https://www.youtube.com/watch?v=PrLsXFvwzuA&t=581s)
[PrLsXFvwzuA&t=581s](https://www.youtube.com/watch?v=PrLsXFvwzuA&t=581s). Talk at JuliaCon 2021.
- Widmann, D., Lindsten, F., and Zachariah, D. Calibration
 tests beyond classification. In *International Conference*
on Learning Representations, 2021. URL [https://](https://openreview.net/forum?id=-bxf89v3Nx)
openreview.net/forum?id=-bxf89v3Nx.

A. Proofs

A.1. Proof of Lemma 4.1

Proof. Using the law of iterated expectations, we have

$$\mathbb{E}[S(Q, Y)] = \mathbb{E}_Q [\mathbb{E}_{Y|Q} [S(Q, Y) | Q]] \quad (24)$$

$$= \mathbb{E}_Q [\mathbb{E}_{Y|Q} [S(Q, Y) - S(\mathbb{P}_{Y|Q}, Y) + S(\mathbb{P}_{Y|Q}, Y) | Q]] \quad (25)$$

$$= \mathbb{E}_Q [\mathbb{E}_{Y|Q} [S(Q, Y) - S(\mathbb{P}_{Y|Q}, Y) | Q]] + \mathbb{E}_Q [\mathbb{E}_{Y|Q} [S(\mathbb{P}_{Y|Q}, Y) | Q]] \quad (26)$$

$$= \mathbb{E}_Q [d(\mathbb{P}_{Y|Q}, Q)] + \mathbb{E}_Q [H(\mathbb{P}_{Y|Q})] . \quad (27)$$

It remains to be shown that $\mathbb{E}_Q [H(\mathbb{P}_{Y|Q})] = I(Y; X|Q) + \mathbb{E}_X [H(\mathbb{P}_{Y|X})]$.

Note that $\mathbb{P}_{Y|X} = \mathbb{P}_{Y|Q, X}$ under the very natural assumption that the model and the target are conditionally independent given the feature (i.e. there is no side information). Again using the law of total expectation, we have

$$\mathbb{E}_Q [H(\mathbb{P}_{Y|Q})] = \mathbb{E}_Q [\mathbb{E}_{X|Q} [H(\mathbb{P}_{Y|Q}) | Q]] \quad (28)$$

$$= \mathbb{E}_Q [\mathbb{E}_{X|Q} [H(\mathbb{P}_{Y|Q}) - H(\mathbb{P}_{Y|X}) + H(\mathbb{P}_{Y|X}) | Q]] \quad (29)$$

$$= \mathbb{E}_Q [\mathbb{E}_{X|Q} [H(\mathbb{P}_{Y|Q}) - H(\mathbb{P}_{Y|X}) | Q]] + \mathbb{E}_Q [\mathbb{E}_{X|Q} [H(\mathbb{P}_{Y|X}) | Q]] \quad (30)$$

$$= \mathbb{E}_Q [H(\mathbb{P}_{Y|Q}) - \mathbb{E}_{X|Q} [H(\mathbb{P}_{Y|X}) | Q]] + \mathbb{E}_Q [\mathbb{E}_{X|Q} [H(\mathbb{P}_{Y|X}) | Q]] \quad (31)$$

$$= \underbrace{\mathbb{E}_Q [H(\mathbb{P}_{Y|Q}) - \mathbb{E}_{X|Q} [H(\mathbb{P}_{Y|Q, X}) | Q]]}_{I(Y; X|Q)} + \mathbb{E}_X [H(\mathbb{P}_{Y|X})] , \quad (32)$$

which concludes the proof. \square

A.2. Proof of Proposition 5.2

Proof. In general any model is calibrated if it is in the form of a conditional law $\mathbb{P}_{Y|\phi(X)}$, where ϕ is an arbitrary measurable function. Let $Z = \phi(X)$ and $W = \mathbb{P}_{Y|Z} = \psi(Z)$. Then we have

$$\mathbb{P}_{Y|W} = \mathbb{E}_{Z|W} [\mathbb{P}_{Y|Z} | W] = \mathbb{E}_{Z|W} [W | W] = W . \quad (33)$$

Choosing $\phi(X) = Q|X$ concludes the first part of the proof.

Under the very natural assumption that $Q \perp\!\!\!\perp Y$ given X , we have $\mathbb{P}_{Y|Q, X} = \mathbb{P}_{Y|\tilde{Q}, X} = \mathbb{P}_{Y|X}$. Therefore we only have to show that

$$\mathbb{E}_Q [H(\mathbb{P}_{Y|Q})] = \mathbb{E}_{\tilde{Q}} [H(\mathbb{P}_{Y|\tilde{Q}})] , \quad (34)$$

in order to prove the right hand side of Equation (18). From Definition 5.1 and the first part of the proof we have

$$\mathbb{P}_{Y|Q} = \tilde{Q} = \mathbb{P}_{Y|\tilde{Q}} . \quad (35)$$

Using Equation (35) and the total law of expectation, it is straightforward to verify Equation (34). \square

A.3. Properties of generalized entropy and mutual information

We include Definition A.1, Proposition A.2 and A.3 only for completeness. They can be readily found e.g. in (Gneiting & Raftery, 2007; Dawid & Musio, 2014; Dawid, 1998). The only novelty is Proposition A.4, which follows straightforwardly from Proposition A.3.

Definition A.1. A function $H : \mathcal{M}(\mathcal{Y}) \rightarrow \mathbb{R}$ is *concave*, if for all $\lambda \in [0, 1]$ and any $q_1, q_2 \in \mathcal{M}(\mathcal{Y})$ we have

$$\lambda H(q_1) + (1 - \lambda)H(q_2) \leq H(\lambda q_1 + (1 - \lambda)q_2) , \quad (36)$$

where $\lambda q_1 + (1 - \lambda)q_2$ is understood as a mixture distribution. We call H *strictly concave* if (36) is satisfied with strict inequality for $\lambda \in (0, 1)$.

Proposition A.2. *The generalized entropy (Equation (4)) is concave, and is strictly concave, when the underlying proper scoring rule is strictly proper.*

Proof. Let $q_1, q_2 \in \mathcal{M}(\mathcal{Y})$ and $\lambda \in [0, 1]$ be arbitrary. Define $q = \lambda q_1 + (1 - \lambda)q_2$. Observe that

$$\mathbb{E}_{M \sim q_1} [S(q_1, M)] \leq \mathbb{E}_{M \sim q_1} [S(q, M)] \quad (37)$$

$$\mathbb{E}_{M \sim q_2} [S(q_2, M)] \leq \mathbb{E}_{M \sim q_2} [S(q, M)] , \quad (38)$$

since S is (strictly) proper. Now add the (37) and (38) inequalities together multiplied by weights λ and $(1 - \lambda)$ respectively and observe that:

$$\lambda \mathbb{E}_{M \sim q_1} [S(q_1, M)] + (1 - \lambda) \mathbb{E}_{M \sim q_2} [S(q_2, M)] \leq \mathbb{E}_{M \sim q} [S(q, M)] , \quad (39)$$

since by the linearity of the expectation operator we have $\lambda \mathbb{E}_{M \sim q_1} + (1 - \lambda) \mathbb{E}_{M \sim q_2} = \mathbb{E}_{M \sim q}$.

□

Let us state an important property of the entropy function H , which enables us to measure the dependence of two random variables via a general scoring rule.

Proposition A.3. *Let ξ and η be jointly distributed random variables. The generalized entropy is monotone, that is*

$$H(\mathbb{P}_\xi) \geq \mathbb{E}_{\eta \sim \mathbb{P}_\eta} [H(\mathbb{P}_{\xi|\eta})] . \quad (40)$$

When using a strictly proper scoring rule, there is equality in (40) iff ξ and η are independent.

Proof. Observe that $\mathbb{P}_\xi = \mathbb{E}_{\eta \sim \mathbb{P}_\eta} [\mathbb{P}_{\xi|\eta}]$ is a convex mixture. Since H is (strictly) concave, by the Jensen inequality we have

$$H(\mathbb{E}_{\eta \sim \mathbb{P}_\eta} [\mathbb{P}_{\xi|\eta}]) \geq \mathbb{E}_{\eta \sim \mathbb{P}_\eta} [H(\mathbb{P}_{\xi|\eta})] . \quad (41)$$

If there is an event with nonzero \mathbb{P}_η probability, where $\mathbb{P}_{\xi|\eta}$ differs from \mathbb{P}_ξ , then by the strict concavity of H , we will get a strict inequality in (41). □

Based on Proposition A.3, one define the generalized mutual information of random variables ξ, η as

$$I(\xi; \eta) = H(\mathbb{P}_\xi) - \mathbb{E}_{\eta \sim \mathbb{P}_\eta} [H(\mathbb{P}_{\xi|\eta})] , \quad (42)$$

i.e., via the amount of expected entropy reduction of ξ , if we condition on η . If S is strictly proper, then $I(\xi, \eta) = 0$ implies that ξ is independent of η . (Dawid & Musio, 2014; Dawid, 1998)

For convenience, we will introduce the generalized *conditional* mutual information induced by S as

$$I(\xi; \eta|\zeta) = \mathbb{E}_{\zeta \sim \mathbb{P}_\zeta} [H(\mathbb{P}_{\xi|\zeta}) - \mathbb{E}_{\eta \sim \mathbb{P}_{\eta|\zeta}} [H(\mathbb{P}_{\xi|(\eta, \zeta)})]] . \quad (43)$$

Proposition A.4. *Given a strictly proper scoring rule, the generalized conditional mutual information (Equation (43)) is nonnegative and characterizes conditional independence, i.e.*

$$I(\xi; \eta|\zeta) = 0 \iff \xi|\zeta \perp\!\!\!\perp \eta|\zeta \quad \text{with } \mathbb{P}_\zeta \text{ probability 1} . \quad (44)$$

Proof. Conditioning on fixed events $\{\zeta = \zeta_0\}$, Proposition A.3 can be applied point-wise. Taking the expectation with respect to \mathbb{E}_ζ results in having the non-negativity and characterization of conditional independence \mathbb{P}_ζ almost everywhere. □

B. Distributional pre-image problem

Conditional kernel mean embeddings only estimate $\hat{\mu}_{Y|X} \approx \mathbb{E}[\psi(Y)|X]$, which is a representation of $\hat{\mathbb{P}}_{Y|X}$ not necessarily easy to work with. To be able to quantify the model error $\mathbb{E}[S(\hat{\mathbb{P}}_{Y|X}, Y)]$ or have predictive quantiles $\hat{\mathbb{P}}(Y \leq t)$, one often needs a more exact form of $\hat{\mathbb{P}}_{Y|X}$. This problem is called the distributional pre-image problem, since we are interested in finding the distribution $q \in \mathcal{M}(\mathcal{Y})$ whose kernel mean embedding is $\hat{\mu}_{Y|X=x}$. There are several different approaches to

solve this problem (Muandet et al., 2017; Chen et al., 2012; Schuster et al., 2020). In order to minimize computational complexity and approximation bias, we have chosen the following approximate pre-image approach.

The approximate distributional pre-image solution starts with choosing a family of parametrized distributions $\{q_\theta \mid \theta \in \Theta\} \subset \mathcal{M}(\mathcal{Y})$ and define the approximate pre-image $\tilde{\mathbb{P}}_{Y|X=x}$ as

$$\tilde{\mathbb{P}}_{Y|X=x} = \arg \min_{\theta \in \Theta} \|\hat{\mu}_{Y|X=x} - \mathbb{E}_{M \sim q_\theta} [\psi(M)]\|_{\mathcal{H}_l}^2. \quad (45)$$

If the parameterized family is too rich, then the optimization (45) can be challenging to solve, and if it is too restrictive then one introduces a significant approximation error to the predictions. We made the following practical choice: Let $n = |\mathcal{D}_{cal}|$, let $\Theta = \Delta_n$, i.e. the n dimensional probability simplex, and let $q_\theta = \sum_{i=1}^n \theta_i \delta_{y_i}$ be the empirical distribution supported on the observations in the calibration set with weight θ_i for the point-mass δ_{y_i} . This is a reasonable choice since extending the support of $\tilde{\mathbb{P}}_{Y|X}$ beyond $\{y_i \mid i \in \mathcal{D}_{cal}\}$ requires prior knowledge (or assumptions) about $\mathbb{P}_{Y|X}$.

With this choice (45) is easy to show³ to be equivalent with

$$\theta^* = \arg \min_{\theta \in \Delta_n} (\theta - \beta)^T L (\theta - \beta), \quad (46)$$

where β are the weights in the conditional kernel mean embedding estimate $\hat{\mu}_{Y|X=x} = \sum_{i=1}^n \beta_i \psi(y_i)$, L is the kernel matrix $[l(y_i, y_j)]_{i,j=1}^n$ and $\tilde{\mathbb{P}}_{Y|X=x}$ results to be $\sum_{i=1}^n \theta_i^* \delta_{y_i}$.

Although (46) is a convex problem and therefore can be solved efficiently for each observation $X = x$, solving simultaneously for all observations in \mathcal{D}_{test} is computationally challenging. Consequently, we decided to further approximate θ^* with

$$\tilde{\theta}^* = \arg \min_{\theta \in \Delta_n} (\theta - \beta)^T I_n (\theta - \beta), \quad (47)$$

which is essentially the Euclidean projection of β to Δ_n , for which there is an $\mathcal{O}(n)$ time algorithm (Duchi et al., 2008). Using $\tilde{\theta}^*$ instead of θ^* should be considered an implementation choice, which is not unprecedented; for example, the authors of (Cevic et al., 2022) also used clipped and renormalized CKME weights for inference in their benchmark section.

C. Benchmark

C.1. Synthetic data set

We sampled $n = 20\,000$ i.i.d. feature samples from $U[-1, 1]$ and then generated the corresponding target variable from the 2 component mixture model

$$\frac{1}{2} \mathcal{N}(x + \sin(3x/20), \sigma^2) + \frac{1}{2} \mathcal{N}(x - \sin(3x/20), \sigma^2), \quad (48)$$

where the variance of both compenets were $\sigma^2 = \frac{1}{10^2}$.

C.2. Detailed Benchmark Results

Find detailed benchmark results of the relative CRPS scores in Table 1. Figure 3 contains aggregated results of acceptance rate (at $\alpha = 5\%$) for the hypothesis of PIT-calibration. See Figure 4, 5, 6 and 7 for detailed results of calibration hypothesis tests (where the relative CRPS results are also plotted, for easier comparison.)

³Using the reproducing property of l

Table 1. CRPS scores relative to the base model trained only on the test set (None(T)). Best relative score per row is marked with **bold** values. Standard deviations per split are displayed for each entry. The numbers after the dataset name indicate the size of the test set $|\mathcal{D}_{test}|$.

BASE MODEL	RE-CALIBRATION DATA SET	NONE(T)	NONE(T+C)	CKME	PIT	GPBETA
GDN	YACHT(31)	1.000±0.000	0.690 ±0.298	1.276±0.537	0.997±0.077	0.905±0.198
	HOUSING(51)	1.000±0.000	0.955 ±0.134	1.150±0.170	1.000±0.033	0.995±0.096
	ENERGY(77)	1.000±0.000	0.962±0.679	0.767 ±0.224	0.981±0.048	0.984±0.048
	CONCRETE(103)	1.000±0.000	0.968 ±0.139	1.045±0.047	0.999±0.014	1.000±0.021
	WINE(160)	1.000±0.000	0.996±0.018	0.901 ±0.026	0.999±0.009	1.000±0.011
	KIN8NM(819)	1.000±0.000	0.971 ±0.026	0.994±0.012	0.996±0.005	1.001±0.006
	POWER(957)	1.000±0.000	0.982 ±0.024	1.004±0.007	0.999±0.003	1.001±0.004
	NAVAL(1193)	1.000±0.000	1.104±0.449	0.417 ±0.105	0.882±0.086	0.896±0.090
	BIKE(1738)	1.000±0.000	0.950 ±0.076	0.977±0.023	1.000±0.007	1.003±0.006
	SINUS(2000)	1.000±0.000	1.000±0.002	0.900 ±0.003	0.905±0.002	1.033±0.045
PROTEIN(4573)	1.000±0.000	0.975±0.088	0.937 ±0.008	0.995±0.005	0.991±0.008	
MDN	YACHT(31)	1.000±0.000	0.960±0.365	1.326±0.591	1.058±0.130	0.870 ±0.131
	HOUSING(51)	1.000±0.000	0.944 ±0.114	1.175±0.149	1.007±0.030	1.027±0.064
	ENERGY(77)	1.000±0.000	0.794±0.308	0.594 ±0.178	1.019±0.046	1.153±0.071
	CONCRETE(103)	1.000±0.000	0.934 ±0.088	1.032±0.038	1.003±0.018	1.006±0.021
	WINE(160)	1.000 ±0.000	1.005±0.022	1.015±0.023	1.020±0.018	1.128±0.027
	KIN8NM(819)	1.000±0.000	0.962 ±0.041	1.001±0.013	0.997±0.008	0.999±0.013
	POWER(957)	1.000±0.000	0.986 ±0.018	1.001±0.008	1.000±0.002	1.017±0.006
	NAVAL(1193)	1.000±0.000	1.002±0.344	0.455 ±0.102	0.907±0.053	0.933±0.054
	BIKE(1738)	1.000±0.000	0.947 ±0.038	0.991±0.016	1.001±0.004	1.029±0.009
	SINUS(2000)	1.000±0.000	0.999 ±0.003	1.006±0.004	1.000±0.002	1.150±0.045
PROTEIN(4573)	1.000±0.000	0.957 ±0.042	0.985±0.005	1.000±0.002	1.079±0.013	
BNN	YACHT(31)	1.000±0.000	0.888±0.171	1.614±0.477	0.931±0.117	0.848 ±0.141
	HOUSING(51)	1.000±0.000	0.994 ±0.047	1.188±0.143	1.001±0.022	1.018±0.039
	ENERGY(77)	1.000±0.000	0.946 ±0.175	1.077±0.187	0.977±0.038	0.973±0.059
	CONCRETE(103)	1.000±0.000	0.960 ±0.074	1.065±0.039	1.001±0.015	1.011±0.022
	WINE(160)	1.000±0.000	0.989±0.027	0.899 ±0.029	1.001±0.014	1.005±0.016
	KIN8NM(819)	1.000±0.000	0.968 ±0.017	1.015±0.012	0.999±0.004	1.006±0.009
	POWER(957)	1.000±0.000	0.981 ±0.019	1.004±0.007	0.999±0.002	1.004±0.005
	NAVAL(1193)	1.000±0.000	0.975±0.217	0.376 ±0.067	0.870±0.046	0.889±0.046
	BIKE(1738)	1.000±0.000	0.977 ±0.040	1.003±0.010	1.000±0.002	1.005±0.003
	SINUS(2000)	1.000±0.000	0.999±0.002	0.900 ±0.002	0.907±0.004	1.030±0.034
PROTEIN(4573)	1.000±0.000	1.002±0.051	0.947 ±0.006	0.998±0.002	0.993±0.008	
DRF	YACHT(31)	1.000±0.000	0.814±0.046	0.405 ±0.086	0.850±0.160	0.885±0.147
	HOUSING(51)	1.000±0.000	0.949 ±0.022	0.984±0.088	0.983±0.034	0.964±0.057
	ENERGY(77)	1.000±0.000	0.809±0.025	0.473 ±0.078	0.950±0.040	0.985±0.057
	CONCRETE(103)	1.000±0.000	0.942±0.013	0.871 ±0.045	0.943±0.022	0.899±0.049
	WINE(160)	1.000±0.000	0.981 ±0.006	0.988±0.025	1.404±0.122	1.126±0.024
	KIN8NM(819)	1.000±0.000	0.979±0.004	0.882 ±0.015	0.989±0.004	1.003±0.016
	POWER(957)	1.000±0.000	0.975±0.002	0.969 ±0.009	0.992±0.002	1.017±0.005
	NAVAL(1193)	1.000±0.000	0.893±0.006	0.560 ±0.016	0.792±0.018	0.882±0.033
	BIKE(1738)	1.000±0.000	0.955±0.006	0.918 ±0.010	0.966±0.006	1.009±0.007
	SINUS(2000)	1.000±0.000	0.999±0.003	0.997 ±0.004	0.999±0.002	1.131±0.025
PROTEIN(4573)	1.000±0.000	0.973±0.001	0.923 ±0.007	0.984±0.000	1.092±0.012	

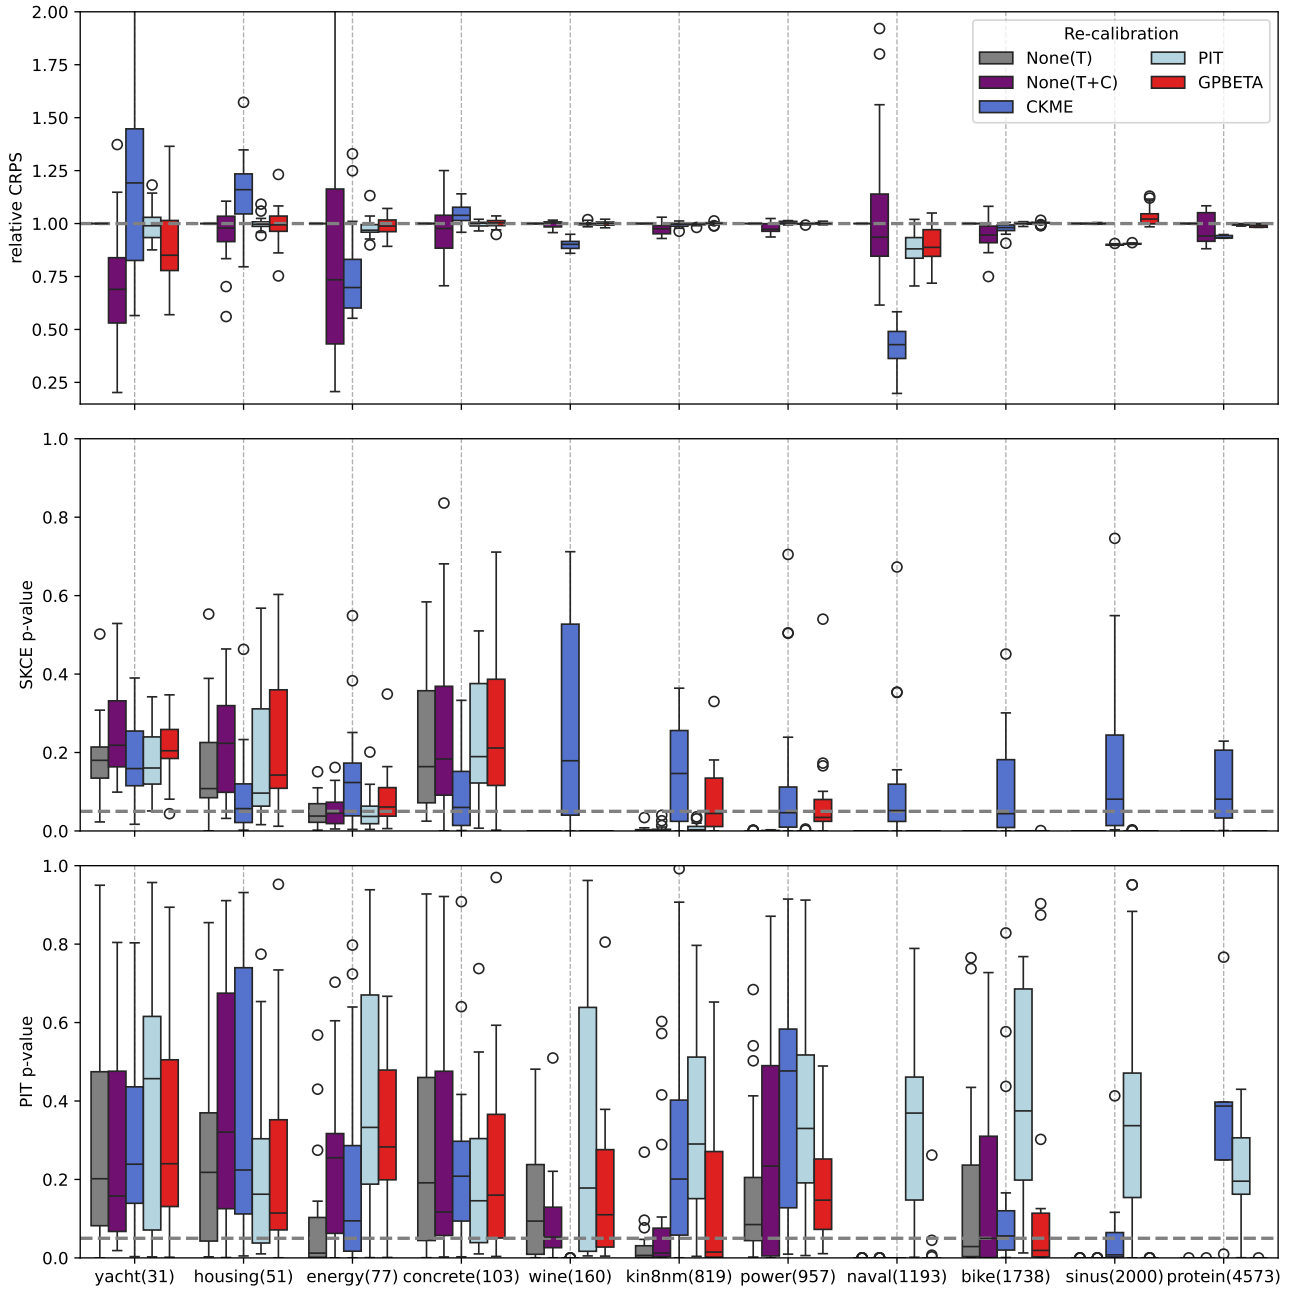


Figure 4. Detailed benchmark results for base model GDN.

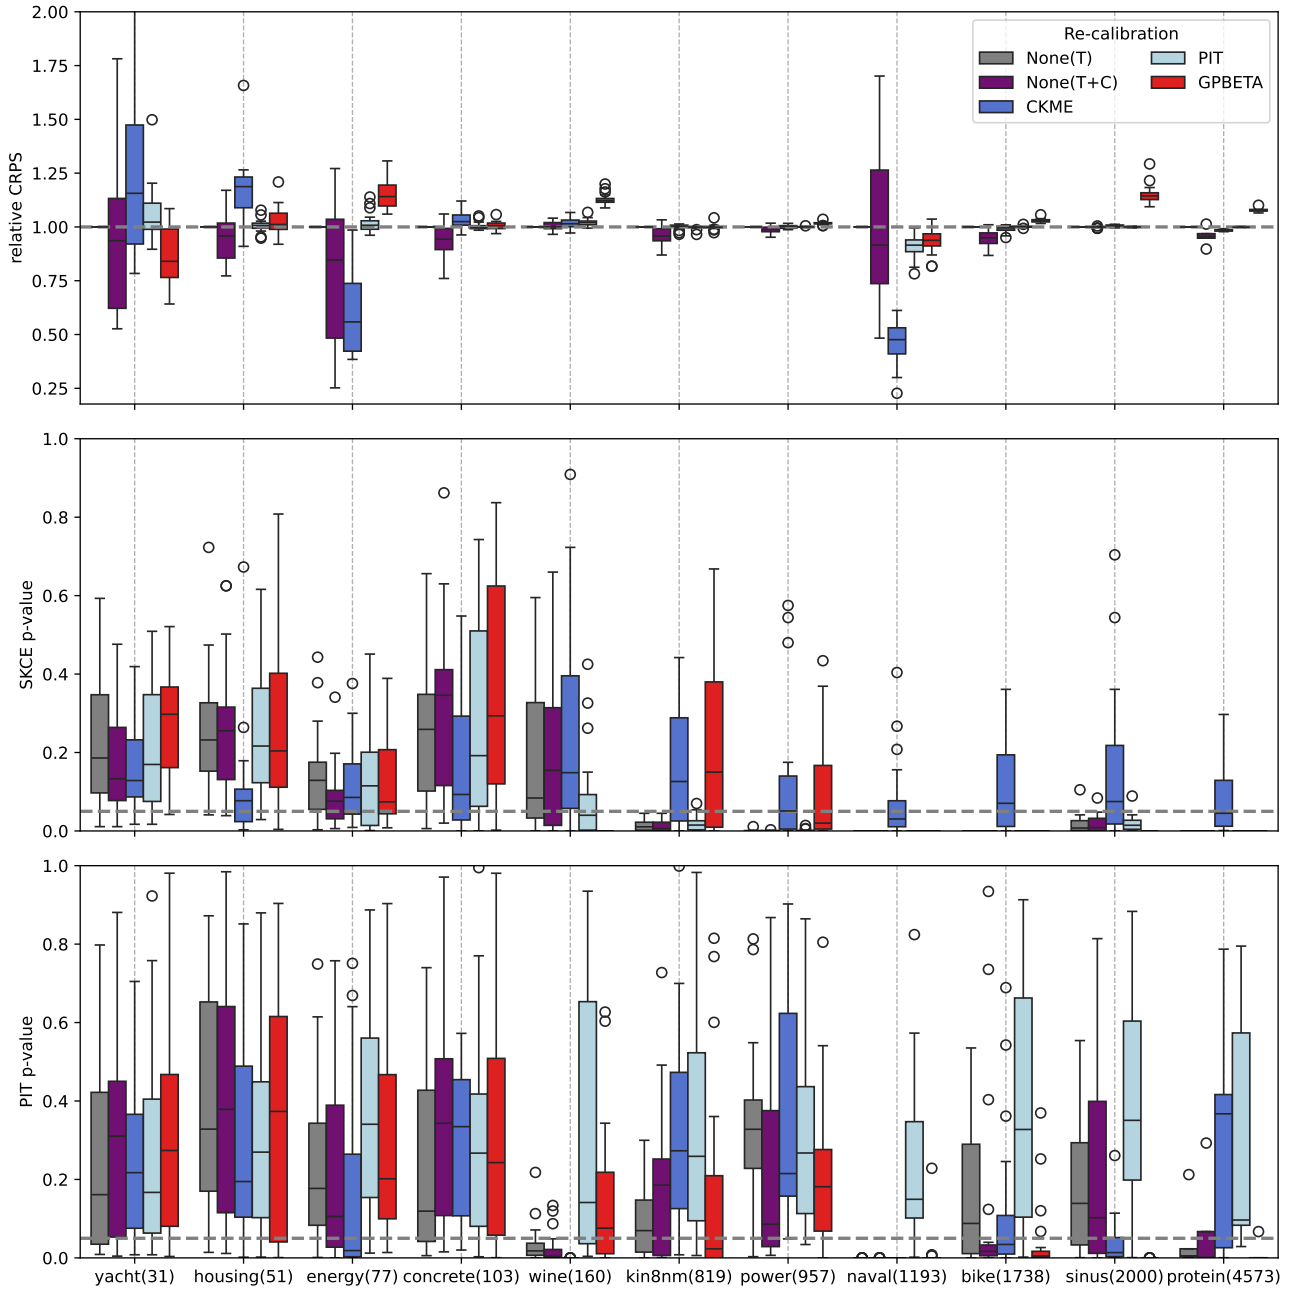


Figure 5. Detailed benchmark results for base model MDN.

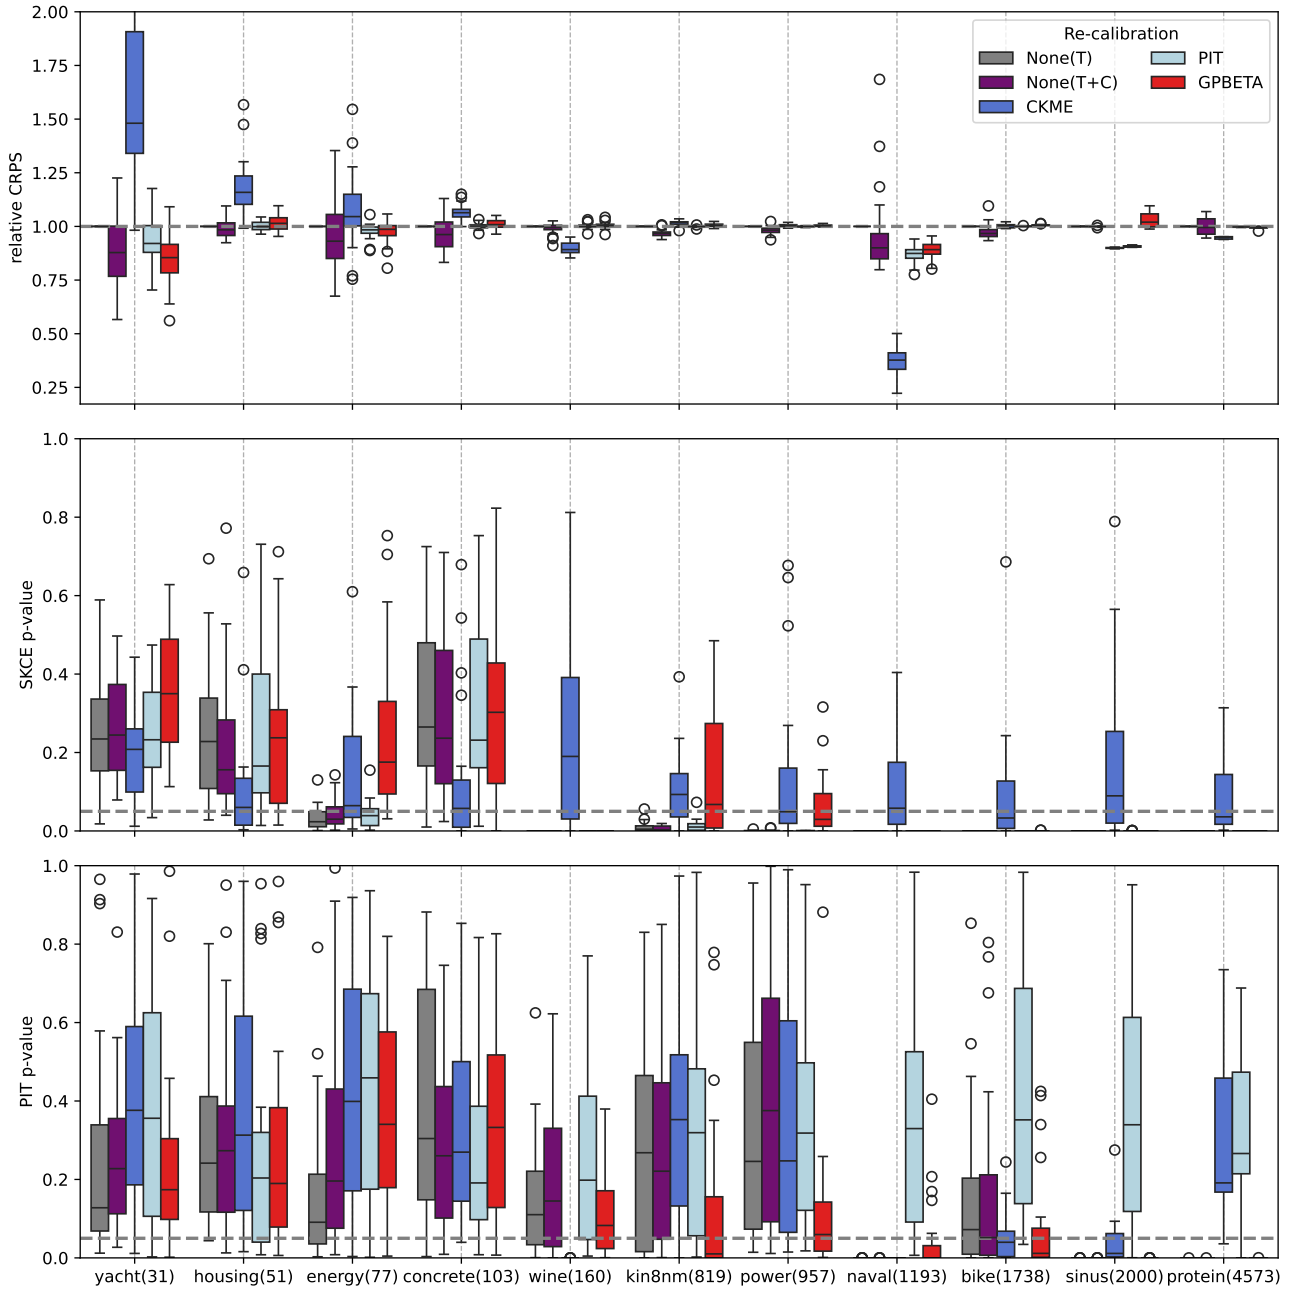


Figure 6. Detailed benchmark results for base model BNN.

Nonparametric Distribution Regression Re-calibration

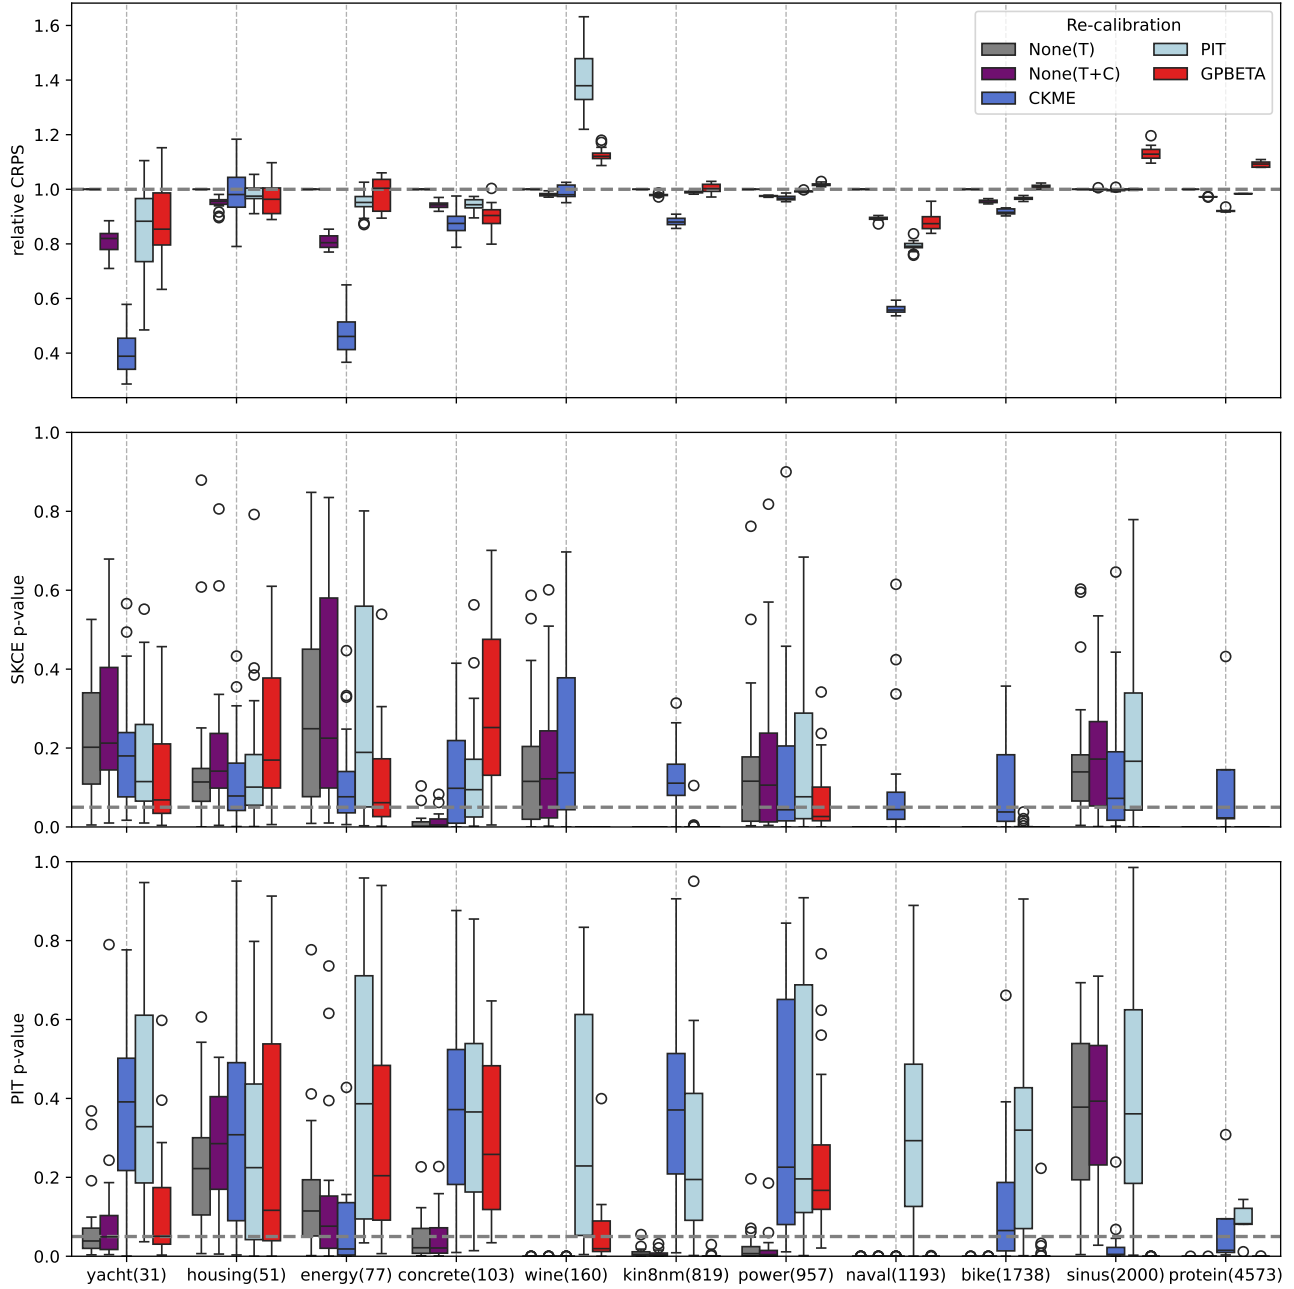


Figure 7. Detailed benchmark results for base model DRF.