

# 000 001 002 003 004 005 CONFEX: UNCERTAINTY-AWARE COUNTERFACTUAL 006 EXPLANATIONS WITH CONFORMAL GUARANTEES 007 008 009

010 **Anonymous authors**  
011  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053  
054  
055  
056  
057  
058  
059  
060  
061  
062  
063  
064  
065  
066  
067  
068  
069  
070  
071  
072  
073  
074  
075  
076  
077  
078  
079  
080  
081  
082  
083  
084  
085  
086  
087  
088  
089  
090  
091  
092  
093  
094  
095  
096  
097  
098  
099  
100

Paper under double-blind review

## ABSTRACT

Counterfactual explanations (CFXs) provide human-understandable justifications for model predictions, enabling actionable recourse and enhancing interpretability. To be reliable, CFXs must avoid regions of high predictive uncertainty, where explanations may be misleading or inapplicable. However, existing methods often neglect uncertainty or lack principled mechanisms for incorporating it with formal guarantees. We propose CONFEX, a novel method for generating uncertainty-aware counterfactual explanations using Conformal Prediction (CP) and Mixed-Integer Linear Programming (MILP). CONFEX explanations are designed to provide local coverage guarantees, addressing the issue that CFX generation violates exchangeability. To do so, we develop a novel localised CP procedure that enjoys an efficient MILP encoding by leveraging an offline tree-based partitioning of the input space. This way, CONFEX generates CFXs with rigorous guarantees on both predictive uncertainty and optimality. We evaluate CONFEX against state-of-the-art methods across diverse benchmarks and metrics, **demonstrating that in many cases, our approach more robust and plausible explanations compared to competing uncertainty-aware generators.**

## 1 INTRODUCTION

Machine learning models are deployed in high-stakes decision-making scenarios like loan approvals, medical diagnoses, and employment screening. In these contexts, algorithmic recourse—providing actionable feedback to individuals influenced by these decisions—is not just a technical concern but also an ethical and legal imperative. Although the legal status of “right to explanations” under the EU’s General Data Protection Regulation (GDPR) remains contested (Wachter et al., 2017; Selbst & Barocas, 2018), there is growing consensus that individuals should be offered meaningful information about algorithmic decisions that impact them (Edwards & Veale, 2017; Binns et al., 2018).

Counterfactual explanations (CFX) were formally introduced by Wachter et al. (2017) as a method for algorithmic recourse. CFXs answer questions like: “What minimal changes to my input features would have altered the model’s decision desirably?”, and Wachter’s formalisation focuses on finding counterfactual explanations that are minimally close to the original point (*factual instance*) or have sparse feature changes. These criteria of closeness and sparseness have been extended in later methods to other desiderata such as diversity, causality, actionability, and plausibility, to generate explanations that work better as a recourse path and are distinguished from adversarial examples.

However, most existing CFX methods fail to account for the inherent uncertainty in both data and model predictions. This is problematic because explanations that ignore uncertainty may lead to false confidence in suggested changes, potentially resulting in ineffective recourse actions when deployed in practice. Uncertainty quantification in CFX is thus crucial for generating reliable and actionable insights.

We introduce CONFEX, an uncertainty-aware CFX generator that builds on *Conformal Prediction (CP)* (Vovk et al., 2022; Angelopoulos et al., 2023). CP is a popular uncertainty quantification framework that offers distribution-free and finite-sample coverage guarantees. It works by using calibration data to construct prediction regions that contain the true (unknown) outcome with a user-specified probability. CP does not require assumptions on the data distribution and the underlying model, except that the calibration data and the test point must be exchangeable. The core idea of our CONFEX method is to constrain the search space for CFXs only to those points leading to a

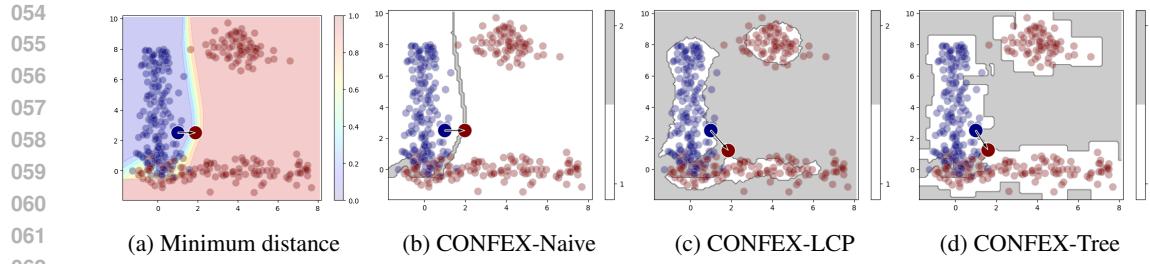


Figure 1: Counterfactuals produced for the same factual instance (marked in blue) for a MLP classifier using approaches MILP-MinDist, CONFEX-Naive, CONFEX-LCP, CONFEX-Tree. CONFEX approaches use bandwidth as 35% of the median pairwise distance between calibration points, and alpha as 2%.

singleton prediction region  $\{y^+\}$ , i.e., points that yield the desired outcome  $y^+$  with a high degree of certainty, since non-singleton CP regions represent uncertain predictions.

To illustrate our methods, Fig. 1a displays CFXs produced over a synthetic 2D dataset inspired from Poyiadzi et al. (2020). We can observe that counterfactuals produced by the minimal distance approach and by a naive application of CP to the CFX generation problem, called CONFEX-Naive (Section 3), fail to be plausible with respect to the data distribution.

These issues with naively applying CP to CFX generation stem from the fact that the generated (test-time) CFX may not be exchangeable with the calibration points, thereby affecting the validity of CP’s guarantees. We solve this by imposing stricter coverage requirements for CP: we build prediction regions that approximately<sup>1</sup> attain *local* (aka *test-conditional*) guarantees, i.e., the target coverage probability is achieved for *any* test point. In contrast, normally, CP guarantees are marginal, i.e., the coverage probability is averaged over the joint calibration and test distribution.

Our CONFEX method relies on a *Mixed-Integer Linear Programming (MILP)* encoding of the optimisation problem, which not only guarantees optimality of solutions but also ensures satisfaction of the CP constraints. We present two methods for incorporating local coverage constraints. The first is *localised CP* (Guan, 2023), which frames conditional coverage as a covariate shift problem (Tibshirani et al., 2019). However, it requires encoding and solving calibration quantiles in MILP, which is computationally expensive and scales poorly with the dataset size. The second, more efficient, method is a KD-tree-based encoding of local calibration quantiles. For this method, we use regression trees, which can be efficiently encoded in MILP.

In summary, our main contributions are:

1. a mathematical formulation for distribution-free uncertainty-aware counterfactual explanations, the first to apply conformal prediction in a principled manner (i.e., by addressing the exchangeability problem via test-conditional coverage, **retaining formal guarantees**);
2. a novel localised CP procedure, with an efficient MILP encoding, for generation of CFXs, which can be used more generally to incorporate (test-conditional) CP uncertainty constraints in any search problem;
3. an extensive experimental evaluation demonstrating that our CONFEX method outperforms competing **uncertainty-aware** generators by providing more **certain**, plausible and stable explanations, as well as enjoying formal guarantees on uncertainty.

## 2 BACKGROUND AND PROBLEM FORMULATION

**Counterfactual Explanations** Let  $\hat{f} : \mathcal{X} \rightarrow \mathcal{Y}$  denote a trained classifier for which we seek to generate counterfactual explanations. Given an instance  $x_0 \in \mathcal{X}$  such that  $\hat{f}(x_0) \neq y^+$ , the goal is to identify a counterfactual instance  $x'$  such that  $\hat{f}(x') = y^+$ . Wachter et al. (2017) frame this as an

<sup>1</sup>Exact conditional guarantees for CP are known to be impossible unless the inputs are discrete (Vovk, 2012; Barber et al., 2020).

108 optimisation problem and solve it via gradient descent.

$$110 \quad x_{\text{cf}} \in \arg \min_{x'} \max_{\lambda} \left( \lambda \text{yloss} \left( \hat{f} (x'), y^+ \right) + \text{dist} (x_0, x') \right). \quad (1)$$

112 The loss function aims to find an explanation that changes the predicted class to the target class  
 113 (first term), while also ensuring that the explanation is close to the input instance (second term).  
 114 Closeness is often defined as an  $L_p$  norm, which can be weighted based on the observed data (e.g. the  
 115 inverse median absolute deviation), or to reflect domain knowledge (Dandl et al., 2020). However,  
 116 by optimising solely for closeness, this formulation often leads to counterfactual explanations that  
 117 resemble adversarial examples and may not be actionable or robust.

118 Desirable properties of CFXs include *validity* (prediction flips to  $y^+$ ), *proximity* (closeness to the  
 119 factual instance), *sparsity* (few feature changes), *plausibility* (realistic and likely under the data  
 120 distribution), *actionability* (only mutable features are altered), *causality* (identified counterfactual  
 121 satisfies causal relationships) and *robustness* (stability under input perturbations); see (Verma et al.,  
 122 2020; Karimi et al., 2021).

123 Uncertainty-aware CFX methods show promise for enhancing the robustness and plausibility of  
 124 CFXs. In this line of work, Schut et al. (2021) propose minimising predictive entropy across an  
 125 ensemble of models to consider the effect of uncertain regions. Bayesian approaches, such as CLUE  
 126 (Antorán et al., 2020), leverage predictive uncertainty from Bayesian neural networks to generate  
 127 epistemically informative counterfactuals.

128 **Conformal Prediction** CP is a distribution-free inference framework that complements any pre-  
 129 dictive model with rigorous uncertainty quantification. CP outputs prediction sets guaranteed to  
 130 contain the true (unknown) outcome with a user-specified probability  $1 - \alpha$  without relying on  
 131 asymptotic or parametric assumptions (Vovk et al., 2022; Angelopoulos et al., 2023). To construct  
 132 these sets, CP performs the following steps:

133 1. **Calibration:** use a held-out calibration dataset  $\mathcal{D}_{\text{cal}} = \{(x_i, y_i)\}_{i=1}^n$  to find the critical value  
 134  $q_{1-\alpha}$  (i.e., the  $1 - \alpha$  quantile) of a chosen test statistic called the (*non-conformity*) *score*  $s(x, y)$ ,  
 135 which is normally chosen to quantify the deviation between the model prediction  $\hat{f}(x)$  and the  
 136 ground truth  $y$ . This step is performed only once, offline. Formally,

$$139 \quad q_{1-\alpha} = Q_{1-\alpha} \left( \sum_{i=1}^n \frac{1}{n+1} \delta_{s(x_i, y_i)} + \frac{1}{n+1} \delta_{+\infty} \right), \quad (2)$$

140 where  $Q_{1-\alpha}$  is the  $1 - \alpha$  quantile function and  $\delta_v$  is the Dirac distribution centered at  $v$ .

141 2. **Inference:** for a test input  $x^*$ , construct a prediction region  $C(x^*)$  by including all labels  $y$  whose  
 142 score is below the critical value (i.e., such that  $s(x^*, y) \leq q_{1-\alpha}$ ).

143 The CP procedure provides the following marginal guarantee for an unseen test point  $(x^*, y^*)$ :

$$147 \quad \mathbb{P}_{\mathcal{D}_{\text{cal}}, (x^*, y^*)} (y^* \in C_{1-\alpha}(x^*)) \geq 1 - \alpha. \quad (3)$$

148 The above holds in finite sample regimes (as opposed to asymptotic) under the mild condition of ex-  
 149 changeability (a weaker assumption than IID), i.e., the joint distribution of calibration and test points  
 150 is invariant under permutations. By marginal guarantees, we mean that the coverage probability of  
 151 equation 3 is achieved on average over the joint calibration and test distribution.

153 **CP and CFXs** To our knowledge, there exist only two methods which apply conformal prediction  
 154 to CFX generation: ECCCo (Altmeyer et al., 2024) and CPICF (Adams et al., 2025).

155 CPICF (Adams et al., 2025) assumes an alternative “individualised” setting, where an institution  
 156 holds a private black-box classifier and aims to provide CFXs to individuals without disclosing the  
 157 classifier. The knowledge of each individual is modelled by their own classifier, and the organisation  
 158 produces a CFX to reduce uncertainty in the global classifier via CP. This is a fundamentally different  
 159 setting to ours, furthermore CPICF’s formulation does not retain any formal CP guarantees.

160 In the standard setting, ECCCo extends Wachter’s formulation (equation 1) with two additional  
 161 terms: one that optimises the energy of the identified counterfactual to enhance plausibility, and

162 one that minimises uncertainty through the smooth conformal set size loss of Stutz et al. (2022).  
 163 However, ECCCo has the following drawbacks: 1) it incorporates conformal prediction, but in a  
 164 way that does not address exchangeability issues, which we detail in Section 3.1; 2) the procedure  
 165 does not guarantee CP regions will have the required size (e.g., singletons); 3) it relies on energy-  
 166 based training to obtain plausible CFXs. As we will show, our approach instead induces plausible  
 167 CFXs solely by using CP constraints, formulating these constraints to enforce local validity (thereby  
 168 solving the exchangeability issues), and thanks to the MILP formulation, it ensures satisfaction of  
 169 the set size constraints whilst being optimally close.

170 Both ECCCo and CPICF fail to retain formal guarantees on generated counterfactuals, and mention  
 171 that further analysis on the role of CP in CFXs is required. In the context of recourse recommendations,  
 172 without uncertainty guarantees, the CFX method may suggest CFXs where the model is  
 173 uncertain. This may mislead the recipient into making changes that do not actually alter the out-  
 174 come. Instead we want to produce reliable explanations for every individual, backed up by formal  
 175 guarantees.

176 **Mixed Integer Linear Programming (MILP) and CFXs** MILP provides a framework for for-  
 177 mulating and deriving CFXs as a constraint-solving problem. The problem is of finding a point  $x'$   
 178 which minimises the distance to the original instance  $x_0$  whilst being classified as  $y^+$ .

$$181 \quad x_{\text{cf}} \in \arg \min_{x'} \text{dist}(x_0, x') \quad \text{s.t. } \hat{f}(x') = y^+ \quad (4)$$

183 We refer to this method as MILP-MinDist, and it serves as a baseline for our CONFEX method.  
 184

185 To allow the encoding, the model  $f$  must be representable in MILP; this is the case for e.g. linear  
 186 classifiers and multilayer perceptrons with ReLU activations, as well as non-differentiable models  
 187 such as decision trees. Neural network layers like sigmoid or softmax are not linearly representable,  
 188 but can be omitted from the MILP encoding if used at the last layer since we can identify if  $f(x_{\text{cf}}) =$   
 189  $y^+$  based on the logits alone.

190 When presented to an MILP solver, this approach is guaranteed to yield a valid and optimal CFX,  
 191 if such an explanation exists. Gradient-based methods, on the other hand, are incomplete, meaning  
 192 that they may fail to find valid CFXs or may return suboptimal solutions.

193 We note that properties like causality and actionability can be incorporated in equation 4 through  
 194 MILP constraints on the input variables; similarly, a set of diverse explanations (as opposed to an  
 195 individual one) can be generated by repeatedly solving the problem and adding constraints or objec-  
 196 tive function terms to block or penalize explanations similar to those already identified (Kanamori  
 197 et al., 2020). By adding such constraints, our method can accommodate these desiderata as well.

198 **Problem Formulation** We aim to find CFXs that modify the factual input to the minimum extent  
 199 necessary to yield the desired label  $y^+$  with high probability. Below is the formal problem statement.

200 **Problem 1 (Uncertainty-aware CFX).** Given a factual input  $x_0$  and an error level  $\alpha \in (0, 1)$ , an  
 201 *uncertainty-aware counterfactual explanation*  $x_{\text{cf}}$  is a solution to the below optimisation problem:

$$203 \quad x_{\text{cf}} \in \arg \min_{x'} \text{dist}(x_0, x') \quad \text{s.t. } \mathbb{P}_{Y|X=x'}(Y = y^+) \geq 1 - \alpha, \quad (5)$$

205 where  $\mathbb{P}_{Y|X=x'}$  is the conditional distribution of labels  $Y$  given  $X = x'$ . Hence, our base method  
 206 targets just the minimal distance and low uncertainty desirable CFX properties.  
 207

### 209 3 CFXs WITH CP CONSTRAINTS: A NAIVE ATTEMPT

211 We first present a naive approach to apply conformal prediction to minimise the uncertainty in  
 212 the generated CFX, which we call CONFEX-Naive. This approach extends MILP-MinDist (see  
 213 equation 4) by restricting the search space to points yielding the singleton CP region  $\{y^+\}$ , i.e.,  
 214 points attaining the target class and with a high degree of certainty:

$$215 \quad x_{\text{cf}} \in \arg \min_{x'} \text{dist}(x_0, x') \quad \text{s.t. } C_{1-\alpha}(x') = \{y^+\} \quad (6)$$

216 Note that the above constraint is equivalent to the constraints  $s(x', y^+) \leq q_{1-\alpha}$  and  
 217  $\bigwedge_{y \neq y^+} s(x', y) > q_{1-\alpha}$ . The quantile  $q_{1-\alpha}$  is pre-computed on the held-out calibration set.  
 218

219 For multi-layer perceptrons, we use the following log-likelihood ratio as the score function

$$220 \quad s(x, y) = \log \left( \frac{\max_{y' \neq y} p(x)_{y'}}{p(x)_y} \right), \quad (7)$$

222 where  $p(x)_y$  is the softmax probability of  $y$  predicted by the model  $f$  for input  $x$ . When the correct  
 223 class is predicted, the ratio is below 1 and we obtain a negative score. When the model is wrong,  
 224 the ratio is positive and the score grows bigger as the model confidence on  $y$  decreases relative to  
 225 that on the predicted class. Importantly, equation 7 can be equivalently expressed in a linear form as  
 226  $s(x, y) = -l(x)_y + \max_{y' \neq y} l(x)_{y'}$ , where  $l(x)$  is the predicted vector of logits, making it efficiently  
 227 representable in MILP.

228 **Relation with MILP-MinDist** We note that our score function is well-formed, i.e.,  $s(x, y)$  is  
 229 lowest when  $y$  is the label predicted by the model  $f$  (and, in particular,  $s(x, y)$  increases as the  
 230 softmax probability of  $y$  decreases). Thus, when a CP prediction region returns the singleton  $\{y^+\}$ ,  
 231 then  $y^+$  is the class with the lowest score, i.e., the class predicted by  $f$ . That is, for any  $\alpha \in (0, 1)$ ,  
 232  $C_{1-\alpha}(x) = \{y^+\} \rightarrow f(x) = y^+$ . This implies that the feasible set of CONFEX is a subset of that  
 233 of MILP-MinDist, and so, CONFEX explanations can never attain smaller (better) distances than  
 234 CFX-base. Importantly, since the above property holds for any  $\alpha$ , it also holds for any choice of  
 235 quantile  $q_{1-\alpha}$ . This property also applies to the localised CP methods described later, which define  
 236 a different quantile value.

### 237 3.1 NEED FOR CONDITIONAL GUARANTEES

238 A visual example of using CONFEX-Naive to generate a counterfactual explanation is shown in  
 239 Figure 1 (plot b). We observe that when adding the singleton set size constraint, the obtained coun-  
 240 terfactual explanation is further from the decision boundary compared to MILP-MinDist (plot a). This  
 241 is is desirable since the identified CFX would resemble less an adversarial example. However,  
 242 the counterfactual explanation the identified CFX is somewhat counterintuitive: it lies in an area  
 243 without local datapoints, i.e., away from the data support (see plot d). Since the CP constraints  
 244 enforce low-uncertainty predictions, we would expect to find the CFX in a region where datapoints  
 245 unambiguously belong to the target class, and not in regions near the decision boundary, where  
 246 multiple classes overlap, or with no or little data support.

247 The main issue is that CONFEX-Naive can return CFXs that are not exchangeable with the calibra-  
 248 tion points, violating CP’s marginal guarantees. Hence, our prediction regions should be valid for  
 249 any choice of test inputs (not just exchangeable ones), requiring the coverage requirements to be  
 250 strengthened to enforce *conditional validity*, i.e., for any choice of  $x = x'$ , the following must hold:

$$252 \quad \mathbb{P}_{\mathcal{D}_{\text{cal}}, (x, y)} (y \in C_{1-\alpha}(x) \mid x = x') \geq 1 - \alpha. \quad (8)$$

253 However, unless the inputs are discrete, the above exact conditional guarantees are known to be  
 254 impossible if we require distribution-free and finite-sample guarantees (Vovk, 2012; Barber et al.,  
 255 2020). To solve this issue, among the several methods recently proposed for CP with approximate  
 256 conditional validity (Jung et al., 2022; Hore & Barber, 2023; Ding et al., 2023; Gibbs et al., 2025;  
 257 Cabezas et al., 2025), we focus on the *localised CP (LCP)* method of Guan (2023), described in  
 258 the next section. Below, we prove that conditional conformal prediction provides a solution to the  
 259 uncertainty-aware CFX problem stated in Problem 1.

260 **Uncertainty-aware CFXs with Conditional Conformal Prediction** Consider the set of accept-  
 261 able points  $A = \{x \mid C_{1-\alpha}(x) = \{y^+\}\}$ . If  $C_{1-\alpha}$  satisfies the conditional guarantees of Equation 8,  
 262 then for every  $x \in A$ , we have that

$$263 \quad \mathbb{P}_{\mathcal{D}_{\text{cal}}, Y \mid X=x} (Y = y^+) \geq 1 - \alpha, \quad (9)$$

264 which is equivalent to stating  $\mathbb{P}_{\mathcal{D}_{\text{cal}}, X, Y} (Y = y^+ \mid X = x) \geq 1 - \alpha$ . This trivially follows  
 265 from the fact that if  $x \in A$ , then  $C_{1-\alpha}(x) = \{y^+\}$ , and so  $\mathbb{P}_{\mathcal{D}_{\text{cal}}, X, Y} (Y = y^+ \mid X = x) =$   
 266  $\mathbb{P}_{\mathcal{D}_{\text{cal}}, X, Y} (Y \in C_{1-\alpha}(x) \mid X = x)$ .

270 Thus, in our optimisation problem, we can use the constraint  $C_{1-\alpha}(x) = \{y^+\}$  (provided  $C_{1-\alpha}$   
 271 offers conditional guarantees) to ensure satisfaction of the uncertainty constraint (9). Note that (9)  
 272 is equal to the chance constraint (5) in Problem 1 except that the probability is over  $\mathcal{D}_{\text{cal}}$  too.  
 273

274 On the other hand, CONFEX-Naive uses standard CP, where  $C_{1-\alpha}$  offers only marginal coverage,  
 275 meaning that the constraint  $C_{1-\alpha}(x) = \{y^+\}$  does not satisfy Eq. (9) but a weaker form of it:  
 276

$$\mathbb{P}_{\mathcal{D}_{\text{cal}}, X, Y}(Y = y^+ \mid X \in A) \geq 1 - \alpha / P(A), \quad (10)$$

277 where  $P(A) = \mathbb{P}_{\mathcal{D}_{\text{cal}}, X, Y}(X \in A)$  is the probability that  $X$  yields a the singleton region  $\{y^+\}$ <sup>2</sup>.  
 278 Therefore, with marginal CP, we cannot attain a principled uncertainty control in CFX generation.  
 279

## 281 4 THE CONFEX APPROACH

284 Our method CONFEX uses Localised Conformal Prediction (LCP) to generate CFXs with more  
 285 principled, local coverage guarantees. We introduce two variants: CONFEX-LCP, which encodes  
 286 LCP constraints via MILP, and CONFEX-Tree, which also provides local guarantees via MILP but  
 287 is more computationally efficient thanks to an offline tree-based representation of the local quantiles.  
 288

### 289 4.1 LOCALISED CONFORMAL PREDICTION (LCP) AND CONFEX-LCP

290 Localised Conformal Prediction (LCP) (Guan, 2023) relaxes strict conditional coverage (see equation  
 291 8) by requiring coverage to hold only within a local neighbourhood around a test input  $x^*$ . To  
 292 achieve this, LCP reweights the calibration points as if they were drawn under the localised  
 293 distribution of  $x^*$ , thereby restoring exchangeability. The reweighted probabilities are computed by a  
 294 *localiser kernel*  $H : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ , which measures how “close”  $x'$  is to  $x$ , with  $H(x, x) = 1$ . In  
 295 our method, we use the  $L_1$ -box kernel

$$H(x, x') = \mathbf{1}(\|x - x'\|_1 \leq h), \quad (11)$$

296 where  $h$  is the kernel bandwidth controlling the degree of localisation. For numerical and ordinal  
 297 features, the  $L_1$  distance is computed after normalisation; for categorical features, we require exact  
 298 matches over all or some categorical features, else  $H(x, x') = 0$ . Other kernels (e.g., based on  
 299 infinity norm or Gaussian smoothing) are also possible.  
 300

301 For a test input  $x^*$ , the local quantile is

$$q_{1-\alpha}^{\text{LCP}}(x^*) = Q_{1-\alpha} \left( \sum_{i=1}^n w_i \delta_{s(x_i, y_i)} + w^* \delta_{+\infty} \right), \quad (12)$$

302 where  $w_i = \frac{H(x^*, x_i)}{W}$  for  $i = 1, \dots, n$  and  $w^* = \frac{H(x^*, x^*)}{W} = \frac{1}{W}$ , with  $W = 1 + \sum_{i=1}^n H(x^*, x_i)$   
 303 being a normalizing factor.  
 304

305 This reweighting step and the resulting prediction region  $C_{1-\alpha}^{\text{LCP}}(x^*) = \{y : s(x^*, y) \leq q_{1-\alpha}^{\text{LCP}}(x^*)\}$   
 306 ensure, for any test point  $x^*$ , the following approximate conditional guarantee:  
 307

$$\mathbb{P}_{\mathcal{D}_{\text{cal}} \sim P_{X, Y}^n, (x, y) \sim P_{X, Y}^*}(y \in C_{\text{LCP}, 1-\alpha}(x)) \geq 1 - \alpha, \quad (13)$$

308 where  $P_{X, Y}^n$  is the (product) distribution of the  $n$  calibration points, and  $P_{X, Y}^* = P_{Y|X} \times P_X^*$  is  
 309 the localised test distribution, with  $P_X^* = P_X \circ H(x^*, X)$  being the distribution of  $X$  obtained by  
 310 applying to  $P_X$  the kernel  $H$  centered at  $x^*$ .  
 311

312 **CONFEX-LCP** We extend CONFEX-Naive by replacing CP regions with LCP regions, yielding  
 313 more principled and adaptive counterfactual generation. Formally,  
 314

$$x_{\text{cf}} \in \arg \min_{x'} \text{dist}(x_0, x') \quad \text{s.t. } C_{1-\alpha}^{\text{LCP}}(x') = \{y^+\}, \quad (14)$$

315 <sup>2</sup>The proof is based on rewriting  $\mathbb{P}(Y \neq y^+ \mid X \in A) = 1 - \mathbb{P}(Y = y^+ \mid X \in A)$  as  
 316  $\frac{\mathbb{P}(Y \neq y^+ \wedge X \in A)}{\mathbb{P}(X \in A)}$  and noticing that the numerator is bounded by  $\alpha$ .  
 317

324 which enforces  $s(x', y^+) \leq q_{1-\alpha}^{LCP}(x')$  and  $s(x', y) > q_{1-\alpha}^{LCP}(x')$  for all  $y \neq y^+$ . Unlike CONFEX-  
 325 Naive, which uses a single global quantile  $\hat{q}$ , here the quantile depends on the candidate  $x'$ , requiring  
 326 explicit encoding in the MILP formulation (see Algorithm 2 in the Appendix). This introduces  
 327 additional variables and big-M constraints linear in the calibration set size. Fig. 1 (plot c) shows a  
 328 CFX computed using CONFEX-LCP.

330 **Properties.** Thanks to the LCP method, CONFEX-LCP computes quantiles using only points lo-  
 331 cal to the test input  $x$ , where locality is defined by the L1 kernel. This yields more adaptive and  
 332 reliable uncertainty estimates than vanilla CP (and CONFEX-Naive), with larger prediction sets in  
 333 sparse or ambiguous regions, whilst ensuring that counterfactual is grounded with the data, i.e., sim-  
 334 ilar (local) individuals which are correctly predicted to be in the target class. We note that features in  
 335 the kernel can be assigned different weights based on domain knowledge. The choice of the kernel  
 336 bandwidth  $h$  is application-specific and it allows us to balance between local and marginal coverage.

#### 337 4.2 CONFEX-TREE: FAST VARIANT OF CONFEX-LCP

339 Due to the increased cost of resolving quantiles using MILP, LCP is infeasible for practical use with  
 340 large calibration sets.

341 In this section, we introduce CONFEX-Tree, an efficient alternative formulation of Localised CP  
 342 which retains formal guarantees. CONFEX-Tree leverages that decision trees are efficiently rep-  
 343 resentable in MILP and uses precomputed local quantiles. While LCP operates at test-time by  
 344 retaining only the calibration points within distance  $h$  of the point, CONFEX-Tree works offline to  
 345 determine locality constraints: it splits the feature space recursively to obtain local neighbourhoods  
 346 of calibration points having kernel width of at most  $h$ .

347 The construction procedure is inspired by kd-trees (Skrodzki, 2019) and detailed in Algorithm 1.  
 348 Each leaf specifies a precomputed local quantile using only calibration points within that leaf. From  
 349 these points, we also compute the midpoint of the smallest enclosing hyper-rectangle. The tree  
 350 construction ensures that no two points in a leaf can have a bigger  $L_\infty$  distance than the kernel  
 351 bandwidth  $h$ . Then, each new test point  $x'$  is assigned to a leaf of the tree and is associated with  
 352 the corresponding quantile if  $x'$  is within  $L_\infty$  distance of  $h/2$  from the midpoint, which means that  
 353 it is within distance of  $h$  from any calibration point of that leaf. To handle categorical features, we  
 354 stratify the dataset by each combination of (all or select) categorical values and generate a tree for  
 355 each stratum (which is equivalent to first splitting on all categorical features).

356 The resulting tree is encoded in MILP and used to provide the quantile value for the test point,  
 357 replacing the LCP regions from CONFEX-LCP. Formally, explanations are derived by solving

$$359 x_{\text{cf}} \in \arg \min_{x'} \text{dist}(x_0, x') \quad \text{s.t. } C_{1-\alpha}^{\text{Tree}}(x') = \{y^+\}, \quad (15)$$

361 where  $C_{1-\alpha}^{\text{Tree}}$  is constructed using the local tree-based quantiles returned by Algorithm 1.

363 **Properties of CONFEX-Tree.** The tree constructed by the CONFEX-Tree defines a partitioning  
 364 of the feature space into disjoint regions  $\{\mathcal{X}_g\}_{g \in \mathcal{G}}$ . Each  $g$  has an associated quantile value  $q_{1-\alpha,g}$   
 365 computed using only calibration points in  $g$ . This results in the following finite-sample group-  
 366 conditional coverage guarantee

$$367 \mathbb{P}(y \in C_{1-\alpha}^{\text{Tree}}(x^*) \mid x^* \in \mathcal{X}_g) \geq 1 - \alpha \quad \text{for all } g \in \mathcal{G}, \quad (16)$$

369 as per Vovk (2012). Note that our method overapproximates the group-conditional quantiles as it  
 370 assigns a quantile of  $\infty$  when  $x^*$  has  $L_\infty$  distance more than  $h/2$  from the midpoint of  $g$ . For this  
 371 reason, it still satisfies the above guarantee.

372 Moreover, by construction, the groups created by CONFEX-Tree are local regions of calibration  
 373 points in the feature space. Hence, we obtain an approximate conditional guarantee, as the tree  
 374 approximates the conditional quantile  $Q_{1-\alpha}(s|x)$  with the granularity of the approximation being  
 375 controlled by the bandwidth  $h$ .

376 Finally, CONFEX-Tree can be viewed as an instance of LCP using the following kernel

$$377 H(x, x') = \mathbf{1}(\|x - x'\|_\infty \leq h \wedge \exists g. x, x' \in \mathcal{X}_g), \quad (17)$$

378

---

**Algorithm 1:** CONFEX-Tree: Tree-based encoding of local quantiles

---

379

**Input** : Calibration set  $\mathcal{D}_{\text{cal}}$ , score function  $s$ , coverage level  $1 - \alpha$ , bandwidth  $h$ 

380

**Output:** Tree-based quantile encoding

381

**Categorical Stratification:**

382

1. Stratify the calibration dataset by each distinct combination of (all or some) categorical feature values.
2. Generate a tree for each group using the Tree Construction procedure over the normalised numerical and ordinal values only.

383

**Tree Construction:**

384

1. If the maximum range along any feature dimension of all calibration points in the node is less than  $h$ , stop and create a leaf node. At each leaf, compute and store:
  - the  $1 - \alpha$  quantile of the scores  $s(x, y)$  of the calibration points assigned to the leaf;
  - the midpoint of the calibration features in the leaf.
2. Otherwise, split the current node along the feature with the maximum spread, using the midpoint of that feature's values as the split point. Recurse on the left and right subsets to build subtrees.

385

**Prediction for test point  $x^*$ :**

386

1. Select the correct tree based on the test point's categorical values.
2. Traverse the tree using  $x'$  until reaching a leaf. Let  $c$  and  $q$  be its stored midpoint and quantile.
3. Reject point if assigned to the leaf but not local: if  $\|x^* - c\|_\infty > h/2$ , return  $\infty$ ; o/w, return  $q$ .

---

400

401

i.e., both points need to belong to the same leaf and have  $L_\infty$  distance bounded by  $h$ . Using this kernel, the guarantees of equation 13 also apply to CONFEX-Tree.

402

To summarise, explanations produced by CONFEX-Tree enjoy a distance optimality guarantee, validity guarantee (Relation with MILP-MinDist), and uncertainty guarantee Eq. (16). Additionally, our local (i.e., approx. conditional) guarantees imply that our CFXs are valid with high probability for any individual, even for out-of-distribution ones. This is preferable to a generator which, over a test set, empirically produces good results over a particular metric, since the distribution may shift at test-time - our method is robust to this.

403

404

## 5 EVALUATION

405

In this section, we evaluate our method against competing CFX methods, assessing the cost (distance), plausibility and sensitivity of CFXs generated by CONFEX-Tree. We explore the impact of varying the kernel bandwidth and the user-specified coverage rate, and we verify the formal coverage guarantees of CONFEX methods. We find that CONFEX consistently produces more stable and plausible CFXs across the benchmarks, provided the kernel bandwidth is appropriately chosen.

406

407

**Experimental setup** For our experiments, two classes of models are considered: multi-layer perceptrons (MLPs) and random forests (RFs). We selected four tabular datasets commonly found in the CFX literature: AdultIncome (Becker & Kohavi, 1996), CaliforniaHousing (Pace & Barry, 1997), GiveMeSomeCredit and GermanCredit (Hofmann, 1994), using a training-calibration-test split of 60%-20%-20% for each.

408

409

To evaluate CONFEX, we compare our efficient tree-based approach CONFEX-Tree (CTree) against competing uncertainty-aware generators: ECCCo (Altmeyer et al., 2024), the only other CFX method which uses CP, and a modified version of Schut (Schut et al., 2021) (called ‘Greedy’ in our table) which uses a single MLP instead of an ensemble, as well as the Wachter et al. (2017) baseline. We also consider plausibility-targeting generators FACE Poyiadzi et al. (2020) and C-CHVAE Pawelczyk et al. (2020). For tree-based models, we compare against the popular methods FeatureTweak (FT) (Tolomei et al., 2017), which searches for possible paths which can change the classification, and FOCUS (Lucic et al., 2021), which optimises for distance over a differentiable relaxation of the tree models. As baselines, we include MILP-MinDist (MinDist) and CONFEX-

	CaliforniaHousing			GermanCredit		
	Distance	Plausibility	Sens ( $10^{-1}$ )	Distance	Plausibility	Sens ( $10^{-1}$ )
<b>Multi-Layer Perceptron</b>						
MinDist	<b>0.03 ± 0.00</b>	$0.30 \pm 0.07$	$42.75 \pm 8.5$	$1.65 \pm 0.18$	$0.54 \pm 0.19$	$0.08 \pm 0.02$
ECCCo	$0.37 \pm 0.02$	$-0.65 \pm 0.05$	$0.26 \pm 0.05$	$0.97 \pm 0.06$	$0.16 \pm 0.11$	$0.06 \pm 0.02$
Greedy	$1.88 \pm 0.27$	$-0.99 \pm 0.02$	$0.14 \pm 0.02$	$0.99 \pm 0.04$	$-0.03 \pm 0.09$	$0.08 \pm 0.02$
Wachter	$0.09 \pm 0.01$	$0.42 \pm 0.08$	$1.66 \pm 0.37$	<b>0.41 ± 0.02</b>	$0.73 \pm 0.05$	$0.33 \pm 0.09$
FACE	$0.21 \pm 0.02$	<b>0.85 ± 0.04</b>	$0.34 \pm 0.03$	$0.69 \pm 0.05$	$0.92 \pm 0.05$	$0.05 \pm 0.01$
C-CHVAE	$1.27 \pm 0.22$	$-0.35 \pm 0.12$	$0.06 \pm 0.03$	$2.45 \pm 0.13$	$0.80 \pm 0.19$	$0.07 \pm 0.01$
CNaive	$0.04 \pm 0.01$	$0.24 \pm 0.08$	$8.41 \pm 2.26$	$2.00 \pm 0.07$	$0.16 \pm 0.28$	$0.05 \pm 0.01$
CTree	$0.55 \pm 0.04$	$0.75 \pm 0.05$	<b>0.05 ± 0.05</b>	$2.31 \pm 0.20$	<b>1.00 ± 0.00</b>	<b>0.01 ± 0.01</b>
<b>Random Forest</b>						
MinDist	<b>0.01 ± 0.00</b>	$0.37 \pm 0.07$	$89.15 \pm 90.9$	$1.65 \pm 0.06$	$0.52 \pm 0.17$	$0.09 \pm 0.01$
FT	$0.12 \pm 0.03$	$0.29 \pm 0.25$	$0.58 \pm 0.16$	$0.50 \pm 0.06$	$0.84 \pm 0.05$	$0.09 \pm 0.01$
FOCUS	$0.11 \pm 0.01$	$0.34 \pm 0.09$	$5.21 \pm 2.31$	<b>0.45 ± 0.14</b>	$0.83 \pm 0.02$	$0.58 \pm 0.25$
FACE	$0.17 \pm 0.01$	<b>0.81 ± 0.02</b>	$0.46 \pm 0.08$	$0.59 \pm 0.06$	$0.88 \pm 0.07$	$0.07 \pm 0.01$
CNaive	$0.03 \pm 0.01$	$0.42 \pm 0.07$	$12.34 \pm 2.94$	$1.62 \pm 0.08$	$0.63 \pm 0.10$	$0.09 \pm 0.01$
CTree	$0.19 \pm 0.02$	$0.61 \pm 0.10$	<b>0.40 ± 0.18</b>	$2.04 \pm 0.16$	<b>1.00 ± 0.00</b>	<b>0.01 ± 0.01</b>

Table 1: Results for CaliforniaHousing and GermanCredit datasets. The best result for each generator over its hyperparameters is reported. See Table 3 in Section A.2 for full results, further discussion, and  $p$ -values for significance of results.

Validity 58% for FT in CaliforniaHousing. For GermanCredit, validity is 50% for FT, 84% for Wachter, 82% for Schut, 84% for ECCCo, 74% for C-CHVAE. This explains why some methods seem to attain smaller distances than MinDist, which is always valid.

Naive (CNaive). As discussed previously, CONFEX-LCP is very expensive due to its “direct” (and inefficient) quantile encoding, hence, we did not conduct extensive experiments for it. Instead, we include a scalability analysis and comparison of CONFEX-LCP and CONFEX-Tree in Section C.

**Metrics** To evaluate the CFXs, we focus on two main dimensions: plausibility and sensitivity. *Plausibility* evaluates whether counterfactuals lie close to the data distribution, and is measured with the Local Outlier Factor (LOF) stratified per target class, with higher scores indicating more realistic examples. *Sensitivity* (Sens) captures robustness to small perturbations of the input instance  $x$ ; counterfactuals with low sensitivity remain consistent under such perturbations.

We run our experiments over five repeats: in each repeat, we train our models and for each model and generator, we compute metrics from 100 generated CFXs for factual points taken from the test set, plus an additional 100 for the sensitivity metric. The metrics obtained are then computed and averaged to ensure statistical reliability. We also record the distance, implausibility, stability, and validity of the method. Further details on the metrics and experimental setup can be found in the appendix.

**Evaluation of conformal guarantees** In the main setup, CFXs are generated for each test instance, but since their ground truth is unknown, coverage cannot be computed. We therefore run an additional simulated setup, identical to CONFEX in that it finds the *closest test point* whose CP region is a singleton comprising the target class. This way, true labels are known and we can compute the empirical coverage  $\mathbb{E}(1(y \in \hat{C}_{1-\alpha}(x)))$  over this resampling of the test set. We measure the gap between the observed coverage and the target  $1 - \alpha$ . Note that this resampling considers only CFX-like points and hence breaks exchangeability. So, we expect CONFEX-Naive to miss the coverage target and the localised procedures to fare better.

**Results discussion** In Table 1, we observe that CONFEX-Tree consistently outperforms competing uncertainty-aware methods by producing, in many cases more plausible and less sensitive explanations. This is in contrast to CONFEX-Naive which shows substantially lower plausibility and higher sensitivity, validating the issues illustrated in Figure 1 and further motivating the use

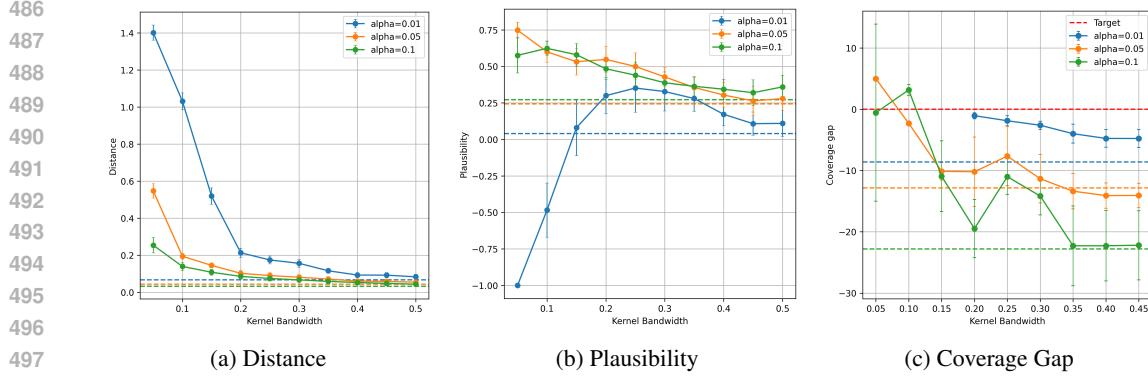


Figure 2: Effect of coverage rate and kernel bandwidth on metrics for CONFEX-Tree on the CaliforniaHousing dataset. CONFEX-Naive is represented by dashed horizontal lines.

of localisation in CP. In terms of plausibility, we find that CONFEX-Tree performs comparably to generator FACE and outperforms C-CHVAE and ECCCo, which explicitly target this metric. In the appendix, we show that our method provides more certain explanations than competing generators, and include results for the AdultIncome and GiveMeSomeCredit datasets.

Fig. 16 illustrates the effect of varying the kernel bandwidth and coverage rate in the CONFEX-Tree method. Increasing the coverage rate  $1 - \alpha$  leads to larger distances, since prediction sets become more conservative and singleton regions less frequent. Larger bandwidths yield shorter distances but at the cost of lower plausibility, as the notion of locality becomes weaker<sup>3</sup>. These observations are consistent with the fact that, as the kernel bandwidth grows, localised CP converges to standard marginal CP, as seen with CONFEX-Naive in the figures.

In the (simulated) CFX setting, the Coverage Gap results confirm that vanilla CP (used by CONFEX-Naive) fails to reach the target coverage, while localised CP with a suitably chosen kernel bandwidth succeeds. For small bandwidths (i.e., “strong” locality), all three choices of  $\alpha$  attain or are close to the target coverage level, but the gap grows as the bandwidth increases and localisation diminishes. For  $\alpha = 0.01$  and small bandwidths, no data is obtained since no test points produced a singleton prediction region (as required by our CONFEX constraints). These figures demonstrate that picking a correct bandwidth is crucial for obtaining good plausibility and coverage guarantees.

## 6 CONCLUSIONS

We introduced a novel MILP-based framework for generating uncertainty-aware counterfactual explanations with formal, distribution-free guarantees. By developing an efficient encoding of localised conformal prediction, we address the critical issue of exchangeability violation in the CFX search process. This allows us to enforce approximate test-conditional guarantees, ensuring the generation of provably reliable, plausible, and robust explanations.

**Limitations** Since our approach uses MILP to solve for CFXs, it will struggle scaling to very large models; gradient-based methods like Wachter and ECCCo are less prone to this problem, but they sacrifice guarantees on CFX validity. Additionally, unlike gradient-based methods, our method can be used on random forest and gradient-boosted trees, which remain competitive on tabular datasets.

Moreover, CP requires a held-out calibration dataset, which may be problematic when data is scarce. Fortunately, CP guarantees hold regardless of the calibration set size (but small sets will lead to more conservative prediction regions).

Picking an appropriate kernel bandwidth is an additional task which requires domain knowledge or evaluation on a validation set, for example, with the coverage gap simulation described in the previous section.

<sup>3</sup>For very small  $\alpha$  (0.01) and small kernel bandwidths, we observe low plausibility: we conjecture this could be due to the CP method localising on outlier points.

540 REFERENCES  
541

542 James M Adams, Gesine Reinert, Lukasz Szpruch, Carsten Maple, and Andrew Elliott. Individualised counter-  
543 factual examples using conformal prediction intervals. *arXiv preprint arXiv:2505.22326*, 2025.

544 Patrick Altmeyer, Mojtaba Farmanbar, Arie van Deursen, and Cynthia CS Liem. Faithful model explanations  
545 through energy-constrained conformal counterfactuals. In *Proceedings of the AAAI Conference on Artificial  
546 Intelligence*, volume 38, pp. 10829–10837, 2024.

547 Anastasios N Angelopoulos, Stephen Bates, et al. Conformal prediction: A gentle introduction. *Foundations  
548 and Trends® in Machine Learning*, 16(4):494–591, 2023.

549 Javier Antorán, Umang Bhatt, Tameem Adel, Adrian Weller, and José Miguel Hernández-Lobato. Getting a  
550 clue: A method for explaining uncertainty estimates. *arXiv preprint arXiv:2006.06848*, 2020.

551 Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. The limits of distribution-  
552 free conditional predictive inference, 2020. URL <https://arxiv.org/abs/1903.04684>.

553 Rina Foygel Barber, Emmanuel J Candès, Aaditya Ramdas, and Ryan J Tibshirani. Conformal prediction  
554 beyond exchangeability. *The Annals of Statistics*, 51(2):816–845, 2023.

555 Barry Becker and Ronny Kohavi. Adult. UCI Machine Learning Repository, 1996.

556 Reuben Binns, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 'it's reducing a  
557 human being to a percentage': Perceptions of justice in algorithmic decisions. In *CHI*, pp. 377. ACM, 2018.

558 Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. Lof: identifying density-based local  
559 outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp.  
560 93–104, 2000.

561 Luben MC Cabezas, Mateus P Otto, Rafael Izbicki, and Rafael B Stern. Regression trees for fast and adaptive  
562 prediction intervals. *Information Sciences*, 686:121369, 2025.

563 Susanne Dandl, Christoph Molnar, Martin Binder, and Bernd Bischl. Multi-objective counterfactual expla-  
564 nations. In *International conference on parallel problem solving from nature*, pp. 448–469. Springer, 2020.

565 Tiffany Ding, Anastasios Angelopoulos, Stephen Bates, Michael Jordan, and Ryan J Tibshirani. Class-  
566 conditional conformal prediction with many classes. *Advances in neural information processing systems*,  
567 36:64555–64576, 2023.

568 Sanghamitra Dutta, Jason Long, Saumitra Mishra, Cecilia Tilli, and Daniele Magazzeni. Robust counterfactual  
569 explanations for tree-based ensembles. In *International conference on machine learning*, pp. 5742–5756.  
570 PMLR, 2022.

571 Lilian Edwards and Michael Veale. Slave to the algorithm? why a 'right to an explanation' is probably not the  
572 remedy you are looking for. volume 16, pp. 1–65, 2017.

573 Isaac Gibbs, John J Cherian, and Emmanuel J Candès. Conformal prediction with conditional guarantees.  
574 *Journal of the Royal Statistical Society Series B: Statistical Methodology*, pp. qkaf008, 2025.

575 Leying Guan. Localized conformal prediction: A generalized inference framework for conformal prediction.  
576 *Biometrika*, 110(1):33–50, 2023.

577 Gurobi, 2022. URL <https://github.com/Gurobi/gurobi-machinelearning>.

578 Hans Hofmann. Statlog (German Credit Data). UCI Machine Learning Repository, 1994.

579 Rohan Hore and Rina Foygel Barber. Conformal prediction with local weights: randomization enables local  
580 guarantees. *arXiv preprint arXiv:2310.07850*, 2023.

581 Christopher Jung, Georgy Noarov, Ramya Ramalingam, and Aaron Roth. Batch multivalid conformal predic-  
582 tion. *arXiv preprint arXiv:2209.15145*, 2022.

583 Kentaro Kanamori, Takuya Takagi, Ken Kobayashi, and Hiroki Arimura. Dace: Distribution-aware counterfac-  
584 tual explanation by mixed-integer linear optimization. In *IJCAI*, pp. 2855–2862, 2020.

585 Amir-Hossein Karimi, Gilles Barthe, Bernhard Schölkopf, and Isabel Valera. A survey of algorithmic recourse:  
586 Definitions, formulations, solutions, and prospects. In *Proceedings of the 2021 International Joint Confer-  
587 ence on Artificial Intelligence (IJCAI)*, 2021.

594 Ed Klotz, Feb 2021. URL [https://cdn.gurobi.com/wp-content/uploads/](https://cdn.gurobi.com/wp-content/uploads/595)  
 596 Models-with-Products-of-Binary-Variables.pdf.

597 Ana Lucic, Harrie Oosterhuis, Hinda Haned, and Maarten de Rijke. Focus: Flexible optimizable counterfactual  
 598 explanations for tree ensembles, 2021. URL <https://arxiv.org/abs/1911.12199>.

599 Kyosuke Morita, 2023. URL <https://github.com/kyosek/CFXplorer>.

600 R. Kelley Pace and Ronald Barry. Sparse spatial autoregressions. *Statistics & Probability Letters*, 33(3):  
 601 291–297, 1997.

602 Martin Pawelczyk, Klaus Broelmann, and Gjergji Kasneci. Learning model-agnostic counterfactual expla-  
 603 nations for tabular data. In *Proceedings of the web conference 2020*, pp. 3126–3132, 2020.

604 Martin Pawelczyk, Sascha Bielawski, Johannes van den Heuvel, Tobias Richter, and Gjergji Kasneci. Carla: A  
 605 python library to benchmark algorithmic recourse and counterfactual explanation algorithms, 2021.

606 Rafael Poyiadzi, Kacper Sokol, Raul Santos-Rodriguez, Tijl De Bie, and Peter Flach. Face: Feasible and  
 607 actionable counterfactual explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and  
 608 Society*, pp. 344–350, 2020.

609 Lisa Schut, Oscar Key, Rory McGrath, Luca Costabello, Bogdan Sacaleanu, Medb Corcoran, and Yarin Gal.  
 610 Generating interpretable counterfactual explanations by implicit minimisation of epistemic and aleatoric  
 611 uncertainties. In *AISTATS*, volume 130 of *Proceedings of Machine Learning Research*, pp. 1756–1764.  
 612 PMLR, 2021.

613 Andrew D. Selbst and Solon Barocas. The intuitive appeal of explainable machines. *Fordham Law Review*, 87:  
 614 1085, 2018.

615 Martin Skrodzki. The kd tree data structure and a proof for neighborhood computation in expected logarithmic  
 616 time. *arXiv preprint arXiv:1903.04936*, 2019.

617 David Stutz, Krishnamurthy, Dvijotham, Ali Taylan Cemgil, and Arnaud Doucet. Learning optimal conformal  
 618 classifiers, 2022. URL <https://arxiv.org/abs/2110.09192>.

619 Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under  
 620 covariate shift. *Advances in neural information processing systems*, 32, 2019.

621 Gabriele Tolomei, Fabrizio Silvestri, Andrew Haines, and Mounia Lalmas. Interpretable predictions of tree-  
 622 based ensembles via actionable feature tweaking. In *Proceedings of the 23rd ACM SIGKDD International  
 623 Conference on Knowledge Discovery and Data Mining*, KDD ’17, pp. 465–474. ACM, August 2017. doi:  
 624 10.1145/3097983.3098039. URL <http://dx.doi.org/10.1145/3097983.3098039>.

625 Srikumar Verma, John P Dickerson, and Kevin Hines. Counterfactual explanations for machine learning: A  
 626 review. *arXiv preprint arXiv:2010.10596*, 2020.

627 Vladimir Vovk. Conditional validity of inductive conformal predictors. In *Asian conference on machine learn-  
 628 ing*, pp. 475–490. PMLR, 2012.

629 Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer  
 630 Nature, 2022.

631 Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black  
 632 box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31:841, 2017.

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648  
649  

## A APPENDIX

650  
651  
**AI Use Declaration** The authors acknowledge the use of Generative AI to minimally polish text.652  
653  

### A.1 RELATED WORKS

654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
Our work integrates three research areas: counterfactual explanations (CFXs), uncertainty quantification in explanations, and the application of conformal prediction (CP) to optimization problems. Counterfactual explanations, introduced by Wachter et al. (2017), provide recourse by identifying minimal feature changes to alter a model’s prediction. While initial work focused on validity and distance, the field has expanded to include desiderata like plausibility and actionability (Verma et al., 2020; Karimi et al., 2021). Methodologies have also diversified from gradient-based optimization to tree-specific algorithms (Tolomei et al., 2017; Lucic et al., 2021) and constraint-based methods using Mixed-Integer Linear Programming (MILP) (Kanamori et al., 2020). However, a critical limitation of many approaches is their failure to account for model uncertainty, which can result in misleading or brittle explanations (Schut et al., 2021). To address this, prior works have employed Bayesian methods (Antorán et al., 2020) or model ensembles (Schut et al., 2021). CONFEX contributes a novel, principled alternative by using Conformal Prediction. More relevant is ECCCo (Altmeyer et al., 2024), which uses a loss term based on the conformal set size (Stutz et al., 2022) but crucially does not address the violation of the exchangeability assumption inherent in the CFX search process.668  
669  

### A.2 FURTHER DISCUSSION OF TABLE 1

670  
671  
672  
673  
674  
675  
676  
677  
For GermanCredit, whilst Wachter obtained the closest counterfactuals, had a validity rate of 84%, demonstrating how gradient-based methods may fail to correctly change prediction to the target class. ECCCo (84%) and FeatureTweak (50%) also suffered validity issues. On the other hand, MILP-MinDistalways found a valid counterfactual, including satisfying correct categorical and ordinal encoding unlike some of the competing tree generators, and this is reflected with an increased distance. Note that in all figures, kernel bandwidth is measured as a multiple of the median pairwise distance between all points in the dataset.678  
679  
Extended versions of Table 1 are included as the following two tables.680  
681  
682  
683  
In each cell, the first line is the mean value of the metric over 5 repeats, the second line is the standard deviation over those repeats, and the third line is the the  $p$ -value for a two-tailed (paired per repeat) t-test to check whether the mean value of the generator’s metric is significantly different from the mean value of the best performing generator. Note that the values in the table, including  $p$ -values, are over valid results only: see the list below.684  
685  
686  
687  
688  
689  
690  
691  

- For CaliforniaHousing, validity is 58% for FeatureTweak.
- For GermanCredit, validity is 50% for FT, 84% for Wachter, 82% for Schut, 84% for ECCCo, 74% for C-CHVAE.
- For GiveMeSomeCredit, validity is 71% for Wachter, 80% for Schut, 50% for FeatureTweak.
- For AdultIncome, Validity 80% for Wachter, 85% for ECCCo, 54% for C-CHVAE, 61% for FeatureTweak

692  
693  
694  
695  
696  
697  
698  
699  
700  
701  
Full results, including details on the Certainty metric “Cert” and for full results.

		CaliforniaHousing				GermanCredit			
		Dist	Plaus	Sens	Cert	Dist	Plaus	Sens	Cert
<b>Multi-Layer Perceptron</b>									
		<b>0.03</b>	0.30	42.75	2.44e-4	1.65	0.54	0.08	0.140
702	MinDist	$\pm 0.00$	$\pm 0.07$	$\pm 8.45$	$\pm 2.62e-4$	$\pm 0.18$	$\pm 0.19$	$\pm 0.02$	$\pm 1.23e-2$
703		(-)	(2.81e-5)	(5.47e-4)	(4.83e-3)	(1.19e-4)	(8.54e-3)	(5.03e-3)	(1.08e-4)
704		0.37	-0.65	0.26	0.00e+00	0.97	0.16	0.06	0.173
705	ECCCo	$\pm 0.02$	$\pm 0.05$	$\pm 0.05$	$\pm 0.00e+00$	$\pm 0.06$	$\pm 0.11$	$\pm 0.02$	$\pm 4.51e-2$
706		(9.44e-7)	(1.14e-6)	(4.39e-3)	(4.78e-3)	(9.72e-6)	(1.16e-4)	(7.32e-3)	(2.77e-3)
707		1.88	-0.99	0.14	0.00e+00	0.99	-0.03	0.08	0.101
708	Greedy	$\pm 0.27$	$\pm 0.02$	$\pm 0.02$	$\pm 0.00e+00$	$\pm 0.04$	$\pm 0.09$	$\pm 0.02$	$\pm 3.58e-2$
709		(1.60e-4)	(1.52e-7)	(0.037)	(4.78e-3)	(5.00e-5)	(2.44e-5)	(2.91e-3)	(3.44e-4)
710		0.09	0.42	1.66	3.58e-2	<b>0.41</b>	0.73	0.33	-8.07e-3
711	Wachter	$\pm 0.01$	$\pm 0.08$	$\pm 0.37$	$\pm 1.24e-2$	$\pm 0.02$	$\pm 0.05$	$\pm 0.09$	$\pm 1.25e-2$
712		(7.98e-5)	(1.29e-4)	(8.20e-4)	(4.78e-3)	(-)	(3.76e-4)	(2.35e-3)	(3.19e-5)
713		0.21	<b>0.85</b>	0.34	5.40e-3	0.69	0.92	0.05	0.313
714	FACE	$\pm 0.02$	$\pm 0.04$	$\pm 0.03$	$\pm 4.46e-3$	$\pm 0.05$	$\pm 0.05$	$\pm 0.01$	$\pm 3.71e-2$
715		(4.23e-6)	(-)	(1.12e-3)	(5.32e-3)	(3.74e-4)	(4.07-2)	(1.26e-3)	(1.26e-2)
716		1.27	-0.35	0.06	0.00e+00	2.45	0.80	0.07	8.17e-2
717	C-CHVAE	$\pm 0.22$	$\pm 0.12$	$\pm 0.03$	$\pm 0.00e+00$	$\pm 0.13$	$\pm 0.19$	$\pm 0.01$	$\pm 5.95e-2$
718		(3.49e-4)	(3.37e-5)	(0.79)	(4.78e-3)	(5.28e-6)	(0.10)	(2.66e-4)	(2.41e-4)
719		0.04	0.24	8.41	3.26e-2	2.00	0.16	0.05	0.213
720	CNaive	$\pm 0.01$	$\pm 0.08$	$\pm 2.26$	$\pm 9.71e-3$	$\pm 0.07$	$\pm 0.28$	$\pm 0.01$	$\pm 3.03e-2$
721		(7.10e-4)	(2.68e-5)	(1.73e-3)	(6.67e-3)	(1.26e-6)	(3.79e-3)	(5.47e-3)	(2.76e-4)
722		0.55	0.75	<b>0.05</b>	<b>0.101</b>	2.31	<b>1.00</b>	<b>0.01</b>	<b>0.483</b>
723	CTree	$\pm 0.04$	$\pm 0.05$	$\pm 0.05$	$\pm 3.57e-2$	$\pm 0.20$	$\pm 0.00$	$\pm 0.01$	$\pm 4.42e-2$
724		(7.50e-6)	(3.00e-2)	(-)	(-)	(3.66e-5)	(-)	(-)	(-)
<b>Random Forest</b>									
		<b>0.01</b>	0.37	89.15	1.78e-2	1.65	0.52	0.09	4.84e-2
725	MinDist	$\pm 0.00$	$\pm 0.07$	$\pm 90.9$	$\pm 4.18e-3$	$\pm 0.06$	$\pm 0.17$	$\pm 0.01$	$\pm 2.14e-2$
726		(-)	(2.05e-4)	(0.12)	(2.45e-4)	(4.52e-5)	(4.88e-3)	(2.02e-4)	(2.37e-2)
727		0.12	0.29	0.58	4.68e-3	0.50	0.84	0.09	4.27e-2
728	FT	$\pm 0.03$	$\pm 0.25$	$\pm 0.16$	$\pm 6.00e-3$	$\pm 0.06$	$\pm 0.05$	$\pm 0.01$	$\pm 4.41e-2$
729		(1.27e-3)	(1.12e-2)	(0.29)	(1.32e-4)	(0.52)	(2.17e-3)	(6.76e-4)	(1.09e-2)
730		0.11	0.34	5.21	1.96e-2	<b>0.45</b>	0.83	0.58	0.134
731	FOCUS	$\pm 0.01$	$\pm 0.09$	$\pm 2.31$	$\pm 7.52e-3$	$\pm 0.14$	$\pm 0.02$	$\pm 0.25$	$\pm 7.82e-3$
732		(1.21e-4)	(6.20e-4)	(1.51e-2)	(2.07e-4)	(-)	(7.24e-5)	(1.11e-2)	(6.00e-2)
733		0.17	<b>0.81</b>	0.46	3.79e-2	0.59	0.88	0.07	0.04
734	FACE	$\pm 0.01$	$\pm 0.02$	$\pm 0.08$	$\pm 8.69e-3$	$\pm 0.06$	$\pm 0.07$	$\pm 0.01$	$\pm 0.01$
735		(8.05e-6)	(-)	(0.60)	(4.21e-4)	(4.59e-2)	(3.70e-2)	(3.81e-4)	(7.10e-1)
736		0.03	0.42	12.34	8.34e-2	1.62	0.63	0.09	0.175
737	CNaive	$\pm 0.01$	$\pm 0.07$	$\pm 2.94$	$\pm 1.54e-2$	$\pm 0.08$	$\pm 0.10$	$\pm 0.01$	$\pm 6.52e-2$
738		(3.42e-3)	(5.67e-4)	(1.34e-3)	(5.17e-3)	(1.05e-5)	(1.54e-3)	(4.94e-4)	(1.95e-1)
739		0.19	0.61	<b>0.40</b>	<b>0.143</b>	2.04	<b>1.00</b>	<b>0.01</b>	<b>0.335</b>
740	CTree	$\pm 0.02$	$\pm 0.10$	$\pm 0.18$	$\pm 1.83e-2$	$\pm 0.16$	$\pm 0.00$	$\pm 0.01$	$\pm 0.148$
741		(2.17e-5)	(1.71e-2)	(-)	(-)	(1.35e-4)	(-)	(-)	(-)

Table 2: [Extended version of Table 1](#). CFX generation results for CaliforniaHousing and GermanCredit, including mean and standard deviation of metric value over 5 runs, and  $p$ -value of a t-test to check whether the mean value of the generator's metric is significantly different from the mean value of the best performing generator. See Section A.2 for more detail, including on validities.

750  
751  
752  
753  
754  
755

756	757	GiveMeSomeCredit				AdultIncome																									
		758	Dist	Plaus	Sens	Cert	759	Dist	Plaus	Sens	Cert																				
		760	761	762	763	764	765	766	767	768	769																				
<b>Multi-Layer Perceptron</b>																															
770	771	772	773	774	775	776	777	778	779	780	781	782	783	784	785	786	787	788	789	790	791	792	793	794	795						
MinDist	(-) (3.12e-2)	(9.67e-3)	(6.42e-2)	(6.20e-5)	(1.04e-3)	(0.35)	(1.78e-3)	(-) (2.69e-8)	(7.09e-3)	(5.80e-2)	(2.69e-2)	(1.07e-3)	(0.719)	(1.92e-3)	(-) (4.53e-2)	(6.37e-3)	(8.13e-2)	(7.53e-2)	(2.64e-3)	(4.18e-3)	(0.374)	(1.20e-3)	(-) (3.42e-4)	(4.64e-3)	(1.51e-5)	(5.80e-2)	(-) (0.424)	(2.23e-2)	(7.85e-4)		
ECCCo	(8.87e-3)	(2.69e-8)	(7.09e-3)	(5.80e-2)	(2.69e-2)	(1.07e-3)	(0.719)	0.69	-0.97	0.21	0.00e+00	0.73	-0.05	0.05	0.13	-0.02	1.05	6.90e-3	0.95	0.02	24190	-5.43e-3	0.09	0.93	0.95	0.00e+00	<b>0.43</b>	0.28	0.21	-2.91e-3	
Greedy	(±0.28)	(±0.03)	(±0.06)	(±0.00e+00)	(±0.11)	(±0.07)	(±0.00)	(±0.07)	(±0.03)	(±0.06)	(±4.06e-3)	(±0.08)	(±0.10)	(±48381)	(±2.44e-2)	(±0.07)	(±0.02)	(±0.07)	(±0.05)	(±0.08)	(±1.03e-2)	(±0.08)	(±0.01)	(±0.09)	(±0.05)	(±0.02)	(±0.08)	(±1.03e-2)			
Wachter	(4.53e-2)	(6.37e-3)	(8.13e-2)	(7.53e-2)	(2.64e-3)	(4.18e-3)	(0.374)	(3.42e-4)	(4.64e-3)	(1.51e-5)	(5.80e-2)	(-) (0.424)	(2.23e-2)	(7.85e-4)	(±0.01)	(±0.02)	(±0.07)	(±0.00e+00)	(±0.09)	(±0.05)	(±0.08)	(±1.03e-2)	(±0.01)	(±0.02)	(±0.07)	(±0.05)	(±0.02)	(±0.08)	(±1.03e-2)		
FACE	(±0.01)	(±0.02)	(±0.04)	(±1.52e-2)	(±0.16)	(±0.12)	(±0.02)	(7.66e-6)	(3.49e-2)	(2.76e-4)	(8.98e-2)	(1.99e-4)	(-) (0.428)	(4.32e-2)	(±0.01)	(±0.02)	(±0.04)	(±0.00e+00)	(±0.16)	(±0.12)	(±0.02)	(±1.71e-2)	(±0.01)	(±0.02)	(±0.07)	(±0.05)	(±0.02)	(±0.08)	(±1.71e-2)		
C-CHVAE	(±0.08)	(±0.06)	(±0.02)	(±1.33e-3)	(±0.55)	(±0.15)	(±0.01)	(5.38e-6)	(3.69e-7)	(-) (5.86e-2)	(2.05e-5)	(0.847)	(0.448)	(6.90e-2)	(±0.08)	(±0.06)	(±0.11)	(±4.56e-3)	(±0.07)	(±0.04)	(±0.01)	(±2.71e-2)	(±0.08)	(±0.06)	(±0.02)	(±0.07)	(±0.02)	(±0.08)	(±2.71e-2)		
CNaive	(0.04)	(0.76)	(0.78)	(4.86e-3)	(1.24)	(-0.13)	(0.05)	(7.87e-4)	(2.08e-3)	(6.41e-4)	(6.28e-2)	(2.10e-4)	(1.90e-3)	(0.876)	(0.04)	(0.05)	(0.01)	(±4.56e-3)	(±0.07)	(±0.04)	(±0.01)	(±2.71e-2)	(0.04)	(0.05)	(0.01)	(0.04)	(0.02)	(0.05)	(0.01)		
CTree	(0.24)	<b>0.98</b>	(0.14)	<b>7.75e-2</b>	(1.78)	(-0.02)	<b>0.05</b>	(1.76e-3)	(-)	(1.41e-2)	(± 5.88e-2)	(-)	(2.98e-4)	(2.91e-2)	(-)	(±0.06)	(±0.01)	(±0.03)	(± 5.88e-2)	(±0.19)	(±0.14)	(±0.02)	(±5.06e-2)	(-)	(-)	(-)	(-)	(-)	(-)	(-)	
<b>Random Forest</b>																															
MinDist	(±0.00)	(±0.01)	(±35.89)	(±6.28e-3)	(±0.03)	(±0.07)	(±0.03)	(-)	(4.64e-3)	(5.81e-3)	(5.49e-4)	(7.56e-2)	(1.44e-3)	(5.97e-2)	(7.97e-4)	(±0.01)	(±0.02)	(±0.04)	(±35.89)	(±0.03)	(±0.07)	(±0.03)	(±1.28e-2)	(±0.01)	(±0.02)	(±0.07)	(±0.05)	(±0.02)	(±0.08)	(±1.28e-2)	
FT	(±0.01)	(±0.02)	(±0.16)	(±9.31e-3)	(±0.10)	(±0.11)	(±0.01)	(9.17e-4)	(1.61e-2)	(6.45e-4)	(3.55e-4)	(0.152)	(2.69e-2)	(0.862)	(4.68e-2)	(±0.01)	(±0.02)	(±0.04)	(±9.31e-3)	(±0.10)	(±0.11)	(±0.01)	(±1.22e-1)	(±0.01)	(±0.02)	(±0.07)	(±0.05)	(±0.02)	(±0.08)	(±1.22e-1)	
FOCUS	(±0.00)	(±0.05)	(±0.79)	(±7.81e-3)	(±0.34)	(±0.06)	(±0.04)	(4.38e-5)	(3.76e-2)	(1.29e-2)	(1.46e-3)	(-) (0.795)	(0.0196)	(8.34e-3)	(4.68e-2)	(±0.00)	(±0.05)	(±0.11)	(±7.81e-3)	(±0.34)	(±0.06)	(±0.04)	(±1.44e-2)	(±0.00)	(±0.02)	(±0.07)	(±0.05)	(±0.02)	(±0.08)	(±1.44e-2)	
FACE	(±0.01)	(±0.01)	(±0.08)	(±9.73e-3)	(±0.07)	(±0.07)	(±0.02)	(9.25e-6)	(3.41e-2)	(-) (1.02e-2)	(6.27e-3)	(-) (0.989)	(0.989)	(7.01e-2)	(2.21e-1)	(±0.01)	(±0.01)	(±0.08)	(±9.73e-3)	(±0.07)	(±0.07)	(±0.02)	(±2.63e-2)	(±0.01)	(±0.02)	(±0.07)	(±0.05)	(±0.02)	(±0.08)	(±2.63e-2)	
CNaive	(0.01)	(0.96)	(28.68)	(3.20e-2)	(0.97)	(0.00)	(0.12)	(1.50e-3)	(3.88e-3)	(8.76e-2)	(3.05e-3)	(1.62e-3)	(9.69e-4)	(0.108)	(0.169)	(±0.00)	(±0.01)	(±25.01)	(±1.38e-2)	(±0.09)	(±0.10)	(±0.02)	(±5.17e-2)	(±0.01)	(±0.02)	(±0.07)	(±0.05)	(±0.02)	(±0.08)	(±5.17e-2)	
CTree	(0.07)	<b>0.99</b>	(0.55)	<b>8.56e-2</b>	(1.56)	(0.08)	<b>0.06</b>	(6.85e-5)	(-)	(8.43e-2)	(±1.21e-2)	(-)	(2.85e-3)	(1.17e-2)	(-)	(±0.01)	(±0.01)	(±0.14)	(±1.21e-2)	(±0.21)	(±0.15)	(±0.03)	(±6.12e-2)	(-)	(-)	(-)	(-)	(-)	(-)	(-)	

Table 3: [Extended version of Table 1](#). CFX generation results for GiveMeSomeCredit and AdultIncome, including mean and standard deviation of metric value over 5 runs, and  $p$ -value of a t-test to check whether the mean value of the generator's metric is significantly different from the mean value of the best performing generator. See Section A.2 for more detail, including on validities.

800  
801  
802  
803  
804  
805  
806  
807  
808  
809

810 **B MILP FORMULATION DETAILS**  
811812 The full CONFEX model optimises the following problem:  
813

814  
815 
$$\begin{aligned} & \operatorname{argmin}_{x'} && \|x_0 - x'\|_1 \\ & \text{subject to} && \text{encoding validity constraints (C1-C5)} \\ & && \text{classifier constraints (C6)} \\ & && \text{conformal quantile constraints (Alg 1-2)} \\ & && \text{conformal singleton set constraints (C7-14)} \\ & && \text{optional further constraints (C15)} \end{aligned}$$
  
816  
817  
818  
819  
820  
821  
822

823 where  $x_0$  is the factual instance and  $x'$  is the counterfactual explanation returned by the optimisation.  
824825 **Encoding validity constraints** The obtained counterfactual explanation must follow correct nu-  
826 metric/categorical/ordinal encoding of the dataset.  
827828 Let the indices  $i = 0, 1, 2, N$  of the  $N$ -length vector  $x'$  be partitioned into  $I_{num}$ , indices of numeric  
829 variables,  $I_{ord}$ , indices of ordinally encoded variables, and  $I_{cat}^1, I_{cat}^2, \dots, I_{cat}^C$ , which are index  
830 groups for each of  $C$  one-hot categorically encoded variables. It is possible that some of these sets  
831 are empty.  
832833 For each numeric variable, we require that the variable is within bounds. Let  $l[i]$  and  $u[i]$  represent  
834 the lower and upper bound of feature  $i$ .  
835

836 
$$x'[i] \geq l[i] \quad \text{for all } i \in I_{num} \text{ where } l[i] \neq -\infty \quad (C1)$$

837 
$$x'[i] \leq u[i] \quad \text{for all } i \in I_{num} \text{ where } u[i] \neq +\infty \quad (C2)$$

838 For ordinal features, we must encode that  $x'[i] \in v[i]$ , where  $v[i]$  is the set of possible ordinal values  
839 that  $x'[i]$  can take.  
840841 Add  $|v[i]|$  binary indicator variables  $V_{i,1}, V_{i,2}, \dots, V_{i,|v[i]|}$  for each  $i \in I_{ord}$   
842 corresponding to possible ordinal values  $v_{i,1}, v_{i,2}, \dots, v_{i,|v[i]|}$   
843

844 
$$\sum_{j=1}^{|v[i]|} V_{i,j} = 1 \text{ for all } i \in I_{ord} \quad (C3)$$

845 
$$x'[i] = \sum_{j=1}^{|v[i]|} V_{i,j} v_{i,j} \text{ for all } i \in I_{ord} \quad (C4)$$

846 For each group of one-hot encoded categorical features  $I_{cat}^1, I_{cat}^2, \dots, I_{cat}^C$ , we must encode that  
847  $x'[i] = 1$  for one  $i$  in  $I_{cat}^c$  and  $x'[i] = 0$  for all  $j \in I_{cat}^c, j \neq i$ .  
848849 Add  $|I_{cat}^c|$  binary indicator variables  $C_{c,1}, C_{c,2}, \dots, C_{c,|I_{cat}^c|}$  for each  $c \in \{1, \dots, C\}$   
850 corresponding to each entry in the one-hot feature  $i_1^c, i_2^c, \dots, i_{|I_{cat}^c|}^c \in I_{cat}^c$   
851

852 
$$\sum_{j=1}^{|I_{cat}^c|} i_j^c = 1 \text{ for all } c \in \{1, \dots, C\} \quad (C5)$$
  
853  
854  
855  
856  
857  
858  
859

860 **Classifier constraints** To encode the classifier prediction  $f(x')$  of the factual  $x'$  we repurpose the  
861 core components of the `gurobi-machinelearning` library (originally designed for encoding  
862 regressors) to instead produce an MILP encoder of Neural Network and Random Forest classifiers.  
863In constraint Eq. (C6), the model prediction is constrained to the variable  $y'$ .

864  
865  
866

$$y' = f(x') \quad (C6)$$

867 Note that in the MILP-MinDistmethod, we constrain  $y'$  to be the target class  $y^+$ . This is done by  
868 checking if the output logit for the correct class is larger than all other classes.  
869

870  
871  
872

$$y'_{y^+} > y'_i \text{ for all } i \neq y^+ \quad (C6b)$$

873 However, in the CONFEX method, this explicit constraint Eq. (C6b) is not required as explained in  
874 Section 3, we only require Eq. (C6).  
875

## 876 B.1 MILP ENCODING OF LOCALISED CP

877 **MILP encoding of CONFEX-LCP** The following algorithm Algorithm 2 computes the LCP  
878 quantile value in MILP. To do this, all calibration scores and calibration points must be accessible  
879 to the optimiser. Variables are constrained as distances from the test point to each calibration  
880 point, and another set of variables compute the corresponding weight according to the L1 kernel.  
881 These weights are used alongside calibration scores to identify the desired weighted quantile. This  
882 encoding is linear in the size of the calibration set.  
883

884

---

**885 Algorithm 2:** Localised CP constraints in MILP

---

886 **Input** : Calibration dataset  $\{(x_i, y_i)\}_{i=1}^n$ , corresponding scores  $\{s_i\}_{i=1}^n$ , test input  $x^*$ , L1  
887 localisation kernel with bandwidth  $h$ , level  $\alpha \in (0, 1)$

888 **Output:** Local quantile  $q_{1-\alpha}^{LCP}$

- 889 1 Sort  $\{(x_i, y_i)\}_{i=1}^n$  in ascending order w.r.t. scores.
- 890 2 Add  $n$  real variables  $d_1, \dots, d_n$ .
- 891 3 For  $i = 1, \dots, n$ , add the L1 distance constraint  $d_i = \|x_i - x^*\|_1$ .
- 892 4 Add  $n$  binary variables  $w_1, \dots, w_n$  as the weights induced by the L1 kernel.
- 893 5 For  $i = 1, \dots, n$ , add the constraint  $w_i = \mathbf{1}(d_i \leq h)$ , implemented for arbitrarily large  $M > 0$   
as

$$d_i \leq h + M(1 - w_i) \wedge d_i \geq h - Mw_i$$

- 894 6 Add  $n$  binary variables  $in_1, \dots, in_n$ ; each  $in_i$  keeps track if the score  $s_i$  is below the quantile.
- 895 7 Add integer variables  $W$  and  $W_{1-\alpha}$  denoting, respectively, the sum of all weights and of those  
896 weights whose score is below the quantile.
- 897 8 Add constraints  $W = \sum_{i=1}^n w_i$ ,  $W_{1-\alpha} = \sum_{i=1}^n in_i \cdot w_i$  and  $W_{1-\alpha} \geq \lceil (1 - \alpha)W \rceil$ . The latter  
898 expresses that the scores below the quantile have probability at least  $1 - \alpha$ .
- 899 9 Define  $W'_{1-\alpha} = \sum_{i=1}^n (1 - in_i) \cdot w_i$  and add constraint  $W'_{1-\alpha} \geq \lfloor \alpha W \rfloor$
- 900 10 Solve constraints and return  $s_k$ .
- 901 11  $q_{1-\alpha}^{LCP}$  will be the largest calibration score  $s_i$  for which  $in_i = 1$ . To identify it, add an integer  
902 variable  $k \in \{1, \dots, n\}$ .
- 903 12 For  $i = 1, \dots, n$ , add the constraint  $in_i = \mathbf{1}(i \leq k)$  using a big-M encoding as done in line 5.

---

904

905

906

907

**908 MILP encoding of CONFEX-Tree** Following the Tree Construction procedure in Algorithm 1,  
909 we obtain a family of decision trees. Each tree contains with leaf nodes holding a centre midpoint  
910  $m$  and quantile  $q$ , these are concatenated as a single vector  $[c_1, c_2, \dots, c_D, q]$ .

911

912

913

914

The MILP encoding procedure consists of encoding the decision tree in the MILP problem using the  
915 `gurobi-machinelearning` library, selecting the correct tree  $T$  to use based on the categorical  
916 values of  $x'$ , identifying the leaf of  $T$  corresponding to  $x'$ , to obtain or reject the quantile based on  
917 the midpoint and distance.

918

919

920

Let there be  $N_T$  trees, each corresponding to a particular categorical combination. Note that certain  
921 categorical features can be ignored, or different categorical values can be considered the same, in  
922 order to reduce the number of trees required to be encoded without sacrificing any formal guarantees.  
923 This corresponds to a different notion of similarity in the LCP kernel.

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930

---

**Algorithm 3: CONFEX-Tree constraints in MILP**


---

**Input** : Trees  $T_1, \dots, T_{N_T}$  with corresponding categorical indicators  $\theta_1, \theta_2, \dots, \theta_{N_T}$ ,  
produced by Algorithm 1, bandwidth  $h$ , test input  $x'$

**Output:** CONFEX-Tree local quantile  $q_{1-\alpha}^{\text{Tree}}$

931  
932  
933  
934  
935  
936  
937  
938  
939

**1 Encoding of trees**

- 2 Add  $N_T$  vector variables  $t_1, t_2, \dots, t_{N_T}$  of shape  $(1 + \text{length of } x')$ , for the output of each tree.
- 3 Let  $x'_{\text{noncat}}$  be the vector  $x'$  excluding all categorically encoded entries.
- 4 Constrain  $t_i = T_1(x'_{\text{noncat}})$  for all  $i \in 1, \dots, N_T$

940

**5 Selection of tree**

- 6 Add  $N_T$  binary indicators  $\tau_1, \tau_2, \dots, \tau_{N_T}$  to determine the active tree.
- 7 Constrain  $\sum_{i=1}^{N_T} \tau_i = 1$ 
  - ▷ The following constraints determine if tree  $T_i$  is active by considering values of  $\theta_i$ , which are the indices of one-hot entries in  $x'$  which should be 1 if the tree  $T_i$  is selected.
- 8 For each  $i$ , constrain  $\tau_i \leq x'[j]$  for all  $j \in \theta_i$ 
  - ▷ Ensure  $\theta_i$  is 0 if any one-hot entries corresponding to the tree is 1.  $|\theta_i|$  is the number of categorical features.
- 9 For each  $i$ , constrain  $\tau_i \geq \sum_{j \in \theta_i} x'[j] - |\theta_i| + 1$  ▷ Ensure  $\theta_i$  is 1 if all one-hot entries corresponding to the tree is 0
- 10 Obtain the tree output  $t$  as  $t = \sum_{i=1}^{N_T} t_i \tau_i$

951

**11 Obtaining of quantile**

▷ Note that in the case of no categorical values (only numeric/ordinal), we have only one tree  $T$  and the algorithm can at this line after constraining  $t = T(x'_{\text{noncat}})$

952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

- 12 Index  $t$  as  $t = [c_1, c_2, \dots, c_D, q]$
- 13 Let  $c = [c_1, c_2, \dots, c_D]$
- 14 Constrain  $d$  to distance of  $x^*$  to centre:  $d = \|x^* - c\|_\infty$
- 15 Constrain  $d \leq h/2$ , since otherwise the point would be rejected and the quantile would be  $\infty$ .
- 16 Obtain  $q$  as the local quantile.

---

972 **MILP encoding of singleton conformal prediction set** The local quantile  $q$  is constrained using  
 973 Algorithm 2 or Algorithm 3. This is used to constrain the conformal prediction set  $C_{1-\alpha}(x')$  to a  
 974 singleton set.

975 Add  $|\mathcal{Y}|$  real variables  $s_1, s_2, \dots, s_{|\mathcal{Y}|}$  to represent the score  $s(x, y)$  for each  $y \in \mathcal{Y}$ .  
 976

977 For random forest classifiers, we use the score function  $s(x, y) = 1 - f(x)_y$ .  
 978

$$979 \quad s_i = 1 - y'_i \text{ for } i = 1, \dots, |\mathcal{Y}| \quad (C7)$$

981 For MLP classifiers that output unnormalised logits, we use the score function Eq. (7), which is  
 982  $s(x, y) = -l(x)_y + \max_{y' \neq y} l(x)_{y'}$ , where  $l(x)$  is the predicted vector of logits.  
 983

$$985 \quad \text{Add } m \text{ real variables } m_1, \dots, m_{|\mathcal{Y}|} \quad (C8)$$

$$986 \quad m_i = \max_{j \neq i} y_j, \quad i = 1, \dots, |\mathcal{Y}| \quad (C9)$$

$$988 \quad s_i = -y'_i + m_i, \quad i = 1, \dots, |\mathcal{Y}| \quad (C10)$$

990 In the case of binary classification, which any multiclass CFX problem can be reduced to, our score  
 991 function is the difference between the two logits and we can remove the maximum constraint as  
 992 follows.

$$995 \quad s_1 = y'_2 - y'_1 \text{ for } i = 1 \quad (C11)$$

$$996 \quad s_2 = -s_1 \text{ for } i = 2 \quad (C12)$$

998 Finally, we constrain  $C_{1-\alpha}(x')$  to a singleton set containing the target class,  $\{y^+\}$ .  
 999

$$1001 \quad s_i \leq q \text{ for all } i \neq y^+ \quad (C13)$$

$$1002 \quad s_i > q \text{ for } i = y^+ \quad (C14)$$

1004 **Further constraints** As mentioned in Section 2, further properties such as causality and action-  
 1005 ability can be incorporated into the model by introducing further constraints on  $x'$ .  
 1006

1007 **Notes on MILP** Strict inequalities such as those present in Eq. (C6b), Eq. (C14) and some model  
 1008 encodings Eq. (C6) can not directly be modelled in MILP, this is resolved by adding a small epsilon  
 1009 to one side of the equation.

1010 In Eq. ((C4)), Algorithm 2, Algorithm 3, we observe products of two variables which would usually  
 1011 indicate a quadratic constraint. However, in all of these cases, at least one variable is a binary  
 1012 variable. This allows the solver to linearise it, see Klotz (2021) for further details.  
 1013

1014  
 1015  
 1016  
 1017  
 1018  
 1019  
 1020  
 1021  
 1022  
 1023  
 1024  
 1025

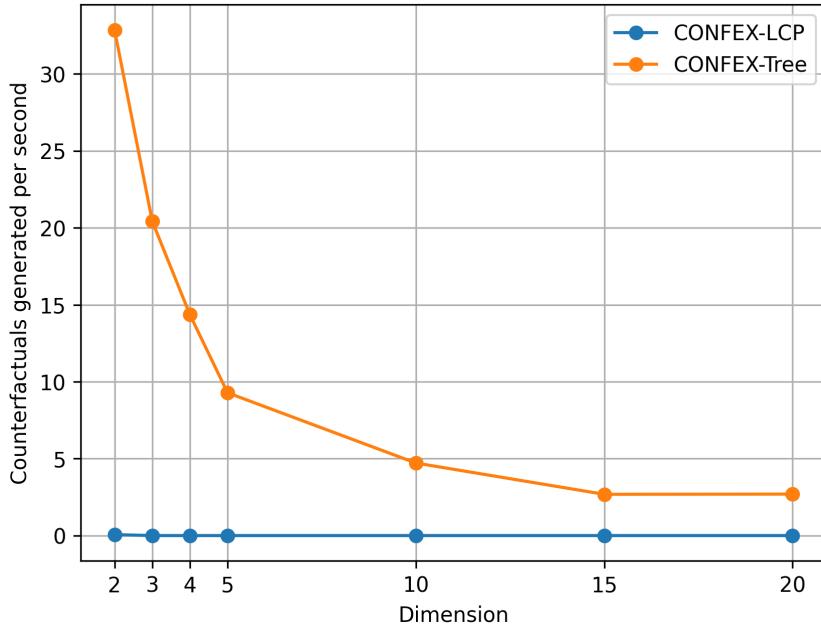
## 1026 C SCALABILITY ANALYSIS

1028 In this section, we empirically analyse the scalability of the CONFEX-LCP and CONFEX-Tree  
 1029 methods. All experiments were conducted on a MacBook Pro, M3 Pro chipset, 18 GB memory.

### 1031 C.1 DATASET DIMENSION

1033 The dimensionality of the dataset affects the number of variables involved in the distances and  
 1034 weights constraints for CONFEX-LCP and the complexity of the tree for CONFEX-Tree. For this  
 1035 experiment, we analyse the effect of changing the dimensionality of the dataset.

1036 We create synthetic datasets of varying dimensionalities by using sklearn's  
 1037 make\_classification method with all features being informative, and we plot the number of  
 1038 counterfactuals generated per second (observed over a 5-minute period) in Fig. 3. Tabular results  
 1039 are available in Table 4. Note that we fix the kernel bandwidth and alpha value, use a MLP model  
 1040 with 50 hidden units, and use a calibration set size of 150.



1062 Figure 3: Counterfactuals generated per second for CONFEX-LCP and CONFEX-Tree, against  
 1063 dimensionality of the dataset.

1065 We find that beyond a dimensionality of 2, CONFEX-LCP is infeasible for use, whilst CONFEX-  
 1066 Tree's generation rate eventually flattens.

1069	Dataset dimension	CONFEX-LCP	CONFEX-Tree
1070	2	0.060	32.853
1071	3	0.000	20.427
1072	4	0.000	14.353
1073	5	0.000	9.280
1074	10	0.000	4.713
1075	15	0.000	2.680
1076	20	0.000	2.693

1077 Table 4: Counterfactuals generated per second for CONFEX-LCP and CONFEX-Tree, against di-  
 1078 mensionality of the dataset.

1080

## C.2 SIZE OF CALIBRATION SET

1081

1082

1083 Varying the size of the calibration set also affects the number of variables involved in constraining  
 1084 distances and weights in CONFEX-LCP as well as affecting the complexity of the tree in CONFEX-  
 1085 Tree procedure.

1086

1087

1088

1089

For this experiment, we use the CaliforniaHousing dataset (which has 8 dimensions) with an MLP  
 model containing 50 hidden units, and fix the kernel bandwidth and alpha value. We vary the size  
 of the calibration set from between 10 and 2000 points. The effect on the rate of counterfactual  
 generation is found in Fig. 4.

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

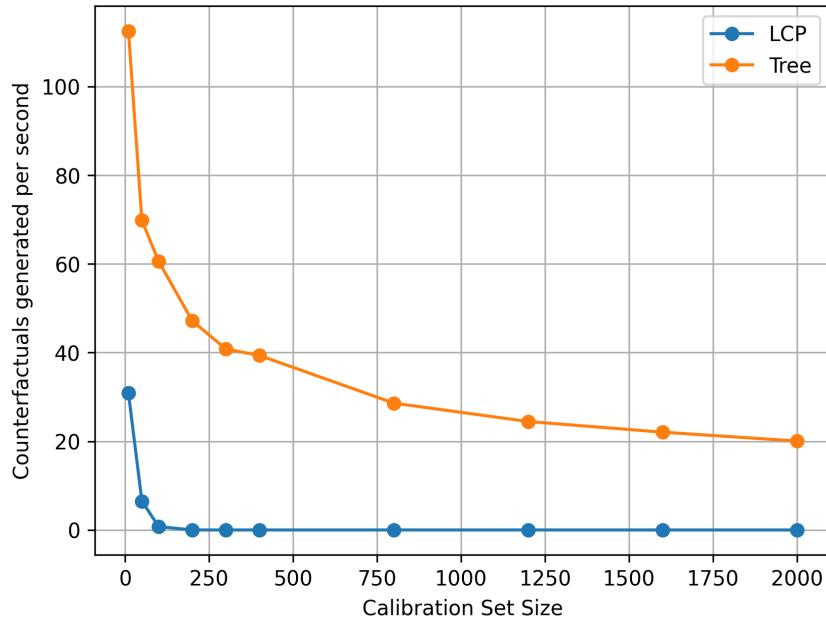
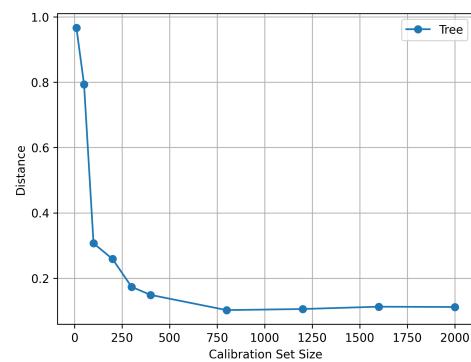
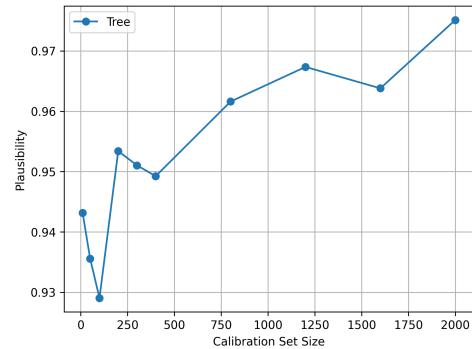


Figure 4: Number of counterfactuals generated per second for CONFEX-LCP and CONFEX-Tree on the CaliforniaHousing dataset, varying the calibration set size.

Fig. 8 shows the effect that an increased calibration set size has on distance and plausibility: we get reduced distances with improved plausibility.



(a) Distance



(b) Plausibility

Figure 5: Effect of varying the calibration set size on distance and plausibility on the CaliforniaHousing dataset, CONFEX-Tree.

1134 C.3 MODEL COMPLEXITY  
1135  
1136

1137 Model outputs are obtained within the MILP formulation through a series of constraints involving  
1138 the input variables and the details of the trained model. As explained in Section B, (C6), we modify  
1139 implementations from the *gurobi-machinelearning* library.

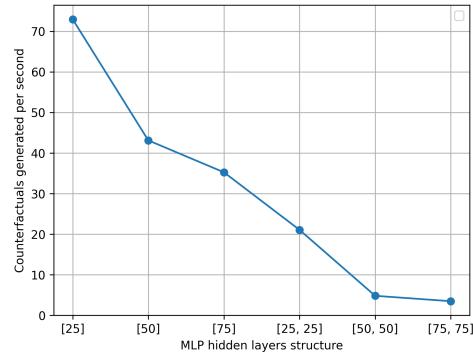
1140 In this section, we see how CONFEX-Tree performs as the model complexity changes. We consider  
1141 three classes of models: multi-layer perceptron, random forests and gradient-boosted trees, varying  
1142 their hyperparameters. Although we use CONFEX-Tree to generate counterfactuals, this analysis  
1143 focusses on the MILP encoding of the classifiers and would apply to other MILP methods as well,  
1144 e.g. MinDist. We use the CaliforniaHousing dataset and fix alpha to 0.1, and bandwidth scale to  
1145 1. The following figures show how the number of CFXs generated per second varies as we change  
1146 hyperparameters. The accuracy of the models over a test set is also shown.

1147

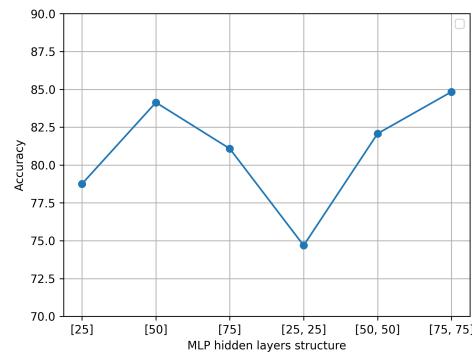
1148

1149

1150



(a) CFXs generated per second



(b) Test Accuracy

1162

1163

1164 Figure 6: Effect of varying MLP model hyperparameters on number of CFXs generated per second,  
1165 and accuracy for CONFEX-Tree.

1166

1167

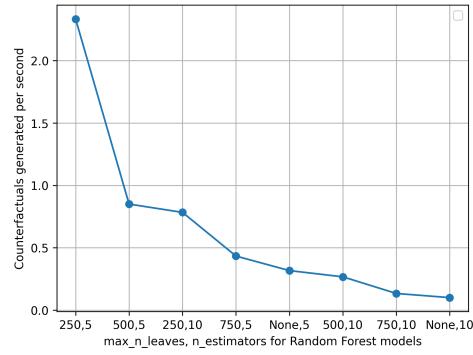
1168

1169

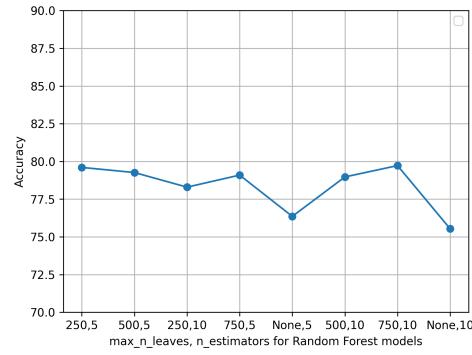
1170

1171

1172



(a) CFXs generated per second

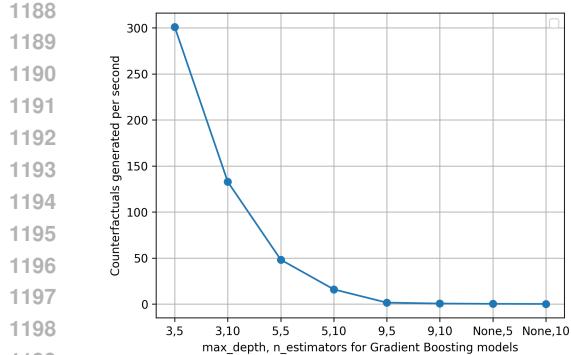


(b) Test Accuracy

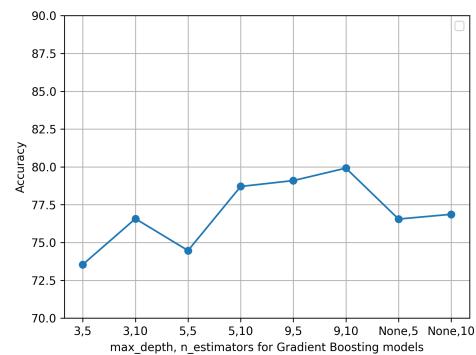
1186

1187

1188 Figure 7: Effect of varying Random Forest model hyperparameters on number of CFXs generated  
1189 per second, and accuracy for CONFEX-Tree.



(a) CFXs generated per second



(b) Test Accuracy

Figure 8: Effect of varying Gradient Boosted Trees hyperparameters on number of CFXs generated per second, and accuracy for CONFEX-Tree.

We find that as the complexity of the model increases, the number of CFXs generated per second decreases. In the worst case tested, Gradient Boosted Trees with no depth limit and 10 estimators, we obtain a reasonable 4 CFXs per second. From the accuracy plots, we can see that in many cases, a less complex model can provide similar accuracy and more complex models exhibit overfitting.

1242 D FURTHER EVALUATION  
12431244 D.1 EXPERIMENTAL SETUP  
12451246 **Generators.** For solving MILP instances, we utilise the Gurobi solver, and utilise the Gurobi  
1247 Machine Learning Gurobi (2022) library to formulate the trained classifiers as constraints. All  
1248 generators, except FOCUS (using the CFXplorer package Morita (2023)) and FeatureTweak/FACE/C-  
1249 CHVAE (implementations ported from CARLA Pawelczyk et al. (2021), and FeatureTweakPy<sup>4</sup>),  
1250 were implemented as part of a Python library to generate CFXs. The details of this library are  
1251 removed for anonymous submission.  
12521253 **Model Configuration.** For all datasets, we used a multilayer perceptron (MLP) with 50 hidden  
1254 units. The batch size was set to 64 for California Housing and German Credit, trained for 100  
1255 epochs, and 256 for GiveMeSomeCredit and Adult Income, trained for 50 epochs. For the random  
1256 forest model, we also evaluated a Random Forest classifier with 5 estimators and number of leaves  
1257 limited to 500 for the GiveMeSomeCredit and AdultIncome models.  
12581259 D.2 METRICS  
12601261 In order to evaluate the quality of the generated counterfactual explanations, we adopt a set of quanti-  
1262 tative metrics that measure different aspects of their usefulness and reliability. Specifically, we focus  
1263 on three core dimensions: *plausibility*, *sensitivity*, and *stability*. In addition, we report auxiliary met-  
1264 rics such as the distance of counterfactuals to the original instance, the proportion of failures, and  
1265 the validity rate of generated explanations. Together, these metrics provide a comprehensive view  
1266 of both the fidelity and robustness of counterfactual explanations.  
12671268 **Plausibility.** A counterfactual explanation should lie close to the underlying data distribution so  
1269 that it represents a realistic and interpretable alternative. To assess this, we measure plausibility  
1270 using the Local Outlier Factor (LOF) (Breunig et al., 2000), which quantifies how isolated a sample  
1271 is with respect to its nearest neighbours. A LOF score of +1 indicates that the counterfactual is  
1272 consistent with observed data, whereas -1 suggest that the counterfactual is implausible. We use the  
1273 scikit-learn implementation of LOF with `novelty=True` and `n_neighbors = 20`, stratified  
1274 by the target class. In practice, we average over 100 test points.  
12751276 **Sensitivity.** Beyond plausibility, we also want to assess whether counterfactuals are *robust* to small  
1277 changes in the input instance. Sensitivity measures how much a counterfactual explanation changes  
1278 when the original instance  $x$  is perturbed within a small neighbourhood. Formally, given an input  $x$   
1279 and its counterfactual  $x_c$ , we uniformly sample a perturbed instance  $x' \sim U_b(x)$  from the  $\ell_2$  ball centred  
1280 around the factual, compute a new counterfactual  $x'_c$ . Sensitivity is then defined as the relative  
1281 deviation between the two counterfactuals, normalised by the cost of the initial counterfactual:  
1282

1283 
$$\text{CFX Sensitivity} = \mathbb{E}_{x' \sim U_b(x)} \left[ \frac{\|x'_c - x_c\|_2}{\|x_c - x\|_2} \right].$$
  
1284

1285 In practice, we sample 4 neighbours from 25 test points to inform our sensitivity metric. Intuitively,  
1286 low sensitivity indicates that the explanation remains stable when the factual input undergoes small  
1287 variations, thereby suggesting robustness and consistency.  
12881289 In our experiments, we choose the budget  $b$  of the uniform sampling to correspond to a ball with  
1290 0.1% of the volume of the feature space.  
1291

1292 
$$V_{\text{ball}} = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2} + 1)} r^d = bV_{\text{total}}$$
  
1293

1294 where  $d$  is the number of non-categorical features in the space. Solving for  $r$ ,  
1295

1296 
$$r = \left( \frac{bV_{\text{total}}}{\pi^{d/2}/\Gamma(\frac{d}{2} + 1)} \right)^{1/d}$$
  
1297

4`https://github.com/upura/featureTweakPy/blob/master/featureTweakPy.py`

1296 This allows the same budget to be used across datasets with differing numbers of features. When  
 1297 sampling neighbours, we do not change categorical values and we fix ordinal values to their closest  
 1298 valid value.  
 1299

1300 **Stability.** Complementary to sensitivity, stability measures how consistent the counterfactual is  
 1301 under perturbations applied directly to the counterfactual itself. That is, we perturb  $x_c$  within a  
 1302 budgeted neighbourhood and evaluate the variance in the model predictions across these perturbed  
 1303 samples. Following an adaptation of (Dutta et al., 2022), stability is computed as:

$$1304 \text{CFX Stability} = \frac{1}{K} \sum_{x' \in N_x} \hat{f}(x')_{y^+} - \sqrt{\frac{1}{K} \sum_{x' \in N_x} \left( \hat{f}(x')_{y^+} - \frac{1}{K} \sum_{x' \in N_x} \hat{f}(x')_{y^+} \right)^2},$$

1305 where  $N_x$  is a set of  $K$  points sampled as  $x' \sim U_b(x_c)$ .  
 1306

1307 where  $\hat{f}(x')_{y^+}$  refers to the predicted probability of the target class. The metric neighbours a large  
 1308 mean value for the predicted probability of sampled neighbours, whilst penalising variations in these  
 1309 values by subtracting the standard deviation to ensure that that mean is not a combination of very  
 1310 high and very low values. Similarly to the Sensivity metric,  $U_b(x_c)$  denotes sampling from the  $\ell_2$   
 1311 ball centred around the counterfactual, computing the radius in the same way, taking the budget to  
 1312 represent 0.1% of the total feature volume.  
 1313

1314 Stability is high when the predictions across perturbed counterfactuals remain close to each other,  
 1315 which indicates that the explanation is not overly sensitive to minor fluctuations in its actualisation.  
 1316

1317 **Certainty.** CONFEX minimises the uncertainty of the counterfactual by constraining the conformal  
 1318 prediction set to be a singleton containing the target class only. Certainty in the counterfactual relates  
 1319 to the property  $\mathbb{P}(y = y^+ | x = x')$ .  
 1320

1321 To quantify the certainty of the counterfactuals in a principled way, we use local conformal  $p$ -values.  
 1322 In a conformal prediction procedure, the conformal  $p$ -value of a point  $(x, y)$  is the proportion of the  
 1323 calibration points with score above  $s(x, y)$ . It is used to determine which labels are included in the  
 1324 prediction set: labels with a  $p$ -value over  $\alpha$  are included and the rest excluded. This is equivalent to  
 1325 checking if  $s(x, y)$  is above the  $1 - \alpha$  quantile of the calibration score (as explained in Section 2).  
 1326

1327 A high  $p$ -value for  $(x, y)$  provides strong evidence that  $y$  is the true label for  $x$ . Hence, for our  
 1328 certainty metric, we compute the average difference between the conformal  $p$ -value for the target  
 1329 class and the max of conformal  $p$ -values for all other classes. If the  $p$ -value of the target class is high  
 1330 and the max  $p$ -value of the other classes are low - indicating a certain prediction with strong evidence  
 1331 in favour of the target class and against others - then our metric will be high. If the prediction is  
 1332 uncertain then the  $p$ -values of all classes will be similar, leading to a lower value of our metric.  
 1333

1334 Note that we use the LCP procedure to compute  $p$ -values for this metric because of its local guarantees.  
 1335 We don't use vanilla CP, because its resulting  $p$ -values would be affected by calibrations points  
 1336 well-away (not local) to the counterfactual point of interest. We do not use the CONFEX-Tree  
 1337 procedure to compute local  $p$ -values since the metric may be seen as tailored to our generator.  
 1338

$$1339 \text{cert}(x') = p_{y^+}(x') - \max_{y \neq y^+} p_y(x') \quad (18)$$

1340 Certainty results are reported in tables found in Section A.2.  
 1341

1342 **Auxiliary metrics.** In addition to the three core dimensions, we report the following supplementary  
 1343 measures:  
 1344

- 1345 • *Distance*: the average L1 distance between the original instance and the counterfactual,  
 1346

$$1347 \text{Distance} = \mathbb{E}(\|x' - x_0\|_1),$$

1348 which quantifies the minimality of the intervention required.  
 1349

- *Validity*: the proportion of counterfactuals that successfully change the prediction to the desired  
 1350 class,  
 1351

$$1352 \text{Validity} = \mathbb{E}(1\{\hat{f}(x') = y^+\}).$$

1350  
 1351 For example, invalidity could be due to numerical artefacts in encoding the models in MILP, or  
 1352 failure for SGD procedures to converge to a flipped class. We report whenever a method a method  
 1353 produces less than 90% validity, and exclude invalid CFXs from the computation of other metrics.  
 1354

- 1355 • *Failure rate*: the proportion of runs where the generator fails to produce a counterfactual, for  
 1356 example due to infeasible constraints in optimisation-based methods such as MILP.
- 1357 • *Implausibility*: The average distance from the counterfactual to the closest 10% of points of the  
 1358 target class, similar to Altmeyer et al. (2024).

1359 **D.2.1 CONDITIONAL COVERAGE RESULTS**

1360 In the additional results, we furthermore evaluate the performance of different conformal CFX gener-  
 1361 ators under four evaluation settings: marginal coverage, class-conditional coverage, random binning,  
 1362 and counterfactual similarity. In the paper we discussed the counterfactual simulation, however we  
 1363 also evaluate the marginal coverage over a test set, average class-conditional coverage, average cov-  
 1364 erage over a random partitioning of the test set into 3 bins. We report the coverage gap (Barber et al.  
 1365 (2023)): the difference between the empirical coverage and target coverage, in percentage points.

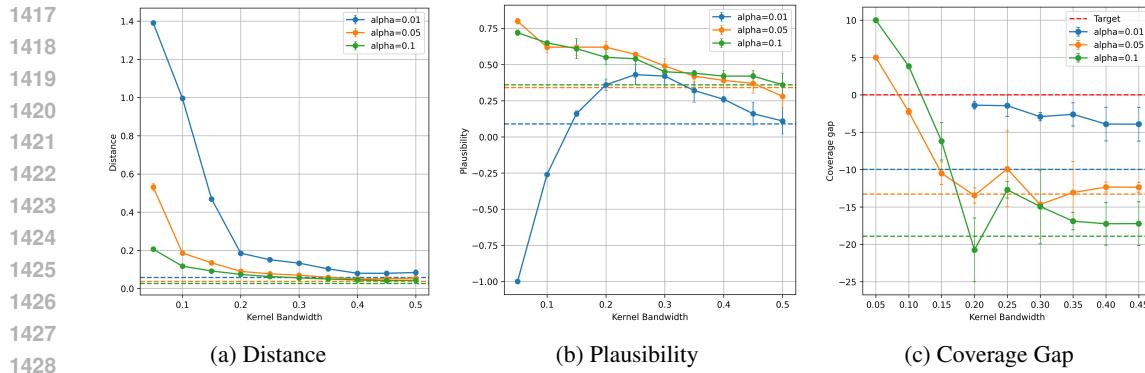
$$1366 \quad \text{CovGap} = 100 \times (\mathbb{P}\{y \in C(x)\} - (1 - \alpha)) \quad (19)$$

1367  
 1368  
 1369  
 1370  
 1371  
 1372  
 1373  
 1374  
 1375  
 1376  
 1377  
 1378  
 1379  
 1380  
 1381  
 1382  
 1383  
 1384  
 1385  
 1386  
 1387  
 1388  
 1389  
 1390  
 1391  
 1392  
 1393  
 1394  
 1395  
 1396  
 1397  
 1398  
 1399  
 1400  
 1401  
 1402  
 1403

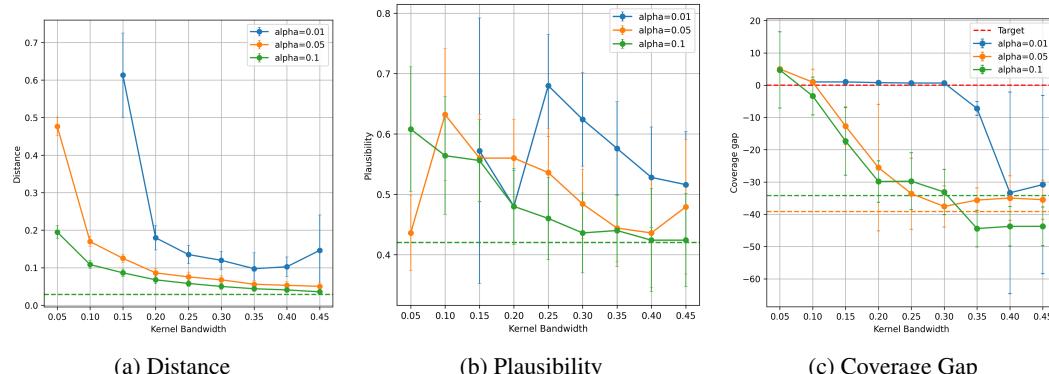
1404 D.3 CALIFORNIA HOUSING  
1405

1406 We use the California Housing dataset Pace & Barry (1997) from the StatLib repository through  
1407 scikit-learn’s `sklearn.datasets.fetch_california_housing` function<sup>5</sup>. The original  
1408 regression problem was changed into a binary classification task by categorizing houses based on  
1409 whether the median income exceeds \$20,000 (42% above, 58% below). The dataset contains 8  
1410 numeric features, which we scaled to the range (0, 1) using MinMax scaling.

1411 Our results demonstrate a nice pattern showing that distance decreases and plausibility decreases  
1412 as the kernel bandwidth increases. CONFEX methods outperform all other methods (except FACE,  
1413 where it comes second) on plausibility and sensitivity.

1414 D.3.1 PLOTS  
1415

1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439  
1440  
1441  
1442  
1443  
1444  
1445



1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457

<sup>5</sup>[https://www.dcc.fc.up.pt/~ltorgo/Regression/cal\\_housing.html](https://www.dcc.fc.up.pt/~ltorgo/Regression/cal_housing.html)

1458 D.3.2 MODEL EVALUATION RESULTS  
1459

1460	1461	1462	1463	1464	1465	1466	1467	1468	1469	1470	1471	1472	1473	Repeat	Accuracy (%)	Precision (%)	F1 Score (%)
1462	1463	repeat0,MLP	83.58	83.61	83.59												
1463	1464	repeat1,MLP	82.95	83.59	82.95												
1464	1465	repeat2,MLP	78.20	80.17	78.06												
1465	1466	repeat3,MLP	81.59	82.37	81.59												
1466	1467	repeat4,MLP	79.31	80.71	79.25												
1467	1468	repeat0,RF	78.05	80.60	77.85												
1468	1469	repeat1,RF	78.10	80.60	77.90												
1469	1470	repeat2,RF	77.59	81.06	77.26												
1470	1471	repeat3,RF	76.02	79.26	75.68												
1471	1472	repeat4,RF	76.36	79.09	76.10												

1474 Table 5: Model evaluation results, CaliforniaHousing.  
1475  
1476  
1477  
1478  
1479  
1480  
1481  
1482  
1483  
1484  
1485  
1486  
1487  
1488  
1489  
1490  
1491  
1492  
1493  
1494  
1495  
1496  
1497  
1498  
1499  
1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511

1512 D.3.3 CFX GENERATION RESULTS  
1513  
1514  
1515  
1516  
1517  
1518  
1519

1520 Generator	1521 Distance	1522 Plausibility	1523 Implausibility	1524 Sensitivity ( $10^{-1}$ )	1525 Stability
<b>MLP</b>					
1522 MinDist	0.03 $\pm$ 0.00	0.30 $\pm$ 0.07	0.21 $\pm$ 0.01	42.75 $\pm$ 8.45	0.02 $\pm$ 0.04
1523 Wachter	0.09 $\pm$ 0.01	0.42 $\pm$ 0.08	0.20 $\pm$ 0.00	1.66 $\pm$ 0.37	0.02 $\pm$ 0.05
1524 Greedy	1.88 $\pm$ 0.27	-0.99 $\pm$ 0.02	0.89 $\pm$ 0.10	0.14 $\pm$ 0.02	0.42 $\pm$ 0.10
1525 ConfexNaive					
1526 $\alpha = 0.01$	0.07 $\pm$ 0.01	0.04 $\pm$ 0.06	0.22 $\pm$ 0.01	3.45 $\pm$ 1.25	0.02 $\pm$ 0.05
1527 $\alpha = 0.05$	0.04 $\pm$ 0.01	0.24 $\pm$ 0.08	0.22 $\pm$ 0.01	8.41 $\pm$ 2.26	0.02 $\pm$ 0.05
1528 $\alpha = 0.1$	0.03 $\pm$ 0.01	0.27 $\pm$ 0.08	0.21 $\pm$ 0.01	14.51 $\pm$ 3.09	0.02 $\pm$ 0.04
1529 ECCCCo					
1530 $\alpha = 0.01$	0.39 $\pm$ 0.02	-0.69 $\pm$ 0.04	0.21 $\pm$ 0.01	0.24 $\pm$ 0.05	0.26 $\pm$ 0.08
1531 $\alpha = 0.05$	0.37 $\pm$ 0.02	-0.65 $\pm$ 0.05	0.21 $\pm$ 0.01	0.26 $\pm$ 0.05	0.25 $\pm$ 0.08
1532 $\alpha = 0.1$	0.37 $\pm$ 0.02	-0.63 $\pm$ 0.04	0.21 $\pm$ 0.01	0.26 $\pm$ 0.05	0.24 $\pm$ 0.08
1533 ConfexTree, $\alpha = 0.01$					
1534 $bw = 0.05$	1.40 $\pm$ 0.04	-1.00 $\pm$ 0.00	0.39 $\pm$ 0.00	0.00 $\pm$ 0.00	1.00 $\pm$ 0.00
1535 $bw = 0.1$	1.03 $\pm$ 0.04	-0.48 $\pm$ 0.19	0.15 $\pm$ 0.01	0.02 $\pm$ 0.01	0.03 $\pm$ 0.04
1536 $bw = 0.15$	0.52 $\pm$ 0.04	0.08 $\pm$ 0.19	0.16 $\pm$ 0.00	0.07 $\pm$ 0.01	0.05 $\pm$ 0.04
1537 $bw = 0.2$	0.21 $\pm$ 0.03	0.30 $\pm$ 0.12	0.16 $\pm$ 0.00	0.40 $\pm$ 0.27	0.05 $\pm$ 0.04
1538 $bw = 0.25$	0.17 $\pm$ 0.02	0.35 $\pm$ 0.17	0.16 $\pm$ 0.00	0.60 $\pm$ 0.32	0.05 $\pm$ 0.05
1539 $bw = 0.3$	0.16 $\pm$ 0.02	0.33 $\pm$ 0.13	0.17 $\pm$ 0.00	1.02 $\pm$ 0.37	0.05 $\pm$ 0.05
1540 $bw = 0.35$	0.12 $\pm$ 0.01	0.28 $\pm$ 0.09	0.18 $\pm$ 0.00	2.07 $\pm$ 0.62	0.04 $\pm$ 0.05
1541 $bw = 0.4$	0.09 $\pm$ 0.01	0.17 $\pm$ 0.08	0.20 $\pm$ 0.00	2.91 $\pm$ 0.99	0.03 $\pm$ 0.05
1542 $bw = 0.45$	0.09 $\pm$ 0.01	0.11 $\pm$ 0.08	0.20 $\pm$ 0.00	2.88 $\pm$ 0.76	0.03 $\pm$ 0.05
1543 $bw = 0.5$	0.08 $\pm$ 0.01	0.11 $\pm$ 0.09	0.20 $\pm$ 0.00	2.14 $\pm$ 0.31	0.08 $\pm$ 0.03
1544 ConfexTree, $\alpha = 0.05$					
1545 $bw = 0.05$	0.55 $\pm$ 0.04	0.75 $\pm$ 0.05	0.15 $\pm$ 0.00	0.05 $\pm$ 0.05	0.07 $\pm$ 0.03
1546 $bw = 0.1$	0.19 $\pm$ 0.02	0.60 $\pm$ 0.07	0.17 $\pm$ 0.01	0.43 $\pm$ 0.12	0.05 $\pm$ 0.05
1547 $bw = 0.15$	0.15 $\pm$ 0.01	0.53 $\pm$ 0.09	0.17 $\pm$ 0.01	0.83 $\pm$ 0.48	0.05 $\pm$ 0.05
1548 $bw = 0.2$	0.10 $\pm$ 0.02	0.55 $\pm$ 0.09	0.18 $\pm$ 0.01	2.15 $\pm$ 0.73	0.03 $\pm$ 0.05
1549 $bw = 0.25$	0.09 $\pm$ 0.02	0.50 $\pm$ 0.09	0.18 $\pm$ 0.01	3.05 $\pm$ 0.79	0.03 $\pm$ 0.04
1550 $bw = 0.3$	0.08 $\pm$ 0.01	0.43 $\pm$ 0.07	0.19 $\pm$ 0.01	4.73 $\pm$ 2.37	0.03 $\pm$ 0.04
1551 $bw = 0.35$	0.07 $\pm$ 0.01	0.36 $\pm$ 0.07	0.19 $\pm$ 0.00	6.03 $\pm$ 0.74	0.03 $\pm$ 0.04
1552 $bw = 0.4$	0.06 $\pm$ 0.01	0.30 $\pm$ 0.08	0.20 $\pm$ 0.00	6.69 $\pm$ 1.37	0.03 $\pm$ 0.04
1553 $bw = 0.45$	0.06 $\pm$ 0.01	0.26 $\pm$ 0.10	0.20 $\pm$ 0.01	6.93 $\pm$ 1.44	0.03 $\pm$ 0.04
1554 $bw = 0.5$	0.05 $\pm$ 0.01	0.28 $\pm$ 0.08	0.20 $\pm$ 0.00	7.16 $\pm$ 0.70	0.07 $\pm$ 0.02
1555 ConfexTree, $\alpha = 0.1$					
1556 $bw = 0.05$	0.25 $\pm$ 0.04	0.58 $\pm$ 0.12	0.17 $\pm$ 0.01	0.20 $\pm$ 0.09	0.06 $\pm$ 0.04
1557 $bw = 0.1$	0.14 $\pm$ 0.02	0.62 $\pm$ 0.05	0.17 $\pm$ 0.01	0.96 $\pm$ 0.75	0.04 $\pm$ 0.04
1558 $bw = 0.15$	0.11 $\pm$ 0.02	0.58 $\pm$ 0.08	0.18 $\pm$ 0.00	2.76 $\pm$ 2.77	0.03 $\pm$ 0.05
1559 $bw = 0.2$	0.09 $\pm$ 0.01	0.48 $\pm$ 0.07	0.19 $\pm$ 0.00	4.88 $\pm$ 2.61	0.03 $\pm$ 0.04
1560 $bw = 0.25$	0.07 $\pm$ 0.01	0.44 $\pm$ 0.09	0.19 $\pm$ 0.00	4.34 $\pm$ 1.72	0.03 $\pm$ 0.04
1561 $bw = 0.3$	0.07 $\pm$ 0.01	0.39 $\pm$ 0.05	0.19 $\pm$ 0.00	6.42 $\pm$ 2.59	0.02 $\pm$ 0.04
1562 $bw = 0.35$	0.06 $\pm$ 0.01	0.36 $\pm$ 0.07	0.19 $\pm$ 0.00	9.80 $\pm$ 0.97	0.02 $\pm$ 0.04
1563 $bw = 0.4$	0.05 $\pm$ 0.01	0.34 $\pm$ 0.07	0.20 $\pm$ 0.00	12.93 $\pm$ 3.75	0.02 $\pm$ 0.04
1564 $bw = 0.45$	0.05 $\pm$ 0.01	0.32 $\pm$ 0.09	0.20 $\pm$ 0.01	14.53 $\pm$ 3.82	0.02 $\pm$ 0.04
1565 $bw = 0.5$	0.04 $\pm$ 0.01	0.36 $\pm$ 0.08	0.20 $\pm$ 0.00	13.87 $\pm$ 0.66	0.07 $\pm$ 0.02
FACE	0.21 $\pm$ 0.02	0.85 $\pm$ 0.04	0.16 $\pm$ 0.00	0.34 $\pm$ 0.03	0.06 $\pm$ 0.05
C-CHVAE	1.27 $\pm$ 0.22	-0.35 $\pm$ 0.12	0.46 $\pm$ 0.15	0.06 $\pm$ 0.03	0.16 $\pm$ 0.03

Table 6: CFX generation results, CaliforniaHousing, MLP. All methods attained full validity.

	Generator	Distance	Plausibility	Implausibility	Sensitivity ( $10^{-1}$ )	Stability
1566	MinDist	$0.01 \pm 0.00$	$0.37 \pm 0.07$	$0.21 \pm 0.00$	$89.15 \pm 90.88$	$0.22 \pm 0.02$
1567	ConfexNaive					
1568	$\alpha = 0.01$	$0.03 \pm 0.01$	$0.42 \pm 0.07$	$0.20 \pm 0.00$	$12.34 \pm 2.94$	$0.23 \pm 0.02$
1569	$\alpha = 0.05$	$0.03 \pm 0.01$	$0.42 \pm 0.07$	$0.20 \pm 0.00$	$12.34 \pm 2.94$	$0.23 \pm 0.02$
1570	$\alpha = 0.1$	$0.03 \pm 0.01$	$0.42 \pm 0.07$	$0.20 \pm 0.00$	$12.34 \pm 2.94$	$0.23 \pm 0.02$
1571	ConfexTree, $\alpha = 0.01$					
1572	$bw = 0.05$	$nan \pm nan$	$nan \pm nan$	$nan \pm nan$	$nan \pm nan$	$nan \pm nan$
1573	$bw = 0.1$	$nan \pm nan$	$nan \pm nan$	$nan \pm nan$	$nan \pm nan$	$nan \pm nan$
1574	$bw = 0.15$	$0.61 \pm 0.11$	$0.57 \pm 0.22$	$0.16 \pm 0.01$	$0.08 \pm 0.03$	$0.27 \pm 0.03$
1575	$bw = 0.2$	$0.18 \pm 0.03$	$0.48 \pm 0.06$	$0.16 \pm 0.00$	$0.59 \pm 0.23$	$0.23 \pm 0.02$
1576	$bw = 0.25$	$0.14 \pm 0.02$	$0.68 \pm 0.08$	$0.16 \pm 0.00$	$1.55 \pm 0.76$	$0.23 \pm 0.02$
1577	$bw = 0.3$	$0.12 \pm 0.02$	$0.62 \pm 0.08$	$0.17 \pm 0.00$	$2.13 \pm 0.77$	$0.23 \pm 0.02$
1578	$bw = 0.35$	$0.10 \pm 0.04$	$0.58 \pm 0.08$	$0.17 \pm 0.01$	$3.77 \pm 1.59$	$0.23 \pm 0.02$
1579	$bw = 0.4$	$0.10 \pm 0.03$	$0.53 \pm 0.08$	$0.18 \pm 0.00$	$5.53 \pm 1.39$	$0.22 \pm 0.02$
1580	$bw = 0.45$	$0.15 \pm 0.09$	$0.52 \pm 0.09$	$0.18 \pm 0.01$	$4.78 \pm 1.98$	$0.23 \pm 0.02$
1581	ConfexTree, $\alpha = 0.05$					
1582	$bw = 0.05$	$0.48 \pm 0.02$	$0.44 \pm 0.06$	$0.16 \pm 0.00$	$0.06 \pm 0.05$	$0.26 \pm 0.03$
1583	$bw = 0.1$	$0.17 \pm 0.01$	$0.63 \pm 0.11$	$0.17 \pm 0.00$	$0.47 \pm 0.10$	$0.23 \pm 0.01$
1584	$bw = 0.15$	$0.12 \pm 0.01$	$0.56 \pm 0.07$	$0.17 \pm 0.00$	$0.91 \pm 0.41$	$0.22 \pm 0.01$
1585	$bw = 0.2$	$0.09 \pm 0.01$	$0.56 \pm 0.06$	$0.18 \pm 0.00$	$2.58 \pm 1.75$	$0.22 \pm 0.02$
1586	$bw = 0.25$	$0.08 \pm 0.01$	$0.54 \pm 0.07$	$0.18 \pm 0.00$	$3.57 \pm 1.14$	$0.23 \pm 0.02$
1587	$bw = 0.3$	$0.07 \pm 0.01$	$0.48 \pm 0.06$	$0.19 \pm 0.00$	$4.52 \pm 1.61$	$0.22 \pm 0.02$
1588	$bw = 0.35$	$0.06 \pm 0.01$	$0.44 \pm 0.06$	$0.19 \pm 0.00$	$8.32 \pm 2.21$	$0.22 \pm 0.02$
1589	$bw = 0.4$	$0.05 \pm 0.01$	$0.44 \pm 0.09$	$0.19 \pm 0.00$	$7.12 \pm 2.07$	$0.22 \pm 0.02$
1590	$bw = 0.45$	$0.05 \pm 0.01$	$0.48 \pm 0.11$	$0.19 \pm 0.01$	$8.99 \pm 3.49$	$0.22 \pm 0.02$
1591	ConfexTree, $\alpha = 0.1$					
1592	$bw = 0.05$	$0.19 \pm 0.02$	$0.61 \pm 0.10$	$0.17 \pm 0.00$	$0.40 \pm 0.18$	$0.23 \pm 0.02$
1593	$bw = 0.1$	$0.11 \pm 0.01$	$0.56 \pm 0.10$	$0.18 \pm 0.00$	$1.64 \pm 1.18$	$0.22 \pm 0.02$
1594	$bw = 0.15$	$0.09 \pm 0.01$	$0.56 \pm 0.07$	$0.18 \pm 0.00$	$1.56 \pm 0.89$	$0.22 \pm 0.02$
1595	$bw = 0.2$	$0.07 \pm 0.01$	$0.48 \pm 0.06$	$0.19 \pm 0.00$	$4.61 \pm 2.05$	$0.22 \pm 0.02$
1596	$bw = 0.25$	$0.06 \pm 0.01$	$0.46 \pm 0.07$	$0.19 \pm 0.00$	$5.80 \pm 1.61$	$0.22 \pm 0.02$
1597	$bw = 0.3$	$0.05 \pm 0.01$	$0.44 \pm 0.07$	$0.19 \pm 0.00$	$7.30 \pm 1.29$	$0.22 \pm 0.02$
1598	$bw = 0.35$	$0.04 \pm 0.01$	$0.44 \pm 0.06$	$0.19 \pm 0.00$	$11.46 \pm 1.56$	$0.22 \pm 0.02$
1599	$bw = 0.4$	$0.04 \pm 0.01$	$0.42 \pm 0.09$	$0.19 \pm 0.00$	$13.16 \pm 1.83$	$0.22 \pm 0.02$
1600	$bw = 0.45$	$0.04 \pm 0.01$	$0.42 \pm 0.08$	$0.20 \pm 0.00$	$15.65 \pm 2.91$	$0.22 \pm 0.02$
1601	FeatureTweak	$0.12 \pm 0.03$	$0.29 \pm 0.25$	$0.21 \pm 0.02$	$0.58 \pm 0.16$	$0.24 \pm 0.03$
1602	FOCUS	$0.11 \pm 0.01$	$0.34 \pm 0.09$	$0.20 \pm 0.00$	$5.21 \pm 2.31$	$0.24 \pm 0.02$
1603	FACE	$0.17 \pm 0.01$	$0.81 \pm 0.02$	$0.17 \pm 0.00$	$0.46 \pm 0.08$	$0.24 \pm 0.02$

Table 7: CFX generation results, CaliforniaHousing, RandomForest. Methods with nan values had 100% failures. Validity 58% for FeatureTweak.

1606  
1607  
1608  
1609  
1610  
1611  
1612  
1613  
1614  
1615  
1616  
1617  
1618  
1619

1620 D.3.4 CONFORMAL EVALUATION RESULTS  
1621  
1622  
1623  
1624  
1625  
1626  
1627  
1628  
1629  
1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637  
1638

1639 Generator	1640 Marginal CovGap	1641 Binning CovGap	1642 Class Cond CovGap	1643 Simulated CovGap
<b>MLP</b>				
1641 ConfexNaive				
1642 $\alpha = 0.01$	$0.99 \pm 0.00$	$-0.35 \pm 0.66$	$-0.38 \pm 0.63$	$-8.61 \pm 1.51$
1643 $\alpha = 0.05$	$0.96 \pm 0.02$	$0.22 \pm 1.84$	$0.17 \pm 1.78$	$-12.83 \pm 4.89$
1644 $\alpha = 0.1$	$0.92 \pm 0.02$	$-0.19 \pm 1.76$	$-0.34 \pm 1.71$	$-22.79 \pm 7.76$
1645 ConfexTree, $\alpha = 0.01$				
1646 $bw = 0.05$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
1647 $bw = 0.1$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
1648 $bw = 0.15$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
1649 $bw = 0.2$	$1.00 \pm 0.00$	$0.91 \pm 0.02$	$0.90 \pm 0.02$	$-1.04 \pm 0.49$
1650 $bw = 0.25$	$1.00 \pm 0.00$	$0.81 \pm 0.04$	$0.80 \pm 0.04$	$-1.87 \pm 0.84$
1651 $bw = 0.3$	$1.00 \pm 0.00$	$0.75 \pm 0.03$	$0.74 \pm 0.04$	$-2.60 \pm 0.66$
1652 $bw = 0.35$	$1.00 \pm 0.00$	$0.60 \pm 0.08$	$0.57 \pm 0.08$	$-3.98 \pm 1.54$
1653 $bw = 0.4$	$1.00 \pm 0.00$	$0.34 \pm 0.05$	$0.30 \pm 0.06$	$-4.76 \pm 1.45$
1654 $bw = 0.45$	$1.00 \pm 0.00$	$0.32 \pm 0.05$	$0.28 \pm 0.06$	$-4.77 \pm 1.45$
1655 ConfexTree, $\alpha = 0.05$				
1656 $bw = 0.05$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$
1657 $bw = 0.1$	$1.00 \pm 0.00$	$4.91 \pm 0.00$	$4.90 \pm 0.00$	$-2.29 \pm 0.31$
1658 $bw = 0.15$	$1.00 \pm 0.00$	$4.82 \pm 0.03$	$4.81 \pm 0.03$	$-10.12 \pm 1.42$
1659 $bw = 0.2$	$1.00 \pm 0.00$	$4.67 \pm 0.04$	$4.65 \pm 0.04$	$-10.19 \pm 5.66$
1660 $bw = 0.25$	$0.99 \pm 0.00$	$4.35 \pm 0.08$	$4.32 \pm 0.07$	$-7.64 \pm 4.82$
1661 $bw = 0.3$	$0.99 \pm 0.00$	$3.83 \pm 0.12$	$3.76 \pm 0.11$	$-11.33 \pm 3.95$
1662 $bw = 0.35$	$0.97 \pm 0.01$	$1.55 \pm 0.27$	$1.35 \pm 0.29$	$-13.37 \pm 2.89$
1663 $bw = 0.4$	$0.95 \pm 0.00$	$0.09 \pm 0.47$	$-0.19 \pm 0.53$	$-14.08 \pm 2.06$
1664 $bw = 0.45$	$0.95 \pm 0.00$	$0.07 \pm 0.47$	$-0.22 \pm 0.53$	$-14.06 \pm 1.97$
1665 ConfexTree, $\alpha = 0.1$				
1666 $bw = 0.05$	$1.00 \pm 0.00$	$9.99 \pm 0.01$	$9.99 \pm 0.01$	$-0.56 \pm 14.46$
1667 $bw = 0.1$	$1.00 \pm 0.00$	$9.76 \pm 0.08$	$9.74 \pm 0.08$	$3.16 \pm 0.88$
1668 $bw = 0.15$	$1.00 \pm 0.00$	$9.54 \pm 0.21$	$9.52 \pm 0.22$	$-10.92 \pm 5.77$
1669 $bw = 0.2$	$0.99 \pm 0.00$	$8.98 \pm 0.14$	$8.93 \pm 0.15$	$-19.47 \pm 4.74$
1670 $bw = 0.25$	$0.99 \pm 0.00$	$8.49 \pm 0.25$	$8.42 \pm 0.24$	$-11.00 \pm 2.93$
1671 $bw = 0.3$	$0.98 \pm 0.00$	$7.55 \pm 0.38$	$7.44 \pm 0.35$	$-14.15 \pm 3.08$
1672 $bw = 0.35$	$0.91 \pm 0.01$	$1.87 \pm 0.80$	$1.41 \pm 0.90$	$-22.28 \pm 6.52$
1673 $bw = 0.4$	$0.91 \pm 0.01$	$1.01 \pm 0.82$	$0.51 \pm 0.92$	$-22.27 \pm 5.75$
1674 $bw = 0.45$	$0.91 \pm 0.01$	$0.97 \pm 0.83$	$0.47 \pm 0.93$	$-22.19 \pm 5.64$

Table 8: Conformal evaluation results, CaliforniaHousing, MLP

1674	Generator	Marginal CovGap	Binning CovGap	Class Cond CovGap	Simulated CovGap
<b>RandomForest</b>					
<b>ConfexNaive</b>					
1678	$\alpha = 0.01$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
1679	$\alpha = 0.05$	$0.95 \pm 0.01$	$0.16 \pm 0.53$	$-0.10 \pm 0.56$	$-39.17 \pm 8.15$
1680	$\alpha = 0.1$	$0.95 \pm 0.01$	$5.16 \pm 0.53$	$4.90 \pm 0.56$	$-34.17 \pm 8.15$
<b>ConfexTree, <math>\alpha = 0.01</math></b>					
1681	$bw = 0.05$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
1682	$bw = 0.1$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm nan$
1683	$bw = 0.15$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm nan$
1684	$bw = 0.2$	$1.00 \pm 0.00$	$0.97 \pm 0.02$	$0.97 \pm 0.03$	$0.79 \pm 0.41$
1685	$bw = 0.25$	$1.00 \pm 0.00$	$0.97 \pm 0.03$	$0.97 \pm 0.03$	$0.66 \pm 0.65$
1686	$bw = 0.3$	$1.00 \pm 0.00$	$0.97 \pm 0.03$	$0.97 \pm 0.03$	$0.66 \pm 0.65$
1687	$bw = 0.35$	$1.00 \pm 0.00$	$0.96 \pm 0.07$	$0.96 \pm 0.07$	$-7.24 \pm 2.12$
1688	$bw = 0.4$	$1.00 \pm 0.00$	$0.83 \pm 0.23$	$0.83 \pm 0.23$	$-33.35 \pm 31.24$
1689	$bw = 0.45$	$1.00 \pm 0.00$	$0.83 \pm 0.23$	$0.83 \pm 0.23$	$-30.80 \pm 27.63$
<b>ConfexTree, <math>\alpha = 0.05</math></b>					
1690	$bw = 0.05$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$
1691	$bw = 0.1$	$1.00 \pm 0.00$	$4.93 \pm 0.05$	$4.93 \pm 0.05$	$0.93 \pm 3.99$
1692	$bw = 0.15$	$1.00 \pm 0.00$	$4.88 \pm 0.05$	$4.87 \pm 0.06$	$-12.72 \pm 5.73$
1693	$bw = 0.2$	$1.00 \pm 0.00$	$4.76 \pm 0.10$	$4.75 \pm 0.10$	$-25.53 \pm 19.62$
1694	$bw = 0.25$	$1.00 \pm 0.00$	$4.48 \pm 0.10$	$4.46 \pm 0.10$	$-33.60 \pm 11.02$
1695	$bw = 0.3$	$0.99 \pm 0.00$	$4.00 \pm 0.25$	$3.98 \pm 0.26$	$-37.57 \pm 6.37$
1696	$bw = 0.35$	$0.97 \pm 0.01$	$2.67 \pm 0.31$	$2.56 \pm 0.32$	$-35.59 \pm 3.75$
1697	$bw = 0.4$	$0.97 \pm 0.01$	$2.54 \pm 0.46$	$2.43 \pm 0.47$	$-34.96 \pm 6.88$
1698	$bw = 0.45$	$0.97 \pm 0.01$	$2.51 \pm 0.45$	$2.40 \pm 0.46$	$-35.47 \pm 6.05$
<b>ConfexTree, <math>\alpha = 0.1</math></b>					
1700	$bw = 0.05$	$1.00 \pm 0.00$	$10.00 \pm 0.01$	$10.00 \pm 0.01$	$4.72 \pm 11.81$
1701	$bw = 0.1$	$1.00 \pm 0.00$	$9.81 \pm 0.05$	$9.80 \pm 0.05$	$-3.35 \pm 5.87$
1702	$bw = 0.15$	$1.00 \pm 0.00$	$9.60 \pm 0.23$	$9.57 \pm 0.24$	$-17.34 \pm 10.55$
1703	$bw = 0.2$	$1.00 \pm 0.00$	$9.33 \pm 0.04$	$9.31 \pm 0.04$	$-29.85 \pm 6.39$
1704	$bw = 0.25$	$0.99 \pm 0.01$	$8.51 \pm 0.40$	$8.47 \pm 0.42$	$-29.75 \pm 8.83$
1705	$bw = 0.3$	$0.98 \pm 0.01$	$7.57 \pm 0.56$	$7.53 \pm 0.58$	$-33.10 \pm 7.00$
1706	$bw = 0.35$	$0.93 \pm 0.01$	$3.39 \pm 0.99$	$3.16 \pm 1.02$	$-44.42 \pm 5.68$
1707	$bw = 0.4$	$0.93 \pm 0.01$	$2.95 \pm 1.07$	$2.69 \pm 1.10$	$-43.75 \pm 6.11$
1708	$bw = 0.45$	$0.93 \pm 0.01$	$2.84 \pm 1.02$	$2.58 \pm 1.05$	$-43.73 \pm 5.96$

Table 9: Conformal evaluation results, CaliforniaHousing, RandomForest

1728  
1729

## D.4 GERMAN CREDIT

1730  
1731  
1732  
1733  
1734

We use the German Credit dataset from the UCI Machine Learning Repository Hofmann (1994), with a cleaned version obtained through Kaggle<sup>6</sup>. The preprocessing included: (i) scaling numeric features (Age, Credit amount, Duration) to (0, 1) using MinMax scaling, (ii) ordinal encoding of categorical features (job, savings account, checking account), then normalised. The Purpose feature was dropped.

1735  
1736  
1737

Our results show that distance decreases and plausibility decreases as the kernel bandwidth increases. When the bandwidth is properly tuned, CONFEX methods outperform all other methods on plausibility and sensitivity.

1738

## D.4.1 PLOTS

1739

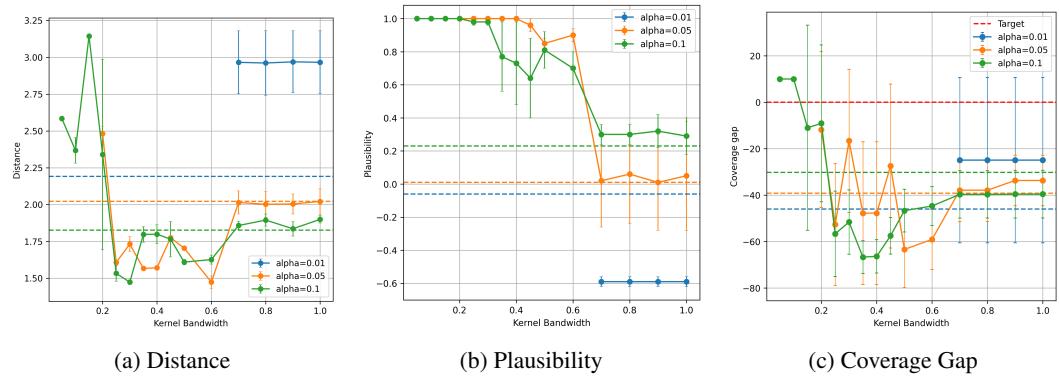
1740  
1741  
1742  
1743  
1744  
1745  
1746  
1747  
1748  
1749  
1750  
1751

Figure 11: Effect of coverage rate and kernel bandwidth on metrics for CONFEX-Tree on the GermanCredit dataset, MLP. CONFEX-Naive is represented by dashed horizontal lines.

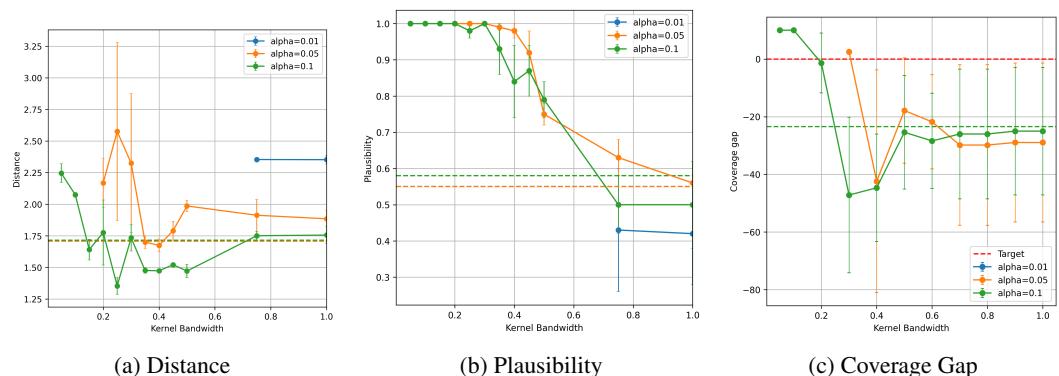
1752  
1753  
1754  
1755  
1756  
1757

Figure 12: Effect of coverage rate and kernel bandwidth on metrics for CONFEX-Tree on the GermanCredit dataset, RandomForest. CONFEX-Naive is represented by dashed horizontal lines.

1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781

<sup>6</sup><https://www.kaggle.com/datasets/uciml/german-credit/data>

1782 D.4.2 MODEL EVALUATION RESULTS  
1783

1784 Repeat	1785 Accuracy (%)	1786 Precision (%)	1787 F1 Score (%)
1788 repeat0,MLP	1789 72.00	1790 72.00	1791 72.00
1792 repeat1,MLP	1793 71.00	1794 70.01	1795 70.39
1796 repeat2,MLP	1797 72.00	1798 72.00	1799 72.00
1800 repeat3,MLP	1801 73.00	1802 75.40	1803 73.74
1804 repeat4,MLP	1805 71.50	1806 71.12	1807 71.29
1808 repeat0,RF	1809 70.00	1810 68.27	1811 68.77
1812 repeat1,RF	1813 69.50	1814 68.31	1815 68.76
1816 repeat2,RF	1817 70.00	1818 68.27	1819 68.77
1820 repeat3,RF	1821 68.50	1822 67.26	1823 67.74
1824 repeat4,RF	1825 72.50	1826 70.28	1827 69.81

1828 Table 10: Model evaluation results, GermanCredit.  
1829  
1830  
1831  
1832  
1833  
1834  
1835

1836 D.4.3 CFX GENERATION RESULTS  
1837  
1838  
1839  
1840  
1841  
1842  
1843  
1844  
1845  
1846

1847 Generator	1848 Distance	1849 Plausibility	1850 Implausibility	1851 Sensitivity ( $10^{-1}$ )	1852 Stability
<b>MLP</b>					
1849 MinDist	$1.65 \pm 0.18$	$0.54 \pm 0.19$	$0.71 \pm 0.06$	$0.08 \pm 0.02$	$0.58 \pm 0.04$
1850 Wachter	$0.41 \pm 0.02$	$0.73 \pm 0.05$	$0.59 \pm 0.02$	$0.33 \pm 0.09$	$0.24 \pm 0.02$
1851 Greedy	$0.99 \pm 0.04$	$-0.03 \pm 0.09$	$0.80 \pm 0.01$	$0.08 \pm 0.02$	$0.68 \pm 0.03$
ConfexNaive					
1853 $\alpha = 0.01$	$2.26 \pm 0.14$	$-0.16 \pm 0.38$	$0.88 \pm 0.06$	$0.03 \pm 0.01$	$0.97 \pm 0.02$
1854 $\alpha = 0.05$	$2.00 \pm 0.07$	$0.16 \pm 0.28$	$0.80 \pm 0.08$	$0.05 \pm 0.01$	$0.83 \pm 0.07$
1855 $\alpha = 0.1$	$1.80 \pm 0.04$	$0.40 \pm 0.22$	$0.72 \pm 0.05$	$0.06 \pm 0.02$	$0.72 \pm 0.09$
ECCCo					
1856 $\alpha = 0.01$	$1.01 \pm 0.07$	$0.12 \pm 0.11$	$0.77 \pm 0.03$	$0.06 \pm 0.02$	$0.73 \pm 0.01$
1857 $\alpha = 0.05$	$1.00 \pm 0.06$	$0.08 \pm 0.14$	$0.77 \pm 0.02$	$0.05 \pm 0.02$	$0.73 \pm 0.01$
1858 $\alpha = 0.1$	$0.97 \pm 0.06$	$0.16 \pm 0.11$	$0.75 \pm 0.02$	$0.06 \pm 0.02$	$0.72 \pm 0.01$
ConfexTree, $\alpha = 0.01$					
1860 $bw = 0.05$	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan
1861 $bw = 0.6$	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan
1862 $bw = 0.8$	$2.97 \pm 0.35$	$-0.58 \pm 0.27$	$0.87 \pm 0.10$	$0.02 \pm 0.01$	$0.99 \pm 0.01$
1863 $bw = 1$	$2.97 \pm 0.35$	$-0.57 \pm 0.27$	$0.87 \pm 0.10$	$0.02 \pm 0.01$	$0.99 \pm 0.01$
1864 $bw = 1.2$	$2.97 \pm 0.21$	$-0.60 \pm 0.04$	$0.84 \pm 0.07$	$0.03 \pm 0.01$	$0.98 \pm 0.00$
1865 $bw = 1.4$	$2.97 \pm 0.21$	$-0.60 \pm 0.04$	$0.84 \pm 0.07$	$0.03 \pm 0.01$	$0.98 \pm 0.00$
1866 $bw = 1.6$	$2.97 \pm 0.21$	$-0.60 \pm 0.04$	$0.84 \pm 0.07$	$0.03 \pm 0.01$	$0.98 \pm 0.00$
ConfexTree, $\alpha = 0.05$					
1867 $bw = 0.1$	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan
1868 $bw = 0.2$	$2.31 \pm 0.20$	$1.00 \pm 0.00$	$0.28 \pm 0.02$	$0.01 \pm 0.01$	$0.88 \pm 0.03$
1869 $bw = 0.4$	$1.60 \pm 0.16$	$1.00 \pm 0.00$	$0.44 \pm 0.02$	$0.03 \pm 0.01$	$0.69 \pm 0.04$
1870 $bw = 0.6$	$1.48 \pm 0.10$	$0.88 \pm 0.04$	$0.59 \pm 0.02$	$0.06 \pm 0.02$	$0.65 \pm 0.07$
1871 $bw = 0.8$	$1.91 \pm 0.10$	$0.26 \pm 0.27$	$0.75 \pm 0.05$	$0.05 \pm 0.01$	$0.79 \pm 0.07$
1872 $bw = 1$	$1.94 \pm 0.09$	$0.23 \pm 0.29$	$0.76 \pm 0.06$	$0.05 \pm 0.01$	$0.80 \pm 0.07$
1873 $bw = 1.2$	$2.00 \pm 0.08$	$0.04 \pm 0.32$	$0.79 \pm 0.05$	$0.05 \pm 0.00$	$0.81 \pm 0.07$
1874 $bw = 1.4$	$2.01 \pm 0.08$	$0.05 \pm 0.31$	$0.80 \pm 0.05$	$0.05 \pm 0.00$	$0.81 \pm 0.06$
1875 $bw = 1.6$	$2.01 \pm 0.08$	$0.05 \pm 0.31$	$0.80 \pm 0.05$	$0.05 \pm 0.00$	$0.81 \pm 0.06$
ConfexTree, $\alpha = 0.1$					
1876 $bw = 0.1$	$2.18 \pm 0.21$	$1.00 \pm 0.00$	$0.25 \pm 0.03$	$0.01 \pm 0.00$	$0.76 \pm 0.11$
1877 $bw = 0.2$	$2.21 \pm 0.71$	$1.00 \pm 0.00$	$0.37 \pm 0.01$	$0.01 \pm 0.00$	$0.66 \pm 0.08$
1878 $bw = 0.4$	$1.67 \pm 0.15$	$0.83 \pm 0.18$	$0.62 \pm 0.05$	$0.07 \pm 0.01$	$0.57 \pm 0.08$
1879 $bw = 0.6$	$1.61 \pm 0.08$	$0.84 \pm 0.13$	$0.62 \pm 0.04$	$0.07 \pm 0.01$	$0.64 \pm 0.07$
1880 $bw = 0.8$	$1.79 \pm 0.12$	$0.45 \pm 0.14$	$0.71 \pm 0.05$	$0.07 \pm 0.01$	$0.73 \pm 0.06$
1881 $bw = 1$	$1.80 \pm 0.10$	$0.47 \pm 0.18$	$0.71 \pm 0.05$	$0.07 \pm 0.01$	$0.74 \pm 0.06$
1882 $bw = 1.2$	$1.86 \pm 0.03$	$0.31 \pm 0.09$	$0.74 \pm 0.01$	$0.07 \pm 0.01$	$0.73 \pm 0.04$
1883 $bw = 1.4$	$1.86 \pm 0.03$	$0.28 \pm 0.10$	$0.75 \pm 0.02$	$0.07 \pm 0.02$	$0.73 \pm 0.05$
1884 $bw = 1.5$	$1.86 \pm 0.03$	$0.28 \pm 0.10$	$0.75 \pm 0.02$	$0.07 \pm 0.02$	$0.73 \pm 0.05$
1885 $bw = 1.6$	$1.86 \pm 0.03$	$0.28 \pm 0.10$	$0.75 \pm 0.02$	$0.07 \pm 0.02$	$0.73 \pm 0.05$
1886 FACE	$0.69 \pm 0.05$	$0.92 \pm 0.05$	$0.45 \pm 0.02$	$0.05 \pm 0.01$	$0.43 \pm 0.01$
1887 C-CHVAE	$2.45 \pm 0.13$	$0.80 \pm 0.19$	$0.52 \pm 0.05$	$0.07 \pm 0.01$	$0.38 \pm 0.06$

1888 Table 11: CFX generation results, GermanCredit, MLP. Methods with nan values had 100% failures.  
1889 Validity 84% for all ECCCo methods, 82% for Greedy, 84% for Wachter, 74% for C-CHVAE.

1890	Generator	Distance	Plausibility	Implausibility	Sensitivity ( $10^{-1}$ )	Stability
<b>RandomForest</b>						
1893	MinDist	$1.65 \pm 0.06$	$0.52 \pm 0.17$	$0.71 \pm 0.05$	$0.09 \pm 0.01$	$0.36 \pm 0.01$
1894	ConfexNaive					
1895	$\alpha = 0.01$	$1.62 \pm 0.08$	$0.63 \pm 0.10$	$0.66 \pm 0.04$	$0.09 \pm 0.01$	$0.43 \pm 0.02$
1896	$\alpha = 0.05$	$1.62 \pm 0.08$	$0.63 \pm 0.10$	$0.66 \pm 0.04$	$0.09 \pm 0.01$	$0.43 \pm 0.02$
1897	$\alpha = 0.1$	$1.62 \pm 0.09$	$0.65 \pm 0.09$	$0.66 \pm 0.04$	$0.09 \pm 0.01$	$0.43 \pm 0.02$
1898	ConfexTree, $\alpha = 0.01$					
1899	$bw = 0.05$	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan
1900	$bw = 0.7$	$2.23 \pm 0.11$	$0.45 \pm 0.09$	$0.61 \pm 0.03$	$0.08 \pm 0.01$	$0.42 \pm 0.02$
1901	$bw = 0.8$	$2.23 \pm 0.11$	$0.46 \pm 0.12$	$0.61 \pm 0.03$	$0.07 \pm 0.01$	$0.42 \pm 0.02$
1902	$bw = 0.9$	$2.22 \pm 0.11$	$0.44 \pm 0.10$	$0.61 \pm 0.03$	$0.07 \pm 0.01$	$0.42 \pm 0.03$
1903	$bw = 1$	$2.22 \pm 0.11$	$0.44 \pm 0.10$	$0.61 \pm 0.03$	$0.07 \pm 0.01$	$0.42 \pm 0.03$
1904	$bw = 1.1$	$2.39 \pm 0.00$	$0.28 \pm 0.00$	$0.67 \pm 0.00$	$0.09 \pm 0.00$	$0.43 \pm 0.00$
1905	$bw = 1.2$	$2.39 \pm 0.00$	$0.28 \pm 0.00$	$0.67 \pm 0.00$	$0.09 \pm 0.00$	$0.43 \pm 0.00$
1906	ConfexTree, $\alpha = 0.05$					
1907	$bw = 0.1$	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan	nan $\pm$ nan
1908	$bw = 0.2$	$2.04 \pm 0.16$	$1.00 \pm 0.00$	$0.35 \pm 0.06$	$0.01 \pm 0.01$	$0.44 \pm 0.26$
1909	$bw = 0.3$	$2.12 \pm 0.43$	$1.00 \pm 0.00$	$0.39 \pm 0.06$	$0.07 \pm 0.02$	$0.38 \pm 0.05$
1910	$bw = 0.4$	$1.73 \pm 0.07$	$0.99 \pm 0.02$	$0.45 \pm 0.07$	$0.07 \pm 0.05$	$0.35 \pm 0.04$
1911	$bw = 0.5$	$1.88 \pm 0.13$	$0.80 \pm 0.08$	$0.54 \pm 0.04$	$0.05 \pm 0.01$	$0.39 \pm 0.03$
1912	$bw = 0.6$	$1.40 \pm 0.13$	$0.82 \pm 0.13$	$0.57 \pm 0.05$	$0.07 \pm 0.01$	$0.39 \pm 0.03$
1913	$bw = 0.7$	$1.63 \pm 0.10$	$0.72 \pm 0.14$	$0.61 \pm 0.07$	$0.11 \pm 0.03$	$0.41 \pm 0.02$
1914	$bw = 0.8$	$1.71 \pm 0.15$	$0.70 \pm 0.16$	$0.63 \pm 0.06$	$0.11 \pm 0.03$	$0.43 \pm 0.03$
1915	$bw = 0.9$	$1.65 \pm 0.08$	$0.74 \pm 0.09$	$0.62 \pm 0.05$	$0.11 \pm 0.02$	$0.43 \pm 0.04$
1916	$bw = 1$	$2.10 \pm 0.64$	$0.59 \pm 0.15$	$0.71 \pm 0.07$	$0.09 \pm 0.02$	$0.45 \pm 0.03$
1917	ConfexTree, $\alpha = 0.1$					
1918	$bw = 0.1$	$2.01 \pm 0.07$	$1.00 \pm 0.00$	$0.30 \pm 0.05$	$0.01 \pm 0.01$	$0.45 \pm 0.24$
1919	$bw = 0.2$	$1.76 \pm 0.17$	$1.00 \pm 0.00$	$0.38 \pm 0.02$	$0.04 \pm 0.02$	$0.25 \pm 0.06$
1920	$bw = 0.3$	$1.48 \pm 0.23$	$0.99 \pm 0.01$	$0.49 \pm 0.04$	$0.10 \pm 0.05$	$0.35 \pm 0.03$
1921	$bw = 0.4$	$1.39 \pm 0.08$	$0.87 \pm 0.09$	$0.55 \pm 0.02$	$0.08 \pm 0.01$	$0.36 \pm 0.02$
1922	$bw = 0.5$	$1.46 \pm 0.10$	$0.76 \pm 0.07$	$0.60 \pm 0.03$	$0.09 \pm 0.01$	$0.37 \pm 0.03$
1923	$bw = 0.6$	$1.49 \pm 0.12$	$0.77 \pm 0.08$	$0.61 \pm 0.01$	$0.13 \pm 0.10$	$0.38 \pm 0.02$
1924	$bw = 0.7$	$1.57 \pm 0.10$	$0.67 \pm 0.14$	$0.64 \pm 0.05$	$0.10 \pm 0.01$	$0.41 \pm 0.01$
1925	$bw = 0.8$	$1.58 \pm 0.09$	$0.66 \pm 0.14$	$0.65 \pm 0.04$	$0.09 \pm 0.01$	$0.42 \pm 0.01$
1926	$bw = 0.9$	$1.61 \pm 0.12$	$0.64 \pm 0.15$	$0.65 \pm 0.05$	$0.09 \pm 0.01$	$0.42 \pm 0.01$
1927	$bw = 1$	$1.63 \pm 0.11$	$0.66 \pm 0.16$	$0.66 \pm 0.05$	$0.09 \pm 0.01$	$0.42 \pm 0.02$
1928	FeatureTweak	$0.50 \pm 0.06$	$0.84 \pm 0.05$	$0.55 \pm 0.03$	$0.09 \pm 0.01$	$0.20 \pm 0.02$
1929	FOCUS	$0.45 \pm 0.14$	$0.83 \pm 0.02$	$0.56 \pm 0.02$	$0.58 \pm 0.25$	$0.26 \pm 0.02$
1930	FACE	$0.59 \pm 0.06$	$0.88 \pm 0.07$	$0.48 \pm 0.01$	$0.07 \pm 0.01$	$0.22 \pm 0.02$

Table 12: CFX generation results, GermanCredit, RandomForest. Methods with nan values had 100% failures. Validity 50% for FeatureTweak.

1931  
1932  
1933  
1934  
1935  
1936  
1937  
1938  
1939  
1940  
1941  
1942  
1943

1944  
1945  
1946  
1947  
1948  
1949  
1950  
1951  
1952  
1953  
1954  
1955  
1956

## D.4.4 CONFORMAL EVALUATION RESULTS

1957  
1958

Generator	Marginal CovGap	Binning CovGap	Class Cond CovGap	Simulated CovGap
<b>MLP</b>				
<b>ConfexNaive</b>				
$\alpha = 0.01$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$\alpha = 0.05$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$nan \pm nan$
$\alpha = 0.1$	$0.93 \pm 0.03$	$-0.81 \pm 0.66$	$2.90 \pm 0.42$	$-23.40 \pm 21.12$
<b>ConfexTree, <math>\alpha = 0.01</math></b>				
$bw = 0.05$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.1$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.2$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.3$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.4$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.5$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.6$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.7$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.8$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 0.9$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
$bw = 1$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
<b>ConfexTree, <math>\alpha = 0.05</math></b>				
$bw = 0.05$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$nan \pm nan$
$bw = 0.1$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$nan \pm nan$
$bw = 0.2$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$nan \pm nan$
$bw = 0.3$	$1.00 \pm 0.00$	$4.93 \pm 0.16$	$4.90 \pm 0.22$	$2.50 \pm nan$
$bw = 0.4$	$0.99 \pm 0.02$	$3.31 \pm 1.53$	$3.80 \pm 1.04$	$-42.39 \pm 38.61$
$bw = 0.5$	$0.95 \pm 0.05$	$2.58 \pm 1.90$	$3.30 \pm 1.44$	$-17.85 \pm 18.25$
$bw = 0.6$	$0.95 \pm 0.05$	$2.42 \pm 1.75$	$3.20 \pm 1.35$	$-21.79 \pm 16.32$
$bw = 0.7$	$0.97 \pm 0.03$	$2.23 \pm 3.11$	$3.20 \pm 2.14$	$-29.83 \pm 27.83$
$bw = 0.8$	$0.97 \pm 0.03$	$2.23 \pm 3.11$	$3.20 \pm 2.14$	$-29.83 \pm 27.83$
$bw = 0.9$	$0.97 \pm 0.03$	$2.23 \pm 3.11$	$3.20 \pm 2.14$	$-28.95 \pm 27.52$
$bw = 1$	$0.97 \pm 0.03$	$2.23 \pm 3.11$	$3.20 \pm 2.14$	$-28.95 \pm 27.52$
<b>ConfexTree, <math>\alpha = 0.1</math></b>				
$bw = 0.05$	$1.00 \pm 0.00$	$10.00 \pm 0.00$	$10.00 \pm 0.00$	$10.00 \pm 0.00$
$bw = 0.1$	$1.00 \pm 0.00$	$10.00 \pm 0.00$	$10.00 \pm 0.00$	$10.00 \pm 0.00$
$bw = 0.2$	$0.99 \pm 0.02$	$8.51 \pm 0.82$	$8.70 \pm 0.76$	$-1.36 \pm 10.37$
$bw = 0.3$	$0.95 \pm 0.05$	$4.33 \pm 0.60$	$5.70 \pm 0.67$	$-47.17 \pm 26.96$
$bw = 0.4$	$0.95 \pm 0.06$	$1.56 \pm 1.82$	$3.90 \pm 1.64$	$-44.65 \pm 18.65$
$bw = 0.5$	$0.91 \pm 0.04$	$-0.34 \pm 0.58$	$2.40 \pm 0.55$	$-25.40 \pm 19.69$
$bw = 0.6$	$0.92 \pm 0.04$	$0.76 \pm 1.33$	$3.30 \pm 1.15$	$-28.40 \pm 16.51$
$bw = 0.7$	$0.93 \pm 0.03$	$-0.08 \pm 1.43$	$3.40 \pm 0.82$	$-26.00 \pm 22.49$
$bw = 0.8$	$0.93 \pm 0.03$	$-0.08 \pm 1.43$	$3.40 \pm 0.82$	$-26.00 \pm 22.49$
$bw = 0.9$	$0.93 \pm 0.03$	$-0.08 \pm 1.43$	$3.40 \pm 0.82$	$-25.00 \pm 22.10$
$bw = 1$	$0.93 \pm 0.03$	$-0.08 \pm 1.43$	$3.40 \pm 0.82$	$-25.00 \pm 22.10$

1997

Table 13: Conformal evaluation results, GermanCredit, MLP

1998	Generator	Marginal CovGap	Binning CovGap	Class Cond CovGap	Simulated CovGap
1999	<b>RandomForest</b>				
2000	<b>ConfexNaive</b>				
2001	$\alpha = 0.01$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2002	$\alpha = 0.05$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$nan \pm nan$
2003	$\alpha = 0.1$	$0.92 \pm 0.04$	$-0.97 \pm 1.23$	$2.75 \pm 0.35$	$-33.75 \pm 12.37$
2004	<b>ConfexTree, <math>\alpha = 0.01</math></b>				
2005	$bw = 0.1$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2006	$bw = 0.2$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2007	$bw = 0.3$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2008	$bw = 0.4$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2009	$bw = 0.5$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2010	$bw = 0.6$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2011	$bw = 0.7$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2012	$bw = 0.8$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2013	$bw = 0.9$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2014	$bw = 1$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$nan \pm nan$
2015	<b>ConfexTree, <math>\alpha = 0.05</math></b>				
2016	$bw = 0.1$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$nan \pm nan$
2017	$bw = 0.2$	$1.00 \pm 0.00$	$5.00 \pm 0.00$	$5.00 \pm 0.00$	$nan \pm nan$
2018	$bw = 0.3$	$1.00 \pm 0.00$	$4.82 \pm 0.25$	$4.75 \pm 0.35$	$2.50 \pm nan$
2019	$bw = 0.4$	$1.00 \pm 0.00$	$4.82 \pm 0.25$	$4.75 \pm 0.35$	$2.50 \pm nan$
2020	$bw = 0.5$	$0.98 \pm 0.04$	$2.13 \pm 1.74$	$3.25 \pm 1.06$	$-16.31 \pm 11.59$
2021	$bw = 0.6$	$0.98 \pm 0.04$	$1.72 \pm 1.16$	$3.00 \pm 0.71$	$-26.15 \pm 2.32$
2022	$bw = 0.7$	$0.95 \pm 0.00$	$0.13 \pm 4.57$	$1.75 \pm 3.18$	$-51.18 \pm 17.93$
2023	$bw = 0.8$	$0.95 \pm 0.00$	$0.13 \pm 4.57$	$1.75 \pm 3.18$	$-51.18 \pm 17.93$
2024	$bw = 0.9$	$0.95 \pm 0.00$	$0.13 \pm 4.57$	$1.75 \pm 3.18$	$-49.43 \pm 20.40$
2025	$bw = 1$	$0.95 \pm 0.00$	$0.13 \pm 4.57$	$1.75 \pm 3.18$	$-49.43 \pm 20.40$
2026	<b>ConfexTree, <math>\alpha = 0.1</math></b>				
2027	$bw = 0.1$	$1.00 \pm 0.00$	$10.00 \pm 0.00$	$10.00 \pm 0.00$	$10.00 \pm nan$
2028	$bw = 0.2$	$1.00 \pm 0.00$	$9.82 \pm 0.25$	$9.75 \pm 0.35$	$7.50 \pm nan$
2029	$bw = 0.3$	$0.98 \pm 0.04$	$5.59 \pm 1.23$	$6.75 \pm 0.35$	$-36.75 \pm 49.85$
2030	$bw = 0.4$	$0.98 \pm 0.04$	$3.31 \pm 0.33$	$5.50 \pm 0.00$	$-36.75 \pm 48.44$
2031	$bw = 0.5$	$0.95 \pm 0.00$	$-0.10 \pm 2.32$	$3.00 \pm 1.41$	$-23.75 \pm 22.27$
2032	$bw = 0.6$	$0.95 \pm 0.00$	$-0.51 \pm 3.41$	$2.75 \pm 2.47$	$-19.50 \pm 12.73$
2033	$bw = 0.7$	$0.92 \pm 0.04$	$-0.79 \pm 1.48$	$3.00 \pm 0.71$	$-34.00 \pm 12.02$
2034	$bw = 0.8$	$0.92 \pm 0.04$	$-0.79 \pm 1.48$	$3.00 \pm 0.71$	$-34.00 \pm 12.02$
2035	$bw = 0.9$	$0.92 \pm 0.04$	$-0.79 \pm 1.48$	$3.00 \pm 0.71$	$-32.75 \pm 13.79$
2036	$bw = 1$	$0.92 \pm 0.04$	$-0.79 \pm 1.48$	$3.00 \pm 0.71$	$-32.75 \pm 13.79$

Table 14: Conformal evaluation results, GermanCredit, RandomForest

2037  
 2038  
 2039  
 2040  
 2041  
 2042  
 2043  
 2044  
 2045  
 2046  
 2047  
 2048  
 2049  
 2050  
 2051

2052  
2053

## D.5 GIVEMESOME CREDIT

This dataset, obtained through Kaggle<sup>7</sup>, contains credit scoring data with 8 numeric features that were scaled to (0, 1) using MinMax scaling.

2056 We find that CONFEX-Tree methods outperform all other methods on plausibility and sensitivity.  
2057 Our results roughly show the same pattern that distance decreases as the kernel bandwidth increases,  
2058 however plausibility remains quite good for all tested (small) choices of kernel bandwidth.

### 2060 D.5.1 PLOTS

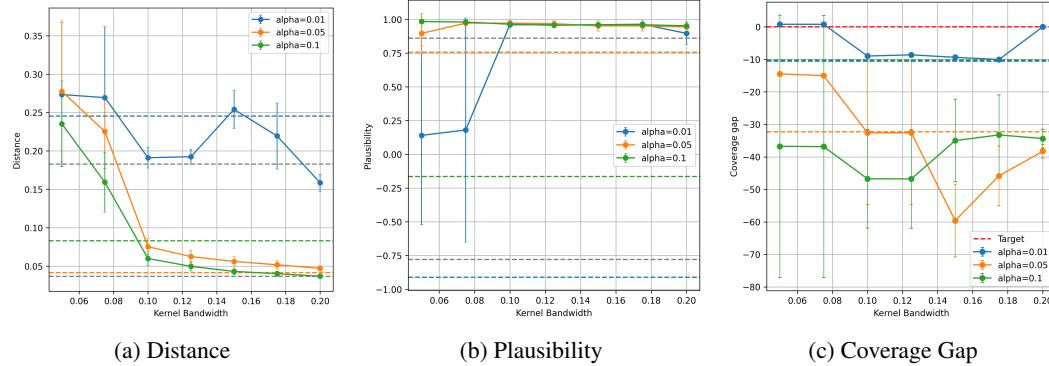


Figure 13: Effect of coverage rate and kernel bandwidth on metrics for CONFEX-Tree on the GiveMeSomeCredit dataset, MLP. CONFEX-Naive is represented by dashed horizontal lines.

2077

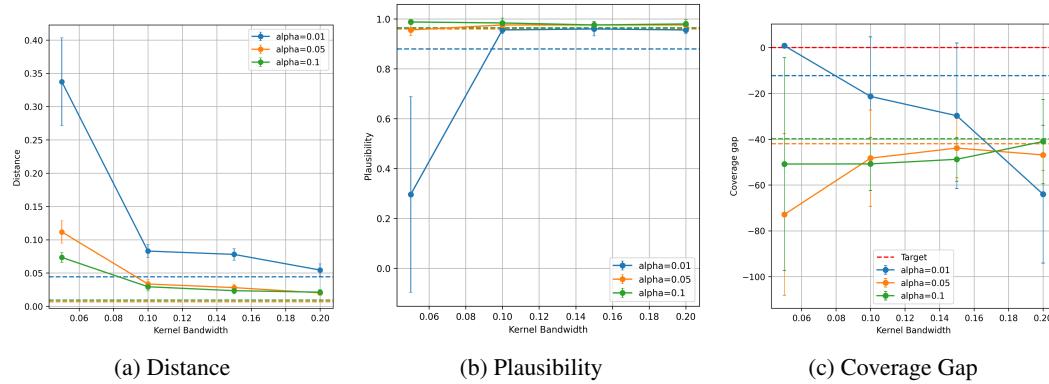


Figure 14: Effect of coverage rate and kernel bandwidth on metrics for CONFEX-Tree on the GiveMeSomeCredit dataset, RandomForest. CONFEX-Naive is represented by dashed horizontal lines.

2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105

<sup>7</sup><https://www.kaggle.com/competitions/GiveMeSomeCredit>

2106 D.5.2 MODEL EVALUATION RESULTS  
2107

2108	2109	Repeat	Accuracy (%)	Precision (%)	F1 Score (%)
2110		repeat0,MLP	93.54	91.82	91.81
2111		repeat1,MLP	93.49	91.79	91.96
2112		repeat2,MLP	93.61	91.98	91.87
2113		repeat3,MLP	93.57	91.89	91.72
2114		repeat4,MLP	93.48	91.73	91.84
2115		repeat0,RF	93.40	91.57	91.78
2116		repeat1,RF	93.40	91.53	91.69
2117		repeat2,RF	93.41	91.61	91.82
2118		repeat3,RF	93.37	91.49	91.68
2119		repeat4,RF	93.53	91.83	91.89

2122 Table 15: Model evaluation results, GiveMeSomeCredit.  
2123  
2124  
2125  
2126  
2127  
2128  
2129  
2130  
2131  
2132  
2133  
2134  
2135  
2136  
2137  
2138  
2139  
2140  
2141  
2142  
2143  
2144  
2145  
2146  
2147  
2148  
2149  
2150  
2151  
2152  
2153  
2154  
2155  
2156  
2157  
2158  
2159

2160 D.5.3 CFX GENERATION RESULTS  
2161  
2162  
2163  
2164  
2165  
2166  
2167  
2168  
2169  
2170  
2171

2172 Generator	2173 Distance	2173 Plausibility	2173 Implausibility	2173 Sensitivity ( $10^{-1}$ )	2173 Stability
<b>MLP</b>					
MinDist	$0.03 \pm 0.00$	$0.93 \pm 0.04$	$0.09 \pm 0.00$	$1.39 \pm 0.57$	$0.18 \pm 0.06$
Wachter	$0.09 \pm 0.01$	$0.93 \pm 0.02$	$0.09 \pm 0.00$	$0.95 \pm 0.07$	$0.18 \pm 0.07$
Greedy	$0.13 \pm 0.07$	$-0.02 \pm 0.38$	$0.14 \pm 0.03$	$1.05 \pm 0.81$	$0.17 \pm 0.06$
<b>ConfexNaive</b>					
$\alpha = 0.01$	$0.25 \pm 0.01$	$-0.91 \pm 0.05$	$0.21 \pm 0.01$	$0.21 \pm 0.05$	$0.18 \pm 0.11$
$\alpha = 0.05$	$0.04 \pm 0.00$	$0.76 \pm 0.05$	$0.09 \pm 0.00$	$0.78 \pm 0.11$	$0.18 \pm 0.06$
$\alpha = 0.1$	$0.08 \pm 0.01$	$-0.16 \pm 0.32$	$0.11 \pm 0.01$	$0.52 \pm 0.09$	$0.16 \pm 0.06$
<b>ECCCo</b>					
$\alpha = 0.01$	$0.71 \pm 0.28$	$-0.98 \pm 0.03$	$0.32 \pm 0.15$	$0.21 \pm 0.06$	$0.20 \pm 0.08$
$\alpha = 0.05$	$0.69 \pm 0.28$	$-0.97 \pm 0.03$	$0.31 \pm 0.14$	$0.21 \pm 0.06$	$0.20 \pm 0.08$
$\alpha = 0.1$	$0.69 \pm 0.28$	$-0.97 \pm 0.03$	$0.31 \pm 0.14$	$0.21 \pm 0.06$	$0.20 \pm 0.07$
<b>ConfexTree, <math>\alpha = 0.01</math></b>					
$bw = 0.05$	$0.27 \pm 0.09$	$0.14 \pm 0.66$	$0.08 \pm 0.00$	$nan \pm nan$	$0.07 \pm 0.05$
$bw = 0.075$	$0.27 \pm 0.09$	$0.18 \pm 0.83$	$0.08 \pm 0.00$	$nan \pm nan$	$0.07 \pm 0.05$
$bw = 0.1$	$0.19 \pm 0.01$	$0.96 \pm 0.02$	$0.07 \pm 0.00$	$0.09 \pm 0.02$	$0.20 \pm 0.08$
$bw = 0.125$	$0.19 \pm 0.01$	$0.96 \pm 0.01$	$0.07 \pm 0.00$	$0.10 \pm 0.03$	$0.20 \pm 0.07$
$bw = 0.15$	$0.25 \pm 0.02$	$0.96 \pm 0.01$	$0.08 \pm 0.00$	$0.09 \pm 0.03$	$0.18 \pm 0.06$
$bw = 0.175$	$0.22 \pm 0.04$	$0.96 \pm 0.01$	$0.08 \pm 0.00$	$0.11 \pm 0.04$	$0.19 \pm 0.07$
$bw = 0.2$	$0.16 \pm 0.01$	$0.90 \pm 0.08$	$0.08 \pm 0.00$	$0.19 \pm 0.02$	$0.18 \pm 0.07$
$bw = 0.25$	$0.29 \pm 0.02$	$0.88 \pm 0.15$	$0.11 \pm 0.01$	$0.10 \pm 0.02$	$0.17 \pm 0.07$
<b>ConfexTree, <math>\alpha = 0.05</math></b>					
$bw = 0.05$	$0.28 \pm 0.09$	$0.90 \pm 0.15$	$0.06 \pm 0.01$	$0.10 \pm 0.09$	$0.25 \pm 0.07$
$bw = 0.075$	$0.23 \pm 0.05$	$0.97 \pm 0.02$	$0.07 \pm 0.01$	$0.14 \pm 0.04$	$0.22 \pm 0.09$
$bw = 0.1$	$0.08 \pm 0.01$	$0.97 \pm 0.02$	$0.08 \pm 0.00$	$0.47 \pm 0.08$	$0.19 \pm 0.07$
$bw = 0.125$	$0.06 \pm 0.01$	$0.97 \pm 0.02$	$0.08 \pm 0.00$	$0.53 \pm 0.09$	$0.19 \pm 0.07$
$bw = 0.15$	$0.06 \pm 0.01$	$0.95 \pm 0.04$	$0.08 \pm 0.00$	$0.56 \pm 0.07$	$0.19 \pm 0.07$
$bw = 0.175$	$0.05 \pm 0.01$	$0.95 \pm 0.04$	$0.08 \pm 0.00$	$0.59 \pm 0.07$	$0.19 \pm 0.07$
$bw = 0.2$	$0.05 \pm 0.00$	$0.94 \pm 0.04$	$0.09 \pm 0.00$	$0.62 \pm 0.13$	$0.18 \pm 0.07$
$bw = 0.25$	$0.05 \pm 0.00$	$0.91 \pm 0.02$	$0.09 \pm 0.00$	$0.70 \pm 0.31$	$0.18 \pm 0.07$
<b>ConfexTree, <math>\alpha = 0.1</math></b>					
$bw = 0.05$	$0.24 \pm 0.06$	$0.98 \pm 0.01$	$0.08 \pm 0.01$	$0.14 \pm 0.03$	$0.22 \pm 0.09$
$bw = 0.075$	$0.16 \pm 0.04$	$0.98 \pm 0.02$	$0.08 \pm 0.01$	$0.27 \pm 0.07$	$0.20 \pm 0.07$
$bw = 0.1$	$0.06 \pm 0.01$	$0.96 \pm 0.01$	$0.08 \pm 0.00$	$0.59 \pm 0.09$	$0.19 \pm 0.07$
$bw = 0.125$	$0.05 \pm 0.01$	$0.96 \pm 0.01$	$0.08 \pm 0.00$	$0.64 \pm 0.08$	$0.19 \pm 0.07$
$bw = 0.15$	$0.04 \pm 0.00$	$0.96 \pm 0.01$	$0.09 \pm 0.00$	$0.71 \pm 0.11$	$0.18 \pm 0.07$
$bw = 0.175$	$0.04 \pm 0.00$	$0.96 \pm 0.01$	$0.09 \pm 0.00$	$0.73 \pm 0.13$	$0.18 \pm 0.07$
$bw = 0.2$	$0.04 \pm 0.00$	$0.95 \pm 0.02$	$0.09 \pm 0.00$	$0.84 \pm 0.33$	$0.18 \pm 0.07$
$bw = 0.25$	$0.03 \pm 0.00$	$0.95 \pm 0.02$	$0.09 \pm 0.00$	$0.85 \pm 0.33$	$0.18 \pm 0.07$
FACE	$0.12 \pm 0.01$	$0.95 \pm 0.02$	$0.09 \pm 0.00$	$0.35 \pm 0.04$	$0.19 \pm 0.07$
C-CHVAE	$1.32 \pm 0.08$	$-0.92 \pm 0.06$	$0.69 \pm 0.05$	$0.09 \pm 0.02$	$0.27 \pm 0.20$

2212 Table 16: CFX generation results, GiveMeSomeCredit, MLP. Validity 71% for Wachter and 80%  
2213 for Schut.

2214	Generator	Distance	Plausibility	Implausibility	Sensitivity ( $10^{-1}$ )	Stability
<b>RandomForest</b>						
2215	MinDist	$0.01 \pm 0.00$	$0.96 \pm 0.01$	$0.09 \pm 0.00$	$96.79 \pm 35.89$	$0.22 \pm 0.07$
2216	ConfexNaive					
2217	$\alpha = 0.01$	$0.04 \pm 0.01$	$0.88 \pm 0.07$	$0.09 \pm 0.00$	$1.28 \pm 0.10$	$0.22 \pm 0.07$
2218	$\alpha = 0.05$	$0.01 \pm 0.00$	$0.96 \pm 0.01$	$0.09 \pm 0.00$	$67.00 \pm 45.36$	$0.22 \pm 0.07$
2219	$\alpha = 0.1$	$0.01 \pm 0.00$	$0.96 \pm 0.01$	$0.09 \pm 0.00$	$28.68 \pm 25.01$	$0.22 \pm 0.07$
2220	ConfexTree, $\alpha = 0.01$					
2221	$bw = 0.05$	$0.34 \pm 0.07$	$0.30 \pm 0.39$	$0.07 \pm 0.01$	nan $\pm$ nan	$0.29 \pm 0.07$
2222	$bw = 0.1$	$0.08 \pm 0.01$	$0.96 \pm 0.01$	$0.07 \pm 0.00$	$0.50 \pm 0.20$	$0.22 \pm 0.07$
2223	$bw = 0.15$	$0.08 \pm 0.01$	$0.96 \pm 0.03$	$0.07 \pm 0.00$	$0.60 \pm 0.22$	$0.22 \pm 0.07$
2224	$bw = 0.2$	$0.05 \pm 0.01$	$0.96 \pm 0.01$	$0.08 \pm 0.00$	$0.73 \pm 0.14$	$0.22 \pm 0.07$
2225	ConfexTree, $\alpha = 0.05$					
2226	$bw = 0.05$	$0.11 \pm 0.02$	$0.96 \pm 0.02$	$0.07 \pm 0.00$	$0.34 \pm 0.08$	$0.22 \pm 0.06$
2227	$bw = 0.1$	$0.03 \pm 0.01$	$0.98 \pm 0.01$	$0.08 \pm 0.00$	$2.90 \pm 2.12$	$0.22 \pm 0.07$
2228	$bw = 0.15$	$0.03 \pm 0.00$	$0.98 \pm 0.01$	$0.08 \pm 0.00$	$3.35 \pm 2.19$	$0.22 \pm 0.07$
2229	$bw = 0.2$	$0.02 \pm 0.00$	$0.98 \pm 0.01$	$0.08 \pm 0.00$	$2.85 \pm 1.09$	$0.22 \pm 0.07$
2230	ConfexTree, $\alpha = 0.1$					
2231	$bw = 0.05$	$0.07 \pm 0.01$	$0.99 \pm 0.01$	$0.08 \pm 0.00$	$0.55 \pm 0.14$	$0.22 \pm 0.06$
2232	$bw = 0.1$	$0.03 \pm 0.01$	$0.98 \pm 0.02$	$0.08 \pm 0.00$	$2.20 \pm 1.93$	$0.22 \pm 0.07$
2233	$bw = 0.15$	$0.02 \pm 0.00$	$0.98 \pm 0.01$	$0.08 \pm 0.00$	$3.71 \pm 2.69$	$0.22 \pm 0.07$
2234	$bw = 0.2$	$0.02 \pm 0.00$	$0.98 \pm 0.02$	$0.09 \pm 0.00$	$4.14 \pm 2.67$	$0.22 \pm 0.07$
2235	FeatureTweak	$0.03 \pm 0.01$	$0.96 \pm 0.02$	$0.09 \pm 0.00$	$1.40 \pm 0.16$	$0.20 \pm 0.06$
2236	FOCUS	$0.05 \pm 0.00$	$0.91 \pm 0.05$	$0.09 \pm 0.00$	$2.27 \pm 0.79$	$0.22 \pm 0.07$
2237	FACE	$0.10 \pm 0.01$	$0.97 \pm 0.01$	$0.08 \pm 0.00$	$0.46 \pm 0.08$	$0.22 \pm 0.07$

Table 17: CFX generation results, GiveMeSomeCredit, RF. ConfexTree with alpha=0.01,bw=0.05 had 78% failures (i.e. one class had no singleton regions). Valdiity 50% for FeatureTweak.

2238  
2239  
2240  
2241  
2242  
2243  
2244  
2245  
2246  
2247  
2248  
2249  
2250  
2251  
2252  
2253  
2254  
2255  
2256  
2257  
2258  
2259  
2260  
2261  
2262  
2263  
2264  
2265  
2266  
2267

2268 D.5.4 CONFORMAL EVALUATION RESULTS  
2269  
2270  
2271  
2272  
2273  
2274  
2275

2276 Generator	2277 Marginal CovGap	2278 Binning CovGap	2279 Class Cond CovGap	2280 Simulated CovGap
<b>MLP</b>				
<b>ConfexNaive</b>				
$\alpha = 0.01$	$0.99 \pm 0.00$	$-5.24 \pm 0.05$	$0.15 \pm 0.01$	$-10.52 \pm 5.95$
$\alpha = 0.05$	$0.96 \pm 0.00$	$-30.19 \pm 0.86$	$0.02 \pm 0.13$	$-32.31 \pm 9.63$
$\alpha = 0.1$	$0.90 \pm 0.00$	$-41.43 \pm 0.21$	$-0.04 \pm 0.04$	$-10.15 \pm 15.42$
<b>ConfexTree, <math>\alpha = 0.01</math></b>				
$bw = 0.05$	$1.00 \pm 0.00$	$0.98 \pm 0.00$	$1.00 \pm 0.00$	$0.80 \pm 0.00$
$bw = 0.075$	$1.00 \pm 0.00$	$0.98 \pm 0.00$	$1.00 \pm 0.00$	$0.80 \pm 0.00$
$bw = 0.1$	$1.00 \pm 0.00$	$-1.84 \pm 0.09$	$0.61 \pm 0.01$	$-8.96 \pm 0.20$
$bw = 0.125$	$1.00 \pm 0.00$	$-1.73 \pm 0.08$	$0.63 \pm 0.01$	$-8.60 \pm 0.17$
$bw = 0.15$	$1.00 \pm 0.00$	$-2.03 \pm 0.09$	$0.59 \pm 0.01$	$-9.35 \pm 0.18$
$bw = 0.175$	$1.00 \pm 0.00$	$-2.15 \pm 0.09$	$0.57 \pm 0.01$	$-10.04 \pm 0.32$
$bw = 0.2$	$1.00 \pm 0.00$	$-3.88 \pm 0.16$	$0.34 \pm 0.02$	$0.04 \pm 0.03$
$bw = 0.25$	$1.00 \pm 0.00$	$-4.21 \pm 0.09$	$0.29 \pm 0.01$	$0.07 \pm 0.03$
$bw = 0.4$	$0.99 \pm 0.00$	$-5.53 \pm 0.10$	$0.11 \pm 0.01$	$-26.76 \pm 39.37$
<b>ConfexTree, <math>\alpha = 0.05</math></b>				
$bw = 0.05$	$1.00 \pm 0.00$	$-0.48 \pm 0.02$	$4.25 \pm 0.00$	$-14.45 \pm 0.11$
$bw = 0.075$	$1.00 \pm 0.00$	$-0.70 \pm 0.02$	$4.22 \pm 0.00$	$-14.99 \pm 0.11$
$bw = 0.1$	$0.98 \pm 0.00$	$-11.87 \pm 0.08$	$2.70 \pm 0.01$	$-32.61 \pm 22.10$
$bw = 0.125$	$0.98 \pm 0.00$	$-11.94 \pm 0.08$	$2.69 \pm 0.01$	$-32.61 \pm 22.10$
$bw = 0.15$	$0.98 \pm 0.00$	$-14.15 \pm 0.13$	$2.36 \pm 0.02$	$-59.62 \pm 11.20$
$bw = 0.175$	$0.98 \pm 0.00$	$-14.64 \pm 0.14$	$2.30 \pm 0.01$	$-45.84 \pm 9.20$
$bw = 0.2$	$0.97 \pm 0.00$	$-18.83 \pm 0.30$	$1.67 \pm 0.03$	$-38.26 \pm 2.07$
<b>ConfexTree, <math>\alpha = 0.1</math></b>				
$bw = 0.05$	$0.99 \pm 0.00$	$-1.06 \pm 0.05$	$8.48 \pm 0.01$	$-36.74 \pm 40.35$
$bw = 0.075$	$0.99 \pm 0.00$	$-1.37 \pm 0.05$	$8.44 \pm 0.01$	$-36.85 \pm 40.31$
$bw = 0.1$	$0.97 \pm 0.00$	$-14.40 \pm 0.13$	$6.63 \pm 0.01$	$-46.74 \pm 15.17$
$bw = 0.125$	$0.97 \pm 0.00$	$-14.62 \pm 0.13$	$6.60 \pm 0.01$	$-46.78 \pm 15.17$
$bw = 0.15$	$0.96 \pm 0.00$	$-18.02 \pm 0.22$	$6.07 \pm 0.02$	$-34.99 \pm 12.69$
$bw = 0.175$	$0.96 \pm 0.00$	$-18.51 \pm 0.23$	$6.01 \pm 0.02$	$-33.24 \pm 12.32$
$bw = 0.2$	$0.96 \pm 0.00$	$-22.20 \pm 0.37$	$5.42 \pm 0.02$	$-34.36 \pm 2.89$

2306  
2307  
2308  
2309  
2310  
2311  
2312  
2313  
2314  
2315  
2316  
2317  
2318  
2319  
2320  
2321 Table 18: Conformal evaluation results, GiveMeSomeCredit, MLP

	Generator	Marginal CovGap	Binning CovGap	Class Cond CovGap	Simulated CovGap
2322	ConfexNaive				
2323	$\alpha = 0.01$	$0.99 \pm 0.00$	$-6.14 \pm 0.33$	$0.03 \pm 0.04$	$-12.25 \pm 6.80$
2324	$\alpha = 0.05$	$0.95 \pm 0.00$	$-30.33 \pm 0.54$	$-0.12 \pm 0.09$	$-41.99 \pm 4.34$
2325	$\alpha = 0.1$	$0.90 \pm 0.00$	$-40.10 \pm 0.33$	$-0.09 \pm 0.07$	$-39.82 \pm 6.80$
2326	ConfexTree, $\alpha = 0.01$				
2327	$bw = 0.05$	$1.00 \pm 0.00$	$0.98 \pm 0.01$	$1.00 \pm 0.00$	$0.84 \pm 0.09$
2328	$bw = 0.1$	$1.00 \pm 0.00$	$-1.90 \pm 0.09$	$0.57 \pm 0.02$	$-21.32 \pm 26.00$
2329	$bw = 0.15$	$1.00 \pm 0.00$	$-1.93 \pm 0.10$	$0.56 \pm 0.02$	$-29.74 \pm 31.74$
2330	$bw = 0.2$	$0.99 \pm 0.00$	$-3.46 \pm 0.06$	$0.35 \pm 0.02$	$-64.01 \pm 30.10$
2331	ConfexTree, $\alpha = 0.05$				
2332	$bw = 0.05$	$0.99 \pm 0.00$	$-0.74 \pm 0.11$	$3.59 \pm 0.06$	$-72.90 \pm 35.28$
2333	$bw = 0.1$	$0.97 \pm 0.00$	$-12.33 \pm 0.29$	$1.52 \pm 0.09$	$-48.29 \pm 20.97$
2334	$bw = 0.15$	$0.97 \pm 0.00$	$-14.46 \pm 0.30$	$1.20 \pm 0.09$	$-43.87 \pm 12.96$
2335	$bw = 0.2$	$0.96 \pm 0.00$	$-19.58 \pm 0.17$	$0.53 \pm 0.07$	$-46.87 \pm 6.75$
2336	ConfexTree, $\alpha = 0.1$				
2337	$bw = 0.05$	$0.96 \pm 0.00$	$-2.32 \pm 0.17$	$5.78 \pm 0.10$	$-50.85 \pm 46.46$
2338	$bw = 0.1$	$0.93 \pm 0.00$	$-16.78 \pm 0.11$	$2.27 \pm 0.09$	$-50.79 \pm 11.62$
2339	$bw = 0.15$	$0.92 \pm 0.00$	$-20.68 \pm 0.12$	$1.64 \pm 0.10$	$-48.78 \pm 9.57$
2340	$bw = 0.2$	$0.91 \pm 0.00$	$-25.84 \pm 0.14$	$0.96 \pm 0.09$	$-41.02 \pm 18.42$
2341					

Table 19: Conformal evaluation results, GiveMeSomeCredit, RandomForest

2342  
2343  
2344  
2345  
2346  
2347  
2348  
2349  
2350  
2351  
2352  
2353  
2354  
2355  
2356  
2357  
2358  
2359  
2360  
2361  
2362  
2363  
2364  
2365  
2366  
2367  
2368  
2369  
2370  
2371  
2372  
2373  
2374  
2375

2376  
2377

## D.6 ADULTINCOME

2378  
2379  
2380  
2381  
2382  
2383

This dataset Becker & Kohavi (1996), obtained through Kaggle<sup>8</sup>, predicts whether an individual's income exceeds \$50,000. We processed the following features: numeric features (Age, Capital Gain, Capital Loss, Hours per week) scaled to (0, 1), ordinal features (education), and categorical features (Workclass, Occupation, Race, Relationship, Gender, Marital status) using one-hot encoding. In CONFEXTree, to avoid splitting over many categorical feature combinations, we consider the first (Workclass) as a feature to split by and do not split the rest.

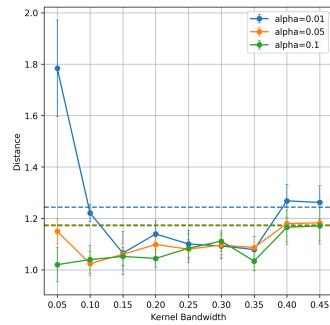
2384  
2385  
2386  
2387

For this dataset we find that CONFEX-Tree methods generally obtain worse plausibility than competing methods (although we have comparable sensitivity). This could be attributed to an insufficient kernel: further tuning to obtain a better kernel (bandwidth, features contributing to the kernel) etc. to better define locality could help with this.

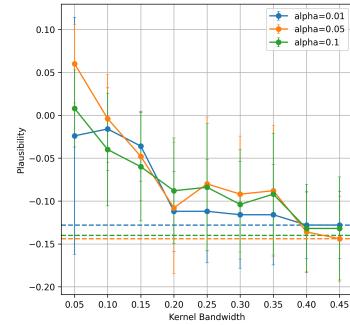
2388  
2389

## D.6.1 PLOTS

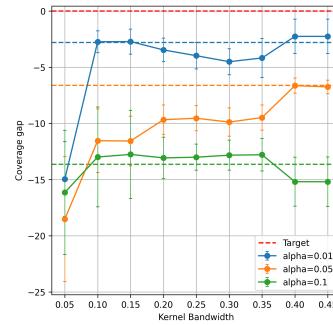
2390



(a) Distance



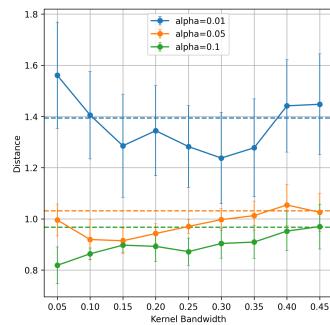
(b) Plausibility



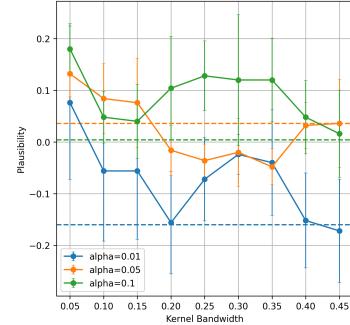
(c) Coverage Gap

2401  
2402

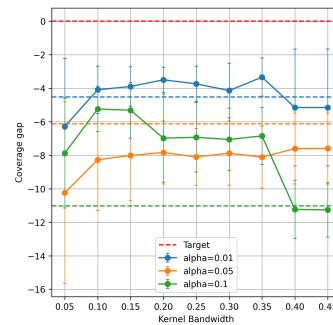
Figure 15: Effect of coverage rate and kernel bandwidth on metrics for CONFEX-Tree on the Adult-Income dataset, MLP. CONFEX-Naive is represented by dashed horizontal lines.

2403  
24042405  
2406

(a) Distance



(b) Plausibility



(c) Coverage Gap

2418  
2419

Figure 16: Effect of coverage rate and kernel bandwidth on metrics for CONFEX-Tree on the Adult-Income dataset, RandomForest. CONFEX-Naive is represented by dashed horizontal lines.

2420  
2421  
2422  
2423  
2424  
2425  
2426  
2427  
2428  
2429

<sup>8</sup><https://www.kaggle.com/datasets/wenruliu/adult-income-dataset>

2430  
2431 D.6.2 MODEL EVALUATION RESULTS

2432 2433	Repeat	Accuracy (%)	Precision (%)	F1 Score (%)
2434	repeat0,MLP	85.41	85.05	85.17
2435	repeat1,MLP	85.04	84.70	84.83
2436	repeat2,MLP	84.89	84.32	84.40
2437	repeat3,MLP	85.03	84.58	84.71
2438	repeat4,MLP	84.96	84.35	84.29
2439	repeat0,RF	85.73	85.20	85.14
2440	repeat1,RF	85.32	84.76	84.72
2441	repeat2,RF	85.70	85.18	85.06
2442	repeat3,RF	85.51	84.96	84.87
2443	repeat4,RF	85.48	84.93	84.89

2444  
2445 Table 20: Model evaluation results, AdultIncome.

2446  
2447  
2448  
2449  
2450  
2451  
2452  
2453  
2454  
2455  
2456  
2457  
2458  
2459  
2460  
2461  
2462  
2463  
2464  
2465  
2466  
2467  
2468  
2469  
2470  
2471  
2472  
2473  
2474  
2475  
2476  
2477  
2478  
2479  
2480  
2481  
2482  
2483

2484 D.6.3 CFX GENERATION RESULTS  
2485  
2486  
2487  
2488  
2489  
2490

2491 Generator	2492 Distance	2493 Plausibility	2494 Implausibility	2495 Sensitivity ( $10^{-1}$ )	2496 Stability
<b>MLP</b>					
2493 MinDist	2494 $1.16 \pm 0.07$	2495 $-0.13 \pm 0.05$	2496 $1.97 \pm 0.04$	2497 $0.11 \pm 0.11$	2498 $0.33 \pm 0.02$
2493 Wachter	2494 $0.39 \pm 0.05$	2495 $0.36 \pm 0.02$	2496 $1.87 \pm 0.01$	2497 $0.14 \pm 0.01$	2498 $0.14 \pm 0.02$
2493 Greedy	2494 $0.95 \pm 0.08$	2495 $0.02 \pm 0.10$	2496 $2.06 \pm 0.04$	2497 $24190.61 \pm 48381.15$	2498 $0.86 \pm 0.03$
<b>ConfexNaive</b>					
2497 $\alpha = 0.01$	2498 $1.24 \pm 0.07$	2499 $-0.13 \pm 0.04$	2500 $1.97 \pm 0.04$	2501 $0.05 \pm 0.01$	2502 $0.44 \pm 0.04$
2497 $\alpha = 0.05$	2498 $1.17 \pm 0.07$	2499 $-0.14 \pm 0.05$	2500 $1.97 \pm 0.04$	2501 $0.05 \pm 0.01$	2502 $0.37 \pm 0.02$
2497 $\alpha = 0.1$	2498 $1.17 \pm 0.06$	2499 $-0.14 \pm 0.04$	2500 $1.97 \pm 0.04$	2501 $0.08 \pm 0.05$	2502 $0.35 \pm 0.02$
<b>ECCCo</b>					
2501 $\alpha = 0.01$	2502 $0.57 \pm 0.01$	2503 $0.13 \pm 0.01$	2504 $1.88 \pm 0.01$	2505 $0.05 \pm 0.00$	2506 $0.37 \pm 0.01$
2501 $\alpha = 0.01$	2502 $0.73 \pm 0.11$	2503 $-0.05 \pm 0.07$	2504 $1.87 \pm 0.06$	2505 $0.05 \pm 0.00$	2506 $0.63 \pm 0.14$
2501 $\alpha = 0.05$	2502 $0.57 \pm 0.00$	2503 $0.12 \pm 0.02$	2504 $1.88 \pm 0.01$	2505 $0.05 \pm 0.00$	2506 $0.37 \pm 0.02$
2501 $\alpha = 0.05$	2502 $0.72 \pm 0.11$	2503 $-0.04 \pm 0.06$	2504 $1.87 \pm 0.06$	2505 $0.05 \pm 0.01$	2506 $0.61 \pm 0.15$
2501 $\alpha = 0.1$	2502 $0.56 \pm 0.01$	2503 $0.12 \pm 0.02$	2504 $1.88 \pm 0.01$	2505 $0.05 \pm 0.00$	2506 $0.37 \pm 0.02$
2501 $\alpha = 0.1$	2502 $0.72 \pm 0.11$	2503 $-0.04 \pm 0.06$	2504 $1.87 \pm 0.06$	2505 $0.05 \pm 0.01$	2506 $0.61 \pm 0.15$
<b>ConfexTree, <math>\alpha = 0.01</math></b>					
2507 $bw = 0.05$	2508 $1.78 \pm 0.19$	2509 $-0.02 \pm 0.14$	2510 $1.80 \pm 0.04$	2511 $0.05 \pm 0.02$	2512 $0.29 \pm 0.02$
2507 $bw = 0.1$	2508 $1.22 \pm 0.03$	2509 $-0.02 \pm 0.05$	2510 $1.84 \pm 0.04$	2511 $0.06 \pm 0.01$	2512 $0.34 \pm 0.02$
2507 $bw = 0.15$	2508 $1.06 \pm 0.08$	2509 $-0.04 \pm 0.04$	2510 $1.86 \pm 0.05$	2511 $0.05 \pm 0.01$	2512 $0.36 \pm 0.04$
2507 $bw = 0.2$	2508 $1.14 \pm 0.05$	2509 $-0.11 \pm 0.05$	2510 $1.93 \pm 0.05$	2511 $0.05 \pm 0.01$	2512 $0.36 \pm 0.03$
2507 $bw = 0.25$	2508 $1.10 \pm 0.05$	2509 $-0.11 \pm 0.06$	2510 $1.93 \pm 0.04$	2511 $0.05 \pm 0.01$	2512 $0.36 \pm 0.03$
2507 $bw = 0.3$	2508 $1.09 \pm 0.05$	2509 $-0.12 \pm 0.06$	2510 $1.94 \pm 0.04$	2511 $0.05 \pm 0.01$	2512 $0.35 \pm 0.03$
2507 $bw = 0.35$	2508 $1.08 \pm 0.05$	2509 $-0.12 \pm 0.06$	2510 $1.93 \pm 0.04$	2511 $0.05 \pm 0.01$	2512 $0.37 \pm 0.03$
2507 $bw = 0.4$	2508 $1.27 \pm 0.06$	2509 $-0.13 \pm 0.04$	2510 $1.97 \pm 0.04$	2511 $0.05 \pm 0.01$	2512 $0.46 \pm 0.04$
2507 $bw = 0.45$	2508 $1.26 \pm 0.07$	2509 $-0.13 \pm 0.04$	2510 $1.97 \pm 0.04$	2511 $0.04 \pm 0.01$	2512 $0.46 \pm 0.04$
<b>ConfexTree, <math>\alpha = 0.05</math></b>					
2516 $bw = 0.05$	2517 $1.15 \pm 0.05$	2518 $0.06 \pm 0.05$	2519 $1.87 \pm 0.03$	2520 $0.08 \pm 0.03$	2521 $0.28 \pm 0.02$
2516 $bw = 0.1$	2517 $1.02 \pm 0.05$	2518 $-0.00 \pm 0.05$	2519 $1.91 \pm 0.05$	2520 $0.07 \pm 0.02$	2521 $0.30 \pm 0.02$
2516 $bw = 0.15$	2517 $1.06 \pm 0.04$	2518 $-0.05 \pm 0.05$	2519 $1.93 \pm 0.04$	2520 $0.07 \pm 0.01$	2521 $0.33 \pm 0.03$
2516 $bw = 0.2$	2517 $1.10 \pm 0.03$	2518 $-0.11 \pm 0.08$	2519 $1.95 \pm 0.04$	2520 $0.06 \pm 0.02$	2521 $0.35 \pm 0.02$
2516 $bw = 0.25$	2517 $1.08 \pm 0.04$	2518 $-0.08 \pm 0.08$	2519 $1.95 \pm 0.03$	2520 $0.07 \pm 0.02$	2521 $0.34 \pm 0.02$
2516 $bw = 0.3$	2517 $1.10 \pm 0.04$	2518 $-0.09 \pm 0.07$	2519 $1.95 \pm 0.03$	2520 $0.07 \pm 0.02$	2521 $0.35 \pm 0.02$
2516 $bw = 0.35$	2517 $1.09 \pm 0.04$	2518 $-0.09 \pm 0.08$	2519 $1.94 \pm 0.03$	2520 $0.06 \pm 0.03$	2521 $0.35 \pm 0.02$
2516 $bw = 0.4$	2517 $1.18 \pm 0.07$	2518 $-0.14 \pm 0.05$	2519 $1.97 \pm 0.04$	2520 $0.05 \pm 0.01$	2521 $0.37 \pm 0.02$
2516 $bw = 0.45$	2517 $1.18 \pm 0.07$	2518 $-0.14 \pm 0.05$	2519 $1.97 \pm 0.04$	2520 $0.05 \pm 0.01$	2521 $0.37 \pm 0.02$
<b>ConfexTree, <math>\alpha = 0.1</math></b>					
2525 $bw = 0.05$	2526 $1.02 \pm 0.07$	2527 $0.01 \pm 0.04$	2528 $1.91 \pm 0.04$	2529 $0.08 \pm 0.01$	2530 $0.27 \pm 0.02$
2525 $bw = 0.1$	2526 $1.04 \pm 0.06$	2527 $-0.04 \pm 0.07$	2528 $1.93 \pm 0.04$	2529 $0.08 \pm 0.02$	2530 $0.31 \pm 0.02$
2525 $bw = 0.15$	2526 $1.05 \pm 0.04$	2527 $-0.06 \pm 0.06$	2528 $1.93 \pm 0.04$	2529 $0.08 \pm 0.02$	2530 $0.32 \pm 0.02$
2525 $bw = 0.2$	2526 $1.04 \pm 0.04$	2527 $-0.09 \pm 0.06$	2528 $1.92 \pm 0.03$	2529 $0.07 \pm 0.02$	2530 $0.33 \pm 0.02$
2525 $bw = 0.25$	2526 $1.08 \pm 0.06$	2527 $-0.08 \pm 0.07$	2528 $1.96 \pm 0.03$	2529 $0.07 \pm 0.02$	2530 $0.34 \pm 0.02$
2525 $bw = 0.3$	2526 $1.11 \pm 0.04$	2527 $-0.10 \pm 0.06$	2528 $1.96 \pm 0.03$	2529 $0.07 \pm 0.03$	2530 $0.34 \pm 0.02$
2525 $bw = 0.35$	2526 $1.03 \pm 0.04$	2527 $-0.09 \pm 0.07$	2528 $1.93 \pm 0.04$	2529 $0.07 \pm 0.03$	2530 $0.34 \pm 0.02$
2525 $bw = 0.4$	2526 $1.17 \pm 0.07$	2527 $-0.13 \pm 0.05$	2528 $1.97 \pm 0.04$	2529 $0.08 \pm 0.07$	2530 $0.35 \pm 0.02$
2525 $bw = 0.45$	2526 $1.17 \pm 0.07$	2527 $-0.13 \pm 0.06$	2528 $1.97 \pm 0.04$	2529 $0.08 \pm 0.07$	2530 $0.35 \pm 0.02$
<b>FACE</b>					
2534 $FACE$	2535 $1.36 \pm 0.16$	2536 $0.34 \pm 0.12$	2537 $1.77 \pm 0.03$	2538 $0.06 \pm 0.02$	2539 $0.32 \pm 0.01$
<b>C-CHVAE</b>					
2534 $C-CHVAE$	2535 $6.39 \pm 0.55$	2536 $0.33 \pm 0.15$	2537 $1.18 \pm 0.18$	2538 $0.04 \pm 0.01$	2539 $0.46 \pm 0.08$

2536 Table 21: CFX generation results, AdultIncome, MLP. Validity 80% for Wachter, 84-85% for all  
2537 ECCCo methods, 89% for Greedy, 54% for C-CHVAE.

2538	Generator	Distance	Plausibility	Implausibility	Sensitivity ( $10^{-1}$ )	Stability
<b>RandomForest</b>						
ConfexNaive						
2542	$\alpha = 0.01$	$1.39 \pm 0.19$	$-0.16 \pm 0.10$	$1.89 \pm 0.07$	$0.14 \pm 0.11$	$0.29 \pm 0.05$
2543	$\alpha = 0.05$	$1.03 \pm 0.08$	$0.04 \pm 0.11$	$1.91 \pm 0.05$	$0.12 \pm 0.05$	$0.26 \pm 0.04$
2544	$\alpha = 0.1$	$0.97 \pm 0.09$	$0.00 \pm 0.10$	$1.91 \pm 0.04$	$0.12 \pm 0.02$	$0.26 \pm 0.03$
ConfexTree, $\alpha = 0.01$						
2545	$bw = 0.05$	$1.56 \pm 0.21$	$0.08 \pm 0.15$	$1.80 \pm 0.04$	$0.06 \pm 0.03$	$0.18 \pm 0.02$
2546	$bw = 0.1$	$1.41 \pm 0.17$	$-0.06 \pm 0.14$	$1.79 \pm 0.05$	$0.11 \pm 0.03$	$0.24 \pm 0.04$
2547	$bw = 0.15$	$1.29 \pm 0.20$	$-0.06 \pm 0.13$	$1.82 \pm 0.06$	$0.09 \pm 0.03$	$0.24 \pm 0.04$
2548	$bw = 0.2$	$1.34 \pm 0.18$	$-0.16 \pm 0.10$	$1.89 \pm 0.07$	$0.12 \pm 0.03$	$0.25 \pm 0.04$
2549	$bw = 0.25$	$1.28 \pm 0.16$	$-0.07 \pm 0.08$	$1.91 \pm 0.07$	$0.12 \pm 0.05$	$0.26 \pm 0.04$
2550	$bw = 0.3$	$1.24 \pm 0.18$	$-0.02 \pm 0.04$	$1.93 \pm 0.06$	$0.10 \pm 0.05$	$0.24 \pm 0.03$
2551	$bw = 0.35$	$1.28 \pm 0.19$	$-0.04 \pm 0.10$	$1.93 \pm 0.06$	$0.14 \pm 0.10$	$0.25 \pm 0.04$
2552	$bw = 0.4$	$1.44 \pm 0.18$	$-0.15 \pm 0.09$	$1.90 \pm 0.07$	$0.08 \pm 0.03$	$0.29 \pm 0.06$
2553	$bw = 0.45$	$1.45 \pm 0.20$	$-0.17 \pm 0.10$	$1.90 \pm 0.07$	$0.09 \pm 0.03$	$0.29 \pm 0.06$
ConfexTree, $\alpha = 0.05$						
2554	$bw = 0.05$	$1.00 \pm 0.06$	$0.13 \pm 0.04$	$1.86 \pm 0.05$	$0.12 \pm 0.04$	$0.19 \pm 0.02$
2555	$bw = 0.1$	$0.92 \pm 0.08$	$0.08 \pm 0.07$	$1.89 \pm 0.05$	$0.12 \pm 0.04$	$0.22 \pm 0.02$
2556	$bw = 0.15$	$0.91 \pm 0.05$	$0.08 \pm 0.09$	$1.90 \pm 0.04$	$0.15 \pm 0.05$	$0.22 \pm 0.02$
2557	$bw = 0.2$	$0.94 \pm 0.04$	$-0.02 \pm 0.05$	$1.93 \pm 0.04$	$0.14 \pm 0.04$	$0.24 \pm 0.03$
2558	$bw = 0.25$	$0.97 \pm 0.03$	$-0.04 \pm 0.03$	$1.94 \pm 0.03$	$0.09 \pm 0.01$	$0.24 \pm 0.03$
2559	$bw = 0.3$	$1.00 \pm 0.04$	$-0.02 \pm 0.07$	$1.95 \pm 0.03$	$0.11 \pm 0.04$	$0.24 \pm 0.02$
2560	$bw = 0.35$	$1.01 \pm 0.06$	$-0.05 \pm 0.03$	$1.95 \pm 0.03$	$0.11 \pm 0.03$	$0.24 \pm 0.02$
2561	$bw = 0.4$	$1.05 \pm 0.08$	$0.03 \pm 0.04$	$1.91 \pm 0.04$	$0.13 \pm 0.04$	$0.26 \pm 0.04$
2562	$bw = 0.45$	$1.03 \pm 0.07$	$0.04 \pm 0.09$	$1.91 \pm 0.04$	$0.15 \pm 0.06$	$0.26 \pm 0.04$
ConfexTree, $\alpha = 0.1$						
2563	$bw = 0.05$	$0.82 \pm 0.07$	$0.18 \pm 0.05$	$1.88 \pm 0.04$	$0.20 \pm 0.06$	$0.20 \pm 0.03$
2564	$bw = 0.1$	$0.86 \pm 0.02$	$0.05 \pm 0.05$	$1.90 \pm 0.03$	$0.18 \pm 0.10$	$0.22 \pm 0.02$
2565	$bw = 0.15$	$0.90 \pm 0.03$	$0.04 \pm 0.07$	$1.92 \pm 0.04$	$0.19 \pm 0.13$	$0.24 \pm 0.02$
2566	$bw = 0.2$	$0.89 \pm 0.06$	$0.10 \pm 0.10$	$1.88 \pm 0.04$	$0.15 \pm 0.08$	$0.25 \pm 0.02$
2567	$bw = 0.25$	$0.87 \pm 0.05$	$0.13 \pm 0.07$	$1.87 \pm 0.03$	$0.15 \pm 0.05$	$0.25 \pm 0.03$
2568	$bw = 0.3$	$0.90 \pm 0.06$	$0.12 \pm 0.13$	$1.89 \pm 0.03$	$0.17 \pm 0.05$	$0.25 \pm 0.02$
2569	$bw = 0.35$	$0.91 \pm 0.06$	$0.12 \pm 0.08$	$1.88 \pm 0.03$	$0.18 \pm 0.10$	$0.25 \pm 0.02$
2570	$bw = 0.4$	$0.95 \pm 0.08$	$0.05 \pm 0.07$	$1.91 \pm 0.04$	$1.92 \pm 3.55$	$0.25 \pm 0.02$
2571	$bw = 0.45$	$0.97 \pm 0.09$	$0.02 \pm 0.08$	$1.92 \pm 0.04$	$2.62 \pm 4.87$	$0.25 \pm 0.03$
2572	FeatureTweak	$0.24 \pm 0.10$	$0.30 \pm 0.11$	$1.84 \pm 0.04$	$0.06 \pm 0.01$	$0.13 \pm 0.01$
2573	FOCUS	$0.58 \pm 0.34$	$0.40 \pm 0.06$	$1.84 \pm 0.04$	$0.29 \pm 0.14$	$0.17 \pm 0.01$
2574	FACE	$1.50 \pm 0.07$	$0.39 \pm 0.07$	$1.74 \pm 0.03$	$0.06 \pm 0.02$	$0.27 \pm 0.01$

Table 22: CFX generation results, AdultIncome, RandomForest. Methods with nan values had 100% failures. Validity 61% for FeatureTweak.

2575  
2576  
2577  
2578  
2579  
2580  
2581  
2582  
2583  
2584  
2585  
2586  
2587  
2588  
2589  
2590  
2591

2592 D.6.4 CONFORMAL EVALUATION RESULTS  
2593  
2594  
2595  
2596  
2597  
2598

2599 Generator	2600 Marginal CovGap	2601 Binning CovGap	2602 Class Cond CovGap	2603 Simulated CovGap
<b>MLP</b>				
ConfexNaive				
$\alpha = 0.01$	$0.99 \pm 0.00$	$-0.88 \pm 0.08$	$-0.01 \pm 0.05$	$-2.81 \pm 0.88$
$\alpha = 0.05$	$0.96 \pm 0.00$	$-2.64 \pm 0.19$	$0.33 \pm 0.09$	$-6.63 \pm 0.94$
$\alpha = 0.1$	$0.90 \pm 0.00$	$-4.74 \pm 0.78$	$0.21 \pm 0.08$	$-13.65 \pm 2.19$
ConfexTree, $\alpha = 0.01$				
$bw = 0.05$	$1.00 \pm 0.00$	$0.10 \pm 0.09$	$0.52 \pm 0.06$	$-14.95 \pm 3.34$
$bw = 0.1$	$0.99 \pm 0.00$	$-0.36 \pm 0.16$	$0.28 \pm 0.09$	$-2.74 \pm 0.97$
$bw = 0.15$	$0.99 \pm 0.00$	$-0.32 \pm 0.17$	$0.30 \pm 0.09$	$-2.72 \pm 1.11$
$bw = 0.2$	$0.99 \pm 0.00$	$-1.03 \pm 0.10$	$-0.08 \pm 0.06$	$-3.46 \pm 1.05$
$bw = 0.25$	$0.99 \pm 0.00$	$-0.97 \pm 0.17$	$-0.06 \pm 0.09$	$-3.98 \pm 1.18$
$bw = 0.3$	$0.99 \pm 0.00$	$-1.07 \pm 0.15$	$-0.13 \pm 0.08$	$-4.51 \pm 1.16$
$bw = 0.35$	$0.99 \pm 0.00$	$-0.93 \pm 0.14$	$-0.04 \pm 0.08$	$-4.18 \pm 1.73$
$bw = 0.4$	$0.99 \pm 0.00$	$-0.76 \pm 0.20$	$0.07 \pm 0.11$	$-2.26 \pm 1.53$
$bw = 0.45$	$0.99 \pm 0.00$	$-0.76 \pm 0.20$	$0.07 \pm 0.11$	$-2.26 \pm 1.53$
ConfexTree, $\alpha = 0.05$				
$bw = 0.05$	$0.98 \pm 0.00$	$0.61 \pm 0.18$	$2.48 \pm 0.05$	$-18.52 \pm 5.56$
$bw = 0.1$	$0.96 \pm 0.00$	$-1.50 \pm 0.16$	$1.11 \pm 0.12$	$-11.54 \pm 2.82$
$bw = 0.15$	$0.96 \pm 0.00$	$-1.92 \pm 0.14$	$0.86 \pm 0.14$	$-11.57 \pm 2.21$
$bw = 0.2$	$0.95 \pm 0.00$	$-2.55 \pm 0.37$	$0.47 \pm 0.14$	$-9.68 \pm 1.30$
$bw = 0.25$	$0.95 \pm 0.00$	$-2.61 \pm 0.31$	$0.35 \pm 0.10$	$-9.54 \pm 1.12$
$bw = 0.3$	$0.95 \pm 0.00$	$-2.71 \pm 0.30$	$0.28 \pm 0.08$	$-9.89 \pm 1.27$
$bw = 0.35$	$0.95 \pm 0.00$	$-2.71 \pm 0.37$	$0.30 \pm 0.14$	$-9.48 \pm 1.13$
$bw = 0.4$	$0.95 \pm 0.01$	$-2.88 \pm 0.27$	$0.14 \pm 0.19$	$-6.64 \pm 0.66$
$bw = 0.45$	$0.95 \pm 0.01$	$-2.88 \pm 0.27$	$0.14 \pm 0.19$	$-6.76 \pm 0.60$
ConfexTree, $\alpha = 0.1$				
$bw = 0.05$	$0.94 \pm 0.00$	$1.33 \pm 0.38$	$4.61 \pm 0.06$	$-16.15 \pm 5.52$
$bw = 0.1$	$0.93 \pm 0.00$	$-0.44 \pm 0.47$	$3.46 \pm 0.05$	$-12.99 \pm 4.44$
$bw = 0.15$	$0.93 \pm 0.00$	$-0.91 \pm 0.53$	$3.16 \pm 0.11$	$-12.76 \pm 3.92$
$bw = 0.2$	$0.92 \pm 0.00$	$-2.51 \pm 0.51$	$1.93 \pm 0.17$	$-13.07 \pm 1.83$
$bw = 0.25$	$0.92 \pm 0.00$	$-2.57 \pm 0.53$	$1.78 \pm 0.13$	$-13.01 \pm 1.16$
$bw = 0.3$	$0.92 \pm 0.00$	$-2.68 \pm 0.55$	$1.72 \pm 0.12$	$-12.83 \pm 1.34$
$bw = 0.35$	$0.92 \pm 0.00$	$-2.81 \pm 0.54$	$1.64 \pm 0.12$	$-12.79 \pm 1.44$
$bw = 0.4$	$0.90 \pm 0.00$	$-5.01 \pm 0.90$	$0.13 \pm 0.17$	$-15.19 \pm 2.17$
$bw = 0.45$	$0.90 \pm 0.00$	$-5.01 \pm 0.90$	$0.13 \pm 0.17$	$-15.19 \pm 2.20$

2633  
2634  
2635  
2636  
2637  
2638  
2639  
2640  
2641  
2642  
2643  
2644  
2645

Table 23: Conformal evaluation results, AdultIncome, MLP

2646	Generator	Marginal CovGap	Binning CovGap	Class Cond CovGap	Simulated CovGap
<b>RandomForest</b>					
<b>ConfexNaive</b>					
2650	$\alpha = 0.01$	$0.99 \pm 0.00$	$-0.90 \pm 0.22$	$0.01 \pm 0.09$	$-5.70 \pm 5.11$
2651	$\alpha = 0.05$	$0.96 \pm 0.00$	$-2.81 \pm 0.08$	$0.35 \pm 0.12$	$-5.49 \pm 2.52$
2652	$\alpha = 0.1$	$0.91 \pm 0.00$	$-5.78 \pm 0.58$	$0.13 \pm 0.14$	$-10.19 \pm 0.48$
<b>ConfexTree, <math>\alpha = 0.01</math></b>					
2653	$bw = 0.05$	$1.00 \pm 0.00$	$0.20 \pm 0.02$	$0.56 \pm 0.01$	$-5.11 \pm 2.96$
2654	$bw = 0.1$	$0.99 \pm 0.00$	$-0.43 \pm 0.08$	$0.23 \pm 0.03$	$-2.95 \pm 0.14$
2655	$bw = 0.15$	$0.99 \pm 0.00$	$-0.45 \pm 0.12$	$0.23 \pm 0.05$	$-2.67 \pm 0.17$
2656	$bw = 0.2$	$0.99 \pm 0.00$	$-0.80 \pm 0.12$	$0.05 \pm 0.04$	$-3.32 \pm 0.27$
2657	$bw = 0.25$	$0.99 \pm 0.00$	$-0.78 \pm 0.13$	$0.04 \pm 0.03$	$-3.89 \pm 0.55$
2658	$bw = 0.3$	$0.99 \pm 0.00$	$-0.87 \pm 0.19$	$-0.00 \pm 0.07$	$-4.68 \pm 0.38$
2659	$bw = 0.35$	$0.99 \pm 0.00$	$-0.79 \pm 0.20$	$0.04 \pm 0.08$	$-3.76 \pm 0.17$
2660	$bw = 0.4$	$0.99 \pm 0.00$	$-0.91 \pm 0.24$	$-0.00 \pm 0.10$	$-7.54 \pm 4.96$
2661	$bw = 0.45$	$0.99 \pm 0.00$	$-0.91 \pm 0.24$	$-0.00 \pm 0.10$	$-7.54 \pm 4.96$
<b>ConfexTree, <math>\alpha = 0.05</math></b>					
2662	$bw = 0.05$	$0.97 \pm 0.01$	$0.03 \pm 0.06$	$1.84 \pm 0.08$	$-6.43 \pm 1.84$
2663	$bw = 0.1$	$0.96 \pm 0.00$	$-1.60 \pm 0.23$	$0.88 \pm 0.11$	$-5.54 \pm 1.04$
2664	$bw = 0.15$	$0.95 \pm 0.00$	$-2.03 \pm 0.04$	$0.58 \pm 0.05$	$-5.39 \pm 0.67$
2665	$bw = 0.2$	$0.95 \pm 0.00$	$-2.59 \pm 0.01$	$0.44 \pm 0.02$	$-6.10 \pm 0.05$
2666	$bw = 0.25$	$0.96 \pm 0.00$	$-2.66 \pm 0.19$	$0.38 \pm 0.10$	$-6.74 \pm 0.05$
2667	$bw = 0.3$	$0.96 \pm 0.00$	$-2.59 \pm 0.12$	$0.43 \pm 0.05$	$-6.17 \pm 0.05$
2668	$bw = 0.35$	$0.96 \pm 0.00$	$-2.52 \pm 0.08$	$0.47 \pm 0.02$	$-6.79 \pm 0.05$
2669	$bw = 0.4$	$0.95 \pm 0.00$	$-2.82 \pm 0.10$	$0.31 \pm 0.18$	$-6.25 \pm 0.14$
2670	$bw = 0.45$	$0.95 \pm 0.00$	$-2.82 \pm 0.10$	$0.31 \pm 0.18$	$-6.24 \pm 0.12$
<b>ConfexTree, <math>\alpha = 0.1</math></b>					
2671	$bw = 0.05$	$0.93 \pm 0.00$	$-0.11 \pm 0.80$	$2.91 \pm 0.35$	$-4.98 \pm 0.44$
2672	$bw = 0.1$	$0.91 \pm 0.00$	$-2.48 \pm 0.49$	$1.38 \pm 0.18$	$-4.62 \pm 1.45$
2673	$bw = 0.15$	$0.91 \pm 0.01$	$-3.14 \pm 0.83$	$1.00 \pm 0.34$	$-4.33 \pm 1.62$
2674	$bw = 0.2$	$0.91 \pm 0.00$	$-4.33 \pm 0.10$	$0.66 \pm 0.05$	$-4.80 \pm 0.47$
2675	$bw = 0.25$	$0.91 \pm 0.00$	$-4.50 \pm 0.07$	$0.54 \pm 0.16$	$-5.34 \pm 0.81$
2676	$bw = 0.3$	$0.91 \pm 0.00$	$-4.67 \pm 0.15$	$0.38 \pm 0.20$	$-5.73 \pm 1.22$
2677	$bw = 0.35$	$0.91 \pm 0.00$	$-4.88 \pm 0.36$	$0.25 \pm 0.33$	$-5.56 \pm 1.45$
2678	$bw = 0.4$	$0.91 \pm 0.00$	$-6.14 \pm 0.35$	$0.04 \pm 0.06$	$-10.77 \pm 3.43$
2679	$bw = 0.45$	$0.91 \pm 0.00$	$-6.14 \pm 0.34$	$0.03 \pm 0.05$	$-10.87 \pm 3.20$

Table 24: Conformal evaluation results, AdultIncome, RandomForest