
Data-SUITE: Data-centric identification of in-distribution incongruous examples

Nabeel Seedat¹ Jonathan Crabbé¹ Mihaela van der Schaar^{1 2 3}

Abstract

Systematic quantification of data quality is critical for consistent model performance. Prior works have focused on out-of-distribution data. Instead, we tackle an understudied yet equally important problem of characterizing incongruous regions of in-distribution (ID) data, which may arise from feature space heterogeneity. To this end, we propose a paradigm shift with Data-SUITE: a data-centric AI framework to identify these regions, independent of a task-specific model. Data-SUITE leverages copula modeling, representation learning, and conformal prediction to build feature-wise confidence interval estimators based on a set of training instances. These estimators can be used to evaluate the congruence of test instances with respect to the training set, to answer two practically useful questions: (1) which test instances will be reliably predicted by a model trained with the training instances? and (2) can we identify incongruous regions of the feature space so that data owners understand the data’s limitations or guide future data collection? We empirically validate Data-SUITE’s performance and coverage guarantees and demonstrate on cross-site medical data, biased data, and data with concept drift, that Data-SUITE best identifies ID regions where a downstream model may be reliable (independent of said model). We also illustrate how these identified regions can provide insights into datasets and highlight their limitations.

1. Introduction

Machine learning models have a well-known reliance on training data quality (Park et al., 2021). Hence, when deploying such models in the real world, the reliability of

predictions depends on the data’s congruence with respect to the training data. Significant literature has focused on identifying data instances that lie out of the training data’s distribution (OOD). This includes label shifts (Ren et al., 2018; Hsu et al., 2020) or input feature shift, where these instances fall out of the support of the training set’s distribution (Zhang et al., 2021). However, a much less studied, yet equally important problem is identifying heterogeneous regions of in-distribution (ID) data.

Data in the wild can be ID yet have heterogeneous regions in feature space. This manifests in varying levels of *incongruence*, in cases of different sub-populations, data biases or temporal changes (Leslie et al., 2021; Gianfrancesco et al., 2018; Obermeyer et al., 2019). We illustrate each of these types of incongruence with real world data (Table 2), in the experiments from Secs. 4.3 and 4.5.

In this paper, we present a *data-centric* framework to characterize such *incongruous* regions of *ID data* and define two groups, namely (i) *inconsistent* and (ii) *uncertain*, with respect to the training distribution. We contextualize the difference based on confidence intervals (CI) (See Sec.3.3 for details). When feature values lie outside of a CI, we term it inconsistent, alternatively we characterize the level of feature uncertainty based on the CI’s width.

At this point, one might ask if the data is ID; why should we worry? Not accounting for these incongruous ID regions of the feature space can be problematic when deploying models in high-stakes settings such as healthcare, where spurious model predictions can be deadly (Saria & Subbaswamy, 2019; Varshney, 2020). That said, even in settings where poor predictions are not risky, consistent exploratory data analysis (EDA) and retroactive auditing of such data is time-consuming for data scientists (Polyzotis et al., 2017; Kandel et al., 2012). Hence, systematically identifying these incongruous regions has immense practical value.

Consequently, we build a framework to empower data scientists to address the previously mentioned challenges related to insightful exploratory data analysis (EDA) and reliable model deployment, anchored by the following desiderata:

(D1) Insightful Data Exploration: Alice has a new dataset \mathcal{D} and wants to explore and gain insights into it with respect

¹Department of Applied Mathematics and Theoretical Physics, University of Cambridge, UK ²The Alan Turing Institute, London, UK ³University of California, Los Angeles, USA. Correspondence to: Nabeel Seedat <ns741@cam.ac.uk>.

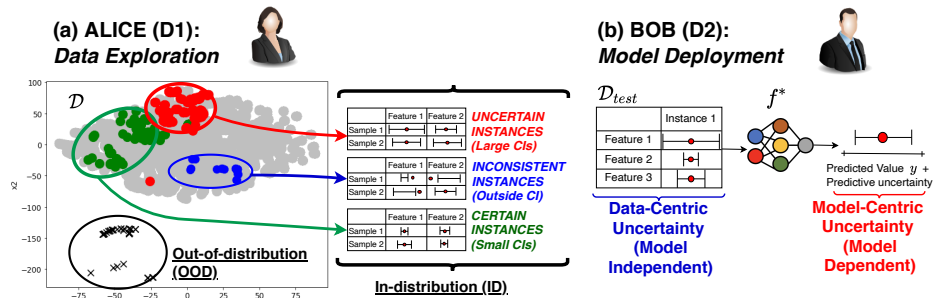


Figure 1. Illustration highlighting two problems Data-SUITE addresses

to a training set, without necessarily training a model. It would be useful if, *independent of a predictive model*, she could both identify the incongruous regions of the feature space (e.g., sub-population bias or under-representation), as well as, obtain easily digestible prototype examples of each region. This could guide where to collect more data and, if this is not possible, to understand the data’s limitations.

(D2) Reliable Model Deployment: Bob has a trained model f^* and now deploys it to another site. For new data $\mathcal{D}_{\text{test}}$, it would be useful if he could identify incongruous regions, for which he should NOT trust f^* to make predictions.

(D3) Practitioner confidence: Both Alice and Bob want to feel confident when using any tool. Guarantees of coverage of predictive intervals (e.g. CIs) could assist in this regard.

These examples, shown in Fig. 1 highlight the need to understand incongruence in data. As we shall discuss in the related work, there has been significant work on uncertainty estimation, with a focus on the uncertainty of a model’s predictions (*model-centric*). Estimating predictive uncertainty can address Bob’s use-case (D2), however since it requires a predictive model it is not naturally suited to Alice’s insights use-case (D1). Further, most predictive uncertainty methods do not provide coverage guarantees (D3).

Therefore, in satisfying all the desiderata, we take a different approach and advocate for a *data-centric* approach, where we model the uncertainty in the data features¹. This is different from *model-centric* predictive uncertainty, as we construct CIs (at feature level), without reference to any downstream model. A benefit of the flexibility is that we can flag instances and draw insights that are not model-specific (i.e. *model independent*).

To ensure clarity, we note that the term *data-centric* is used in the context of *data-centric AI*, which we define as “tools applied to the underlying data used to train and evaluate models”. We note that both paradigms of ML (model and data-centric) rely on models, differing based on how mod-

¹Here, “feature” uncertainty refers to the degree of incongruity with the training distribution, rather than the uncertainty of the measured value (e.g. measurement noise)

els are used. In “model-centric ML”: models are used for predictive tasks and in “data-centric ML”: models are used to study/evaluate the data itself, which does not preclude using algorithms to process the data. Our definition is consistent with (Polyzotis et al., 2017) “data-centric AI as the problem of (...) [A]nd quality monitoring processes for datasets”. This paper fits this as a systematic tool for data quality evaluation based on uncertainty.

In this work, we focus on tabular data, a common format in medicine, finance, manufacturing etc, where data is based on relational databases (Borisov et al., 2021; Yoon et al., 2020). That said, compared to image data, tabular data has an added challenge since specific features may be uncertain while others are not; hence characterizing an instance as a whole is non-trivial.

Contributions. We present **Data Searching for Uncertain and Inconsistent Test Examples (Data-SUITE)**, a *data-centric* framework to identify incongruous regions of data using CIs and make the following contributions:

- Data-SUITE is a paradigm shift from model-centric uncertainty and, to the best of our knowledge, the first to characterize ID regions in a systematic *data-centric, model-independent* manner. Not only is this more flexible, but also enables us to gain insights which are not model-specific.
- Data-SUITE’s pipeline-based approach to construct feature-wise CIs enables specific properties (Sec. 3.2) that permit us to flag uncertain and inconsistent instances, making it possible to identify incongruous data regions.
- Data-SUITE’s performance and properties, such as coverage guarantees, are validated to satisfy D3 (Sec. 4.1).
- Further motivating the paradigm shift, we empirically highlight the performance benefit of a data-centric approach compared to a model-centric approach (Sec. 4.2).
- As a portrayal of reliable model deployment (D2), we show on real-world datasets with different types of incongruence, that Data-SUITE best identifies incongruous data regions, translating to the best performance improvement. (Sec. 4.3).
- Finally, we illustrate with multiple use-cases how Data-SUITE can be used as a model-independent tool to facilitate insightful data exploration, hence satisfying D1 (Sec. 4.5).

2. Related work

This paper primarily engages with the literature on uncertainty quantification and contributes to the nascent area of data-centric AI. We also highlight the key differences of our work with the literature on noisy labels.

Uncertainty quantification. There are numerous Bayesian and non-Bayesian methods for uncertainty quantification, including Gaussian processes (Williams & Rasmussen, 2006), Quantile Regression (Koenker & Hallock, 2001), Bayesian Neural Networks (Ghosh et al., 2018; Graves, 2011), Deep Ensembles (Lakshminarayanan et al., 2017), Dropout (Gal & Ghahramani, 2016; Chan et al., 2020) and Conformal Prediction (Vovk et al., 2005). These methods typically assess predictive uncertainty, i.e., measuring the certainty in the model’s prediction (Seedat & Kanan, 2019). The predominant focus on predictive uncertainty is different from the notion of uncertainty in our setting, which is feature (i.e. data) uncertainty. We specifically highlight that we quantify data uncertainty, independent of a task-specific model. Additionally, the aforementioned methods often do not assess the coverage or provide guarantees of the uncertainty interval (Wasserman, 2004; Alaa & Van Der Schaar, 2020) (i.e., how often the interval contains the true value). The concept of coverage will be outlined further in Secs. 3 and 4.

Data-Centric AI. Ensuring high data quality is a critical but often overlooked problem in ML, where the focus is optimizing models (Sambasivan et al., 2021; Jain et al., 2020). Even when it is considered, the process of assessing datasets is adhoc or artisanal (Sambasivan et al., 2021; Ng et al., 2021). However, there has been recent discussion around data-centric AI (DCAI), which we define as tools applied to the underlying data used to train and evaluate models, independent of the task-specific, predictive models. Our work contributes to this nascent body of work – presenting Data-SUITE, which, to the best of our knowledge, is the first systematic data-centric framework to model uncertainty in datasets. Specifically, we model the uncertainty in the feature (data) values themselves (*data-centric*), which contrasts to modeling the uncertainty in predictions (*model-centric*). We also highlight a tangential of classical data management (Kumar et al., 2017), which does not consider uncertainty and randomness in the data for subsequent analysis. An exception is probabilistic databases (Suciu et al., 2011).

Noisy labels. Learning with noisy data is a widely studied problem, we refer the reader to (Algan & Ulusoy, 2021; Song et al., 2020) for an in depth review. In machine learning, the focus is label noise. We argue that work on noisy labels is not directly related, as the goal is to learn a model robust to the label noise, which is different from our goal of modeling the uncertainty in the features. Additionally, methods are often coupled to the task-specific predictive model, which is different from our model-independent setting.

3. Data-SUITE

In this section, we give a detailed formulation of Data-SUITE^{2,3}. We start with a problem formulation and outline the motivation for working with feature confidence intervals (CIs). Then, we describe how these CIs are built by leveraging copula modelling, representation learning and conformal prediction. Finally, we demonstrate how these CIs permit to flag uncertain and inconsistent instances.

3.1. Preliminaries

We consider a feature space $\mathcal{X} = \prod_{i=1}^{d_X} [a_i, b_i] \subseteq \mathbb{R}^{d_X}$, where $[a_i, b_i]$ is the range for feature i . Note that we make the range of each feature explicit, this will be necessary in the definition of our formalism. We assume that we have a set of $M \in \mathbb{N}^*$ training instances $\mathcal{D}_{\text{train}} = \{x^m \mid m \in [M]\}$ sampled from an unknown distribution \mathbb{P} , where $[M]$ denotes the positive integers between 1 and M . These instances typically correspond to training data for a model on a downstream task, such as classification.

We assume that we are given new test instances $\mathcal{D}_{\text{test}}$. Our purpose is to flag the subset of instances from $\mathcal{D}_{\text{test}}$ that are quantitatively different from instances of $\mathcal{D}_{\text{train}}$ without necessarily being OOD. To that aim, we use $\mathcal{D}_{\text{train}}$ to build CIs $[l_i(x), r_i(x)] \subseteq [a_i, b_i]$ for each feature $i \in [d_X]$ of each test instance $x \in \mathcal{D}_{\text{test}}$. As we will show in Sec. 3.3, these CIs permit to systematically flag test instances whose features are uncertain or inconsistent with respect to $\mathcal{D}_{\text{train}}$. For now, let us motivate the usage of feature CIs: (1) With a model of uncertainty and inconsistency at the feature level, it is possible to identify regions of the feature space \mathcal{X} where bias and/or low coverage occurs with the training data $\mathcal{D}_{\text{train}}$. (2) Since CIs are built with $\mathcal{D}_{\text{train}}$ and without reference to any predictive downstream model, the flagged instances in $\mathcal{D}_{\text{test}}$ are likely to be problematic for any downstream model trained on top of $\mathcal{D}_{\text{train}}$. Hence, we are able to draw conclusions that are not model-specific. These two points are illustrated in our experiments from Sec. 4. Let us now detail how the CIs are built.

3.2. Feature CIs

We now build CIs $[l_i(x), r_i(x)] \subseteq [a_i, b_i]$ for each feature $i \in [d_X]$ of each test instance $x \in \mathcal{D}_{\text{test}}$. It goes without saying that the CIs should satisfy some properties, i.e. (P1) *Coverage*: We would like to guarantee that the feature x_i of an instance $x \sim \mathbb{P}$ lies within the interval such that $\mathbb{E} [1_{x_i \in [l_i(x), r_i(x)]}] \geq 1 - \alpha$ where the significance level $\alpha \in (0, 1)$ can be chosen. In this way, a feature out of the CI hints that x is unlikely to be sampled from \mathbb{P} at the given significance level. This is then considered across all features to characterize the instance (see Sec. 3.3).

²<https://github.com/seedatnabeel/Data-SUITE>

³<https://github.com/vanderschaarlab/mlforhealthlabpub/tree/main/alg/Data-SUITE>

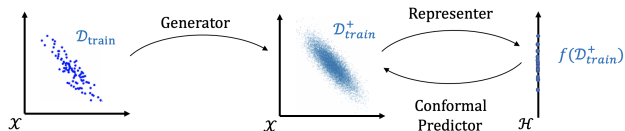


Figure 2. Outline of our framework **Data-SUITE**.

(P2) *Instance-wise*: The CI should be *adaptive at an instance level*. i.e, we do not wish $r_i(x) - l_i(x)$ to be constant w.r.t $x \in \mathcal{X}$. In this way, the CIs permit to order various test instances $x \in \mathcal{D}_{\text{test}}$ according to their uncertainty. This property is particularly desirable in healthcare settings where we wish to quantify variable uncertainty for individual patients, rather than for the population as whole.

(P3) *Feature-wise*: We build CIs $[l_i(x), r_i(x)]$ for each feature $i \in [d_X]$ as opposed to an overall confidence region $R(x) \subset \mathcal{X}$. While less general than the latter approach, feature-wise CIs are more interpretable, allowing attribution of inconsistencies and uncertainty to individual features.

(P4) *Downstream coupling*: Instances with smaller CIs are more reliably predicted by a downstream model trained on $\mathcal{D}_{\text{train}}$. More precisely, our CIs should have a negative correlation between CI width and downstream model performance. In this way, CIs allow to draw conclusions about the incongruence of test instances $x \in \mathcal{D}_{\text{test}}$.

To construct feature CIs that satisfy these properties, we introduce a new framework leveraging copula modeling, representation learning and conformal prediction. The blueprint of our method is presented in Fig. 2. Concretely, our method relies on 3 building blocks: a *generator* that augments the initial training set $\mathcal{D}_{\text{train}}$; a *representer* that leverages the augmented training set $\mathcal{D}_{\text{train}}^+$ to learn a low-dimensional representation $f: \mathcal{X} \rightarrow \mathcal{H}$ of the data and a *conformal predictor* that predicts instance-wise feature CIs $[l_i(x), r_i(x)]$ on the basis of each instance’s representation $f(x) \in \mathcal{H}$. By construction of the CIs, this method fulfills properties (P2) and (P3). As we will see in the following, the conformal predictor’s theoretical properties guarantees (P1). However, we also empirically validate (P1), see Sec 4.1. We then also demonstrate (P4) empirically in Sec. 4. Appendix C.1 quantifies the significance of each block via an ablation study. Let us now detail each block.

Generator. The purpose of the generator is to augment the initial training set $\mathcal{D}_{\text{train}}$ with instances that are consistent with the initial distribution \mathbb{P} . Many data augmentation techniques can be used for this block. Since our focus is on tabular data, we found *copula modeling* to be particularly useful. Copulas leverage Sklar’s theorem (Sklar, 1959) to estimate multivariate distributions with univariate marginal

distributions. In our case, we use vine copulas (Bedford & Cooke, 2001) to build an estimate $\hat{\mathbb{P}}$ for the distribution \mathbb{P} on the basis of $\mathcal{D}_{\text{train}}$. We then build an augmented training set $\mathcal{D}_{\text{train}}^+$ by sampling from the copula density $\hat{\mathbb{P}}$. Interestingly, our method does not need to access $\mathcal{D}_{\text{train}}$ once the copula density $\hat{\mathbb{P}}$ is available. It is perfectly possible to use only instances from $\hat{\mathbb{P}}$ to build the augmented dataset $\mathcal{D}_{\text{train}}^+$. This could be useful for data sharing, if the access to the training set $\mathcal{D}_{\text{train}}$ is restricted to the user. Further details and motivations on copulas is found in Appendix A.2.1. Note that a copula might not be ideal for very high-dimensional (large d_X) data in domains such as computer vision or genomics. In those cases, copula modeling can be replaced by domain-specific augmentation techniques.

Representer. A trivial way to verify the coverage guarantee (P1) would be to use the true values of the features to build the CIs: $[l_i(x), r_i(x)] = [x_i - \delta, x_i + \delta]$ for some $\delta \in \mathbb{R}^+$. The problem with this approach is two-fold: (1) it does not leverage the distribution \mathbb{P} underlying the training set $\mathcal{D}_{\text{train}}$ and (2) it results in an uninformative reconstruction with CIs that does not capture the specificity of each instance, hence contradicting (P2). To provide a more satisfactory solution, we propose to represent the augmented training data $\mathcal{D}_{\text{train}}^+$ with a representation function $f: \mathcal{X} \rightarrow \mathcal{H}$ that maps the data into a lower-dimensional latent representation space $\mathcal{H} \subseteq \mathbb{R}^{d_H}$, $d_H < d_X$. The purpose of this representer is to capture the structure of the low-dimensional manifold underlying $\mathcal{D}_{\text{train}}^+$. At test time, the conformal predictor (detailed next), uses the lower representations $f(x) \in \mathcal{H}$ to estimate a reconstruction interval for each feature x_i . This permits to bring a satisfactory solution to the two aforementioned problems: (1) the CIs are reconstructed in terms of latent factors that are useful to describe the training set $\mathcal{D}_{\text{train}}$ and (2) the predicted CIs vary according to the representation $f(x) \in \mathcal{H}$ of each test instance $x \in \mathcal{D}_{\text{test}}$. In essence, our approach is analogous to autoencoders. As we will explain soon, the crucial difference is the decoding step: our method outputs CIs for the reconstructed input. In this work, we use Principal Component Analysis (PCA), the workhorse for tabular data, to learn the representer f . Note that more general encoder architectures can be used in settings such as computer vision.

Conformal Predictor. We now turn to the core of the problem: estimating feature-wise CIs. As previously mentioned, the CIs $[l_i(x), r_i(x)]$, $i \in [d_X]$ will be *computed on the basis of the latent representation $f(x)$* for each $x \in \mathcal{D}_{\text{test}}$. The idea is simple: for each feature $i \in [d_X]$, we train a regressor $g_i: \mathcal{H} \rightarrow [a_i, b_i]$ to reconstruct an estimate of the initial features x_i from the latent representation $f(x)$ of the associated training instance: $(g_i \circ f)(x) \approx x_i$. We stress that the regressor g_i has no knowledge of the true observed x_i but only of the latent representation $f(x)$, as illustrated

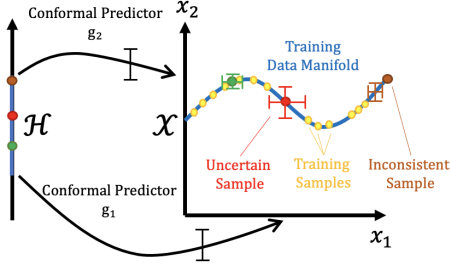


Figure 3. Conformal Predictor in **Data-SUITE**.

in Fig. 3. Of course, the feature regressor by themselves provides point-wise estimates for the features. In order to turn these into CIs, we use conformal prediction as a wrapper around the feature regressor (Vovk et al., 2005).

We formalize our problem in the framework of Inductive Conformal Prediction (for motivations, see Appendix A.2.2). Hence, under the formulation, we start by splitting the augmented training set into a proper training set and a calibration set: $\mathcal{D}_{\text{train}}^+ = \mathcal{D}_{\text{train2}}^+ \sqcup \mathcal{D}_{\text{cal}}^+$. We use the latent representation of the proper training set, to train the feature regressor $g_i, i \in [d_X]$ for the reconstruction task. Then, the latent representation of the calibration set is used to compute the non-conformity score (μ), which estimates how different a new instance looks from other instances.

In practice, we use the absolute error non-conformity score $\mu_i(x) = |x_i - (g_i \circ f)(x)|$. We obtain an empirical distribution of non-conformity scores $\{\mu_i(x) \mid x \in \mathcal{D}_{\text{cal}}^+\}$ over the calibration instances for each feature $i \in [d_X]$. This is used to obtain the critical non-conformity score ϵ , which corresponds to the $\lceil (|\mathcal{D}_{\text{cal}}^+| + 1)(1 - \alpha) \rceil$ -th smallest residual from the set $\{\mu_i(x) \mid x \in \mathcal{D}_{\text{cal}}^+\}$ (Vovk, 2013). We then apply the method to any unseen incoming data to obtain predictive CIs for the data point, i.e. $[l_i(x), r_i(x)] = [(g_i \circ f)(x) - \epsilon, (g_i \circ f)(x) + \epsilon]$. However, in this form the CIs are constant for all instances, where the width of the interval is determined by the residuals of the most difficult instances (largest residuals).

We adapt our conformal prediction framework to obtain the desired adaptive intervals (P2) using a normalized non-conformity function (γ), see Eq. 1 (Boström et al., 2016; Johansson et al., 2015). The numerator is computed as before based on μ , however, the denominator normalizes per instance. We learn the normalizer per feature $i \in [d_X]$.

To do so, we compute the log residuals per feature, for all instances in the respective proper training set $\mathcal{D}_{\text{train2}}$. We produce tuples per feature: $\{(f(x), \ln|x - (g_i \circ f)(x)|) \mid x \in \mathcal{D}_{\text{train2}}\}$. These are used to train a *different* model, $\sigma_i : \mathcal{X} \rightarrow \mathbb{R}^+$ (e.g., MLP), to predict the log residuals. We can then apply σ_i to *test instances* to capture the difficulty in predicting said instance. Note, we apply an exponential to the predicted log residual for the test instance converting

to the true scale and ensuring positive estimates.

$$\gamma_i(x) \equiv \frac{|x_i - (g_i \circ f)(x)|}{\sigma_i(x)}, \quad (1)$$

We can then obtain the critical non-conformity score ϵ applied to the empirical distribution of normalized non-conformity scores $\{\gamma_i(x) \mid x \in \mathcal{D}_{\text{cal}}^+\}$, in the same way as before, based on residuals.

The instance-specific adaptive intervals are then obtained as per Eq. 2, where g is the underlying feature regressor and σ_i is the instance-wise normalizing function.

$$[l_i(x), r_i(x)] = [(g_i \circ f)(x) - \epsilon\sigma_i(x), (g_i \circ f)(x) + \epsilon\sigma_i(x)] \quad (2)$$

Remarks on theoretical guarantees. Under the exchangeability assumption detailed in the Appendix A.2.2, the validity of coverage guarantees (P2) is fulfilled with our definition. In our implementation, we use $\alpha = .05$.

3.3. Identifying Inconsistent and Uncertain Instances

Now that we have CIs $[l_i(x), r_i(x)] \subset [a_i, b_i]$ for each feature $x_i, i \in [d_X]$ of the instance x , we can evaluate if instances from a dataset falls within the predicted range. If it falls outside the predicted range, we characterize the inconsistency (see Definition 3.1)

Definition 3.1 (Inconsistency). Let $x \in \mathcal{D}_{\text{test}}$ be a test instance for which we construct a $(1 - \alpha)$ -CI, $[l_i(x), r_i(x)]$, for each feature $x_i, i \in [d_X]$ for some predetermined $\alpha \in (0, 1)$. For each $x_i, i \in [d_X]$, the *feature inconsistency* is a binary variable indicating if x_i falls out of the CI.

$$\nu_i(x) \equiv 1(x_i \notin [l_i(x), r_i(x)]) \quad (3)$$

The *instance inconsistency* $\nu(x)$ is obtained by averaging over the feature inconsistencies $\nu_i(x)$.

$$\nu(x) \equiv \frac{1}{d_X} \sum_{i=1}^{d_X} \nu_i(x)$$

The instance x is *inconsistent* if the fraction of inconsistent features is above a predetermined threshold⁴ $\lambda \in [0, 1]$: $\nu(x) > \lambda$.

There can also be degrees of uncertainty in the feature value for features that fall within the CI, which can reflect the instance as a whole. Indeed, if the CI $[l_i(x), r_i(x)]$ is large, the feature x_i is likely to fall within its range. Nonetheless, we should keep in mind that large CIs correspond to a large uncertainty for the related feature. This will also typically

⁴In our implementation, we use $\lambda = 0.5$.

happen when the instance $x \in \mathcal{D}_{\text{test}}$ differs from the training set $\mathcal{D}_{\text{train}}$ used to build the CI. We now introduce a quantitative measure that expresses the degree of uncertainty of the instance with respect to $\mathcal{D}_{\text{train}}$ (see Definition 3.2).

Definition 3.2 (Uncertainty). Let $x \in \mathcal{D}_{\text{test}}$ be a test instance for which we construct a $(1 - \alpha)$ CI $[l_i(x), r_i(x)]$ for each feature $x_i, i \in [d_X]$ for a predetermined $\alpha \in (0, 1)$. For each $x_i, i \in [d_X]$, we define *feature uncertainty* $\Delta_i(x)$ as the feature CI width normalized by the feature range:

$$\Delta_i(x) \equiv \frac{r_i(x) - l_i(x)}{b_i - a_i} \quad (4)$$

The *instance uncertainty* $\Delta(x)$ is obtained by averaging over all feature uncertainties:

$$\Delta(x) \equiv \frac{1}{d_X} \sum_{i=1}^{d_X} \Delta_i(x) \in (0, 1].$$

Remark. Instance uncertainties are strictly larger than zero, as feature uncertainties are computed over all features. Hence, this characterization offers a natural split between *certain* and *uncertain* instances if we sort the instances based on uncertainty.

4. Experiments

This section presents a detailed empirical evaluation demonstrating that Data-SUITE satisfies (D1) Insightful Data Exploration, (D2) Reliable Model Deployment and (D3) Practitioner confidence, introduced in Sec. 1. We tackle these in reverse order, as practitioner confidence is a prerequisite for the adoption of D1 and D2. Recall that the notion of uncertainty in Data-SUITE is different from predictive uncertainty (model-centric). We empirically compare these two paradigms using methods for predictive uncertainty. That said, a natural additional question is whether model-centric uncertainty estimation methods can simply be applied in this setting and provide uncertainty estimates for feature values. We benchmark the following widely used Bayesian and non-Bayesian methods (under BOTH model-centric & data-centric paradigms): Bayesian Neural Networks (BNN) (Ghosh et al., 2018), Deep Ensembles (ENS) (Lakshminarayanan et al., 2017), Gaussian Processes (GP) (Williams & Rasmussen, 2006), Monte-Carlo Dropout (MCD) (Gal & Ghahramani, 2016) and Quantile Regression (QR) (Koenker & Hallock, 2001). We also ablate and test Data-SUITE’s constituent components independently: conditional sampling from copula (COP), Conformal Prediction on raw data (CONF) (Vovk et al., 2005; Balasubramanian et al., 2014). For implementation details, see Appendix B.1.

4.1. Validating coverage & comparing properties

We firstly wish to validate the CIs to ensure that the coverage guarantees are satisfied such that users can have confidence

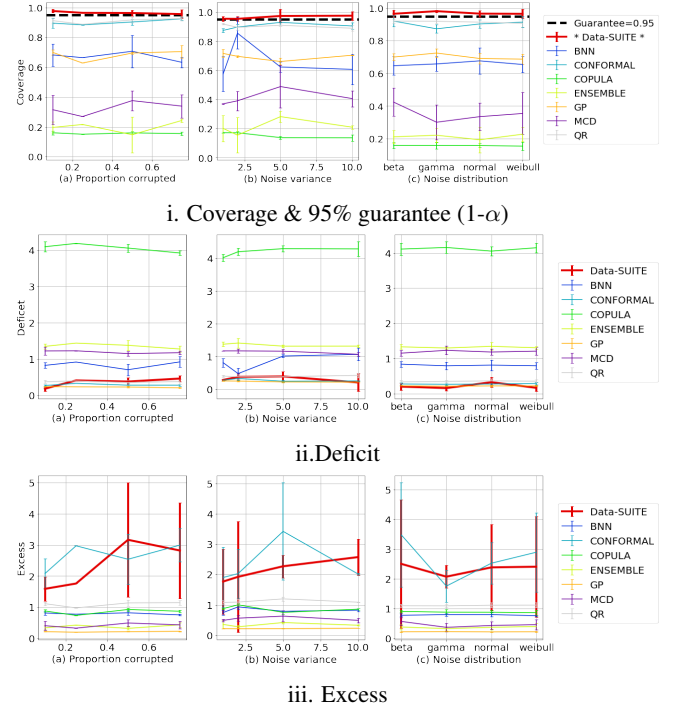


Figure 4. Comparison of methods based on coverage, deficit and excess under various configurations (D_a, D_b, D_c)

that the true value lies within the predicted CIs (D3). We assess the CIs based on the following metrics defined in (Navratil et al., 2020) - (1) Coverage: how often the CI contains the true value, (2) Deficit: extent of CI shortfall (i.e., the severity of the errors) and (3) Excess: extent of CI excess width to capture the true value.

$$\text{Coverage} = \mathbb{E} [1_{x_i \in [l_i, r_i]}] \quad (5)$$

$$\text{Deficit} = \mathbb{E} [1_{x_i \notin [l_i, r_i]} \cdot \min \{|x_i - l_i|, |x_i - r_i|\}] \quad (6)$$

$$\text{Excess} = \mathbb{E} [1_{x_i \in [l_i, r_i]} \cdot \min \{x_i - l_i, r_i - x_i\}] \quad (7)$$

Synthetic data. We assess the properties of different methods using synthetic data as the ground truth values are available, even when encoding incongruence. The synthetic data with features, $\mathbf{X} = [X_1, X_2, X_3]$, is drawn IID from a multivariate Gaussian distribution, parameterized by mean vector μ and a positive definite covariance matrix Σ (details in Appendix B.2). We sample $n = 1000$ points for both $\mathcal{D}_{\text{train}}^{\text{synth}}$ and $\mathcal{D}_{\text{test}}^{\text{synth}}$ and encode incongruence into $\mathcal{D}_{\text{test}}^{\text{synth}}$ using a multivariate additive model $\tilde{\mathbf{X}} = \mathbf{X} + \mathbf{Z}$, where $\mathbf{Z} \in \mathbb{R}^{n \times m}$, is the perturbation matrix. We conduct experiments with different configurations: (1) D_a : Multivariate Gaussian with variance 2 and varying proportion of perturbed instances. (2) D_b : Multivariate Gaussian with varying variance and fixed proportion of perturbed instances (50%) and (3) D_c : Varying distribution $\in \{\text{Beta}, \text{Gamma}, \text{Normal}, \text{Weibull}\}$.

Fig. 4 outlines mean coverage, deficit and excess averaged over five runs for $\mathcal{D}_{\text{test}}^{\text{synth}}$ under the different configurations

(D_a, D_b, D_c). There is a clear variability amongst the different methods, suggesting specific methods are more suitable. Data-SUITE outperforms the other methods based on coverage and deficit across all configurations. We note the methods with poor coverage are typically “incorrectly” confident, i.e. small intervals with low coverage and high deficit. Fig. 4 also demonstrates a meaningful relationship that coverage and deficit are inversely related (high coverage is associated with the low deficit), as with deficit and excess. Although high coverage and low deficit ideally occur with low excess, we observe that high levels of coverage occur in conjunction with high levels of excess. Critically, however, in satisfying **D3**, Data-SUITE maintains the 95% coverage guarantees across all configurations, unlike other methods.

4.2. Synthetic data stratification w/ downstream task

While it is essential to validate a method’s properties, the most useful goal is whether the intervals can be used to identify instances that will be reliably predicted by a downstream predictive model. With Data-SUITE, we stress that this is done in a model-independent manner (i.e. no knowledge of the downstream model). We train a downstream regression model using $\mathcal{D}_{\text{train}}^{\text{synth}}$, where features X_1, X_2 are used to predict X_3 . We first compute a baseline mean squared error (MSE) on a held-out validation set of $\mathcal{D}_{\text{train}}^{\text{synth}}$ and the complete test set $\mathcal{D}_{\text{test}}^{\text{synth}}$ ($\hat{\mathbf{X}} = \mathbf{X} + \mathbf{Z}$). Thereafter, we construct predictive intervals for $\mathcal{D}_{\text{test}}^{\text{synth}}$ using all benchmark methods (either uncertainty intervals or CIs). The intervals are then used to sort instances based on width.

In addition, we answer the question of whether a *data-centric* or *model-centric* approach yields the best performance. For data-centric paradigm, we construct intervals for the features X_1, X_2 , hence instances are categorized in a model-independent manner based on data-level CIs. In contrast, the model-centric paradigm is tightly coupled with a task-specific model, categorizing instances using predictive uncertainty based on prediction X_3 .

We then compute MSE for the 100 most certain instances as ranked by each method (smallest widths). For Data-SUITE, we also compute the MSE for those instances identified as *inconsistent* (outside CIs). The best method is one in which the *certain* sorted instances produce MSE values closest to the clean train MSE (baseline) i.e. has the lowest MSE.

Table 1 shows the MSE for configurations D_a and D_b . As one example of satisfying **D2**, Data-SUITE has the best performance and identifies the top 100 *certain* instances that yields the best downstream model performance, with the lowest MSE across all configurations. In addition, as expected the *inconsistent* instances are unreliably predicted. The poor performance for ablations of Data-SUITE components, suggests the necessity of the inter-connected framework (more in Appendix C.1)

Table 1. MSE based on instance stratification for different methods. Data-SUITE outperforms other methods, whilst data-centric methods in general outperform model-centric methods

		Proportion (D_a)			Variance (D_b)	
PERTURBATION		.1	.25	.5	1	2
<i>Train Data (BASELINE)</i>		.067	.059	.068	.065	.068
Test Data		.222	.513	.889	.275	.889
Data-centric	Data-SUITE (All, Uncertainty)	.069	.122	.197	.104	.197
	Data-SUITE (All, Inconsistent)	.595	1.608	2.322	.791	2.322
	Data-SUITE (CONF)	.125	.396	.846	.293	.846
	Data-SUITE (COP)	.220	.277	.451	.236	.451
	BNN	.192	.216	.704	.173	.704
	ENS	.125	.311	.565	.204	.565
	GP	.112	.153	.296	.158	.296
	MCD	.173	.391	.692	.201	.692
	QR	.116	.228	.635	.193	.635
Model-centric	BNN (Predictive)	.208	.220	.692	.195	.692
	ENS (Predictive)	.143	.226	.625	.257	.625
	GP (Predictive)	.147	.472	.584	.237	.584
	MCD (Predictive)	.206	.255	.684	.213	.684
	QR (Predictive)	.187	.477	.671	.223	.671

Additionally, we see for the same base methods (e.g. BNN, MCD etc), that the data-centric paradigm outperforms the model-centric paradigm in identifying the “best” instances to give the lowest MSE. This result highlights the performance advantage of a flexible, model-independent data-centric paradigm compared to the model-centric paradigm.

4.3. Real dataset stratification w/ downstream task

We now demonstrate how Data-SUITE can be practically used on real data to stratify instances for improved downstream performance (satisfying **D2**). Specifically, to assist with more reliable and performant model deployment across a variety of scenarios. To this end, we select three real-world datasets with different types of incongruence as presented in Table 2. For details see Appendix B.3.

Evaluation. We stratify $\mathcal{D}_{\text{test}}$ into *certain* and *uncertain* instances based on the interval width predicted by each method. e.g. the most uncertain has the largest width.

Are the identified instances OOD? At this point, one might be tempted to assume that the identified instances are simply OOD. i.e. samples that fall out of the support of the training set’s distribution. We show in reality that this is unlikely the case. We apply existing algorithms to detect OOD and outliers (for others see (Yang et al., 2021)): Mahalanobis distance (Lee et al., 2018), SUOD (Zhao et al., 2021), CO-POD (Li et al., 2020) and Isolation Forest (Liu et al., 2012). For each of the detection methods, we compute the overlap between the predicted OOD/Outlier instances and the uncertain and inconsistent instances as identified by Data-SUITE. We found minimal overlap across methods, ranging between 1-18%. Additionally, the OOD detection methods were often unconfident in their predictions, with average confidence scores ranging between 5-50%. Both results suggest the identified uncertain and inconsistent instances are unlikely OOD. For more see Appendix C.2.

Each stratification method will identify different instances

Table 2. Comparison of real-world datasets

Dataset	Incongruence Type	Downstream Task	Stratification
Seer (US) & Cutract (UK)	Geographic UK-US (Cross-site medical)	Predict mortality from prostate cancer	$\mathcal{D}_{train}^{Seer}$: Seer (US), $\mathcal{D}_{test}^{Cutract}$: Cutract (UK)
Adult	Demographic Bias (Gender & Income)	Predict income over \$50K	$\mathcal{D}_{train}^{Adult}$, $\mathcal{D}_{test}^{Adult}$ balanced split
Electricity	Temporal (Consumption patterns)	Predict electricity price rise/fall	$\mathcal{D}_{train}^{Elec}$: 1996, $\mathcal{D}_{test}^{Elec}$: 1997-1998

for each group, hence we aim to quantify which method identifies instances that provide the most improvement to downstream performance. We do this by computing the accuracy of the *certain* and *uncertain* stratifications, on a downstream random forest trained on \mathcal{D}_{train} . Ideally, correct instance stratification results in greater accuracy for *certain* compared to *uncertain* instances. As an overall comparative metric, we compute the *Mean Performance Improvement - MPI* (Eq. 8). MPI is the difference in accuracy (*Acc*) between *certain* and *uncertain* instances, as identified by a specific method, averaged over different threshold proportions P . The best performing method would clearly identify the most appropriate *certain* and *uncertain* instances, which would translate to the largest MPI.

$$MPI = \frac{1}{|P|} \sum_{p \in P} Acc(Cert_p) - Acc(Uncert_p) \quad (8)$$

where $P = \{0.05k \mid k \in [20]\}$, $Cert_p =$ Set of p most *certain* instances, $Uncert_p =$ Set of p most *uncertain* instances.

Fig. 5 illustrates an example of Data-SUITE, applied to the CUTRACT dataset. The metric *MPI* (Eq. 8) is the mean difference between *certain* (green) and *uncertain* (red) curves. The results demonstrate the performance improvement when evaluating with the stratified *certain* and *uncertain* instances (compared to performance evaluated on the baseline \mathcal{D}_{test} or random sampling of instances). The result further demonstrates that the identified *inconsistent* instances have worse performance when compared to *uncertain* instances.

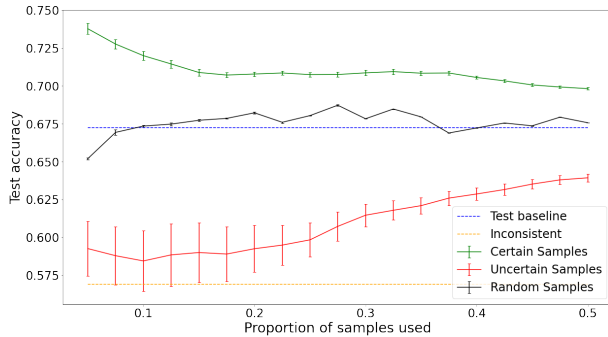


Figure 5. Example on CUTRACT of how Data-SUITE instance stratification can be used to improve downstream performance, contrasted with baseline \mathcal{D}_{test} (blue) or random selection (black).

Table 3 shows the *MPI* scores across methods. In satisfying **D2**, of improving deployed model performance, Data-

Table 3. MPI metric across datasets for different methods

	SEER-CUTTRACT	Adult	Electricity
Data-SUITE	0.11 ± 0.015	0.64 ± 0.03	0.26 ± 0.03
BNN	0.08 ± 0.02	-0.15 ± 0.02	-0.005 ± 0.01
CONFORMAL	0.05 ± 0.01	-0.07 ± 0.07	0.12 ± 0.03
ENSEMBLE	0.01 ± 0.02	-0.03 ± 0.02	-0.02 ± 0.02
GP	0.05 ± 0.04	0.56 ± 0.02	0.04 ± 0.04
MCD	0.01 ± 0.01	-0.16 ± 0.01	0.15 ± 0.03
QR	-0.10 ± 0.03	0.12 ± 0.06	0.15 ± 0.06

SUITE consistently outperforms other methods, providing the greatest performance improvement, with the lowest variability across datasets. The result suggests that Data-SUITE identifies the most appropriate *certain* and *uncertain* instances, accounting for the performance improvement. Overall, the quality of stratification by Data-SUITE has not been matched by any benchmark uncertainty estimation method.

4.4. Data-SUITE usage with diverse downstream models

Data-SUITE operates to identify instances, independent of the downstream task predictive model. Previous experiments use a fixed downstream model (e.g. RF). We now evaluate a more diverse set downstream models, i.e. RF, MLP, GBT and include a ROBUST model- Median-of-Means (MOM) estimators (Lecué et al., 2020)). For each downstream model, we compare the performance on instances identified by Data-SUITE and baselines. We rank them based on MPI (see Eq.8). i.e larger MPI=higher rank. Ideally, the rank should be invariant across models, showing the instances identified by Data-SUITE are impactful, no matter the downstream model.

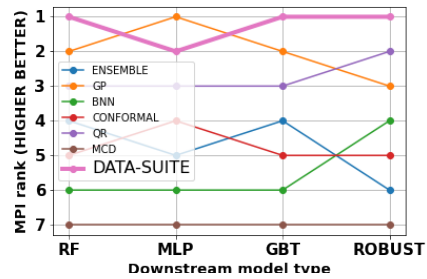


Figure 6. Rank assessment for a diverse set of downstream models, showing Data-SUITE has the most consistent ordering across different models

We show results for the Adult dataset in Fig. 6, with the other datasets included in Appendix C.5. Overall, the results are similar, whereby Data-SUITE’s identification remains the most appropriate, with consistent highest MPI rank, no matter the model. In contrast, baseline methods rank differently depending on the model. This shows that Data-SUITE identifies impactful instances no matter the downstream task-specific model.

4.5. Use-Case: Data-SUITE in the hands of users

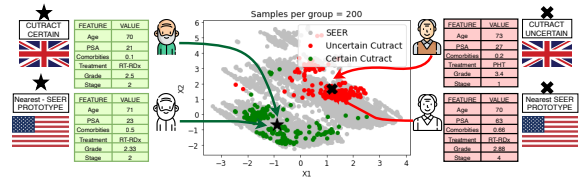
We now demonstrate how users can practically leverage Data-SUITE to better understand their data. We do so by profiling the incongruous regions identified by Data-SUITE and highlight the insights which can be garnered independent of a model. This satisfies **D1**, where the quantitative profiling provides valuable insights that could assist data owners to characterize where to collect more data and if this is not possible, to understand the data’s limitations.

For visual purposes, we embed the identified *certain* and *uncertain* instances into 2-D low-dimensional space as shown in Fig. 7. We clearly see that the *certain* and *uncertain* instances are distinct regions and that they lie in-distribution as evidenced by the embedding projection. This reinforces the quantitative findings of the previous experiment (i.e., that the identified instances are not OOD). We further highlight centroid, “average prototypes” of the *certain* and *uncertain* regions as a digestible example of the region, which can easily be understood by stakeholders. For the SEER-CUTRACT analysis, in addition to prototypes for \mathcal{D}_{test}^{Cut} regions, we can also find the nearest neighbor SEER (USA) prototypes for each instance. Comparing the average and nearest neighbor prototypes assists us to tease out the incongruence between the two geographic sites.

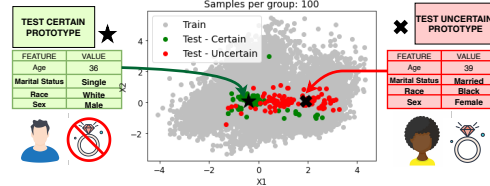
Overall, Fig. 7, in their respective captions, highlights the most valuable insights, quantitatively garnered on the basis of Data-SUITE across all three datasets. We conduct a more detailed analysis of the regions in Appendix C.3, to outline further potential practitioner usage.

5. Discussion

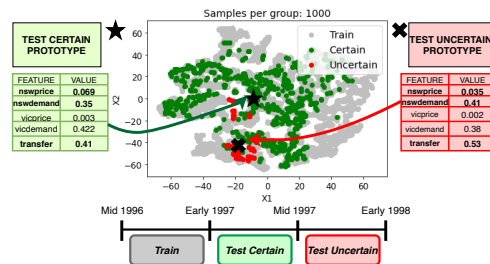
Automation should not replace the expertise and judgment of a data scientist in understanding the data, nor will it replace the ingenuity required to build better models. In this spirit, we developed Data-SUITE and illustrated its capability, across multiple datasets, to empower data scientists to perform more insightful data exploration, as well as, enable more reliable model deployment. We address these use-cases for the understudied problem of in-distribution heterogeneity and propose a flexible data-centric solution. This permits the identification of impactful instances, independent of a task-specific model. Data-SUITE allows to perform stratification of test data into inconsistent and uncertain instances with respect to training data. This stratification has been shown to be in line with downstream performance and to provide valuable insights for profiling incongruent test instances in a rigorous and quantitative way. The quality of this stratification by Data-SUITE is not matched by any benchmark uncertainty estimation method (data-centric or not). The promising result serves two roles. 1. The method could be used by practitioners both to improve data exploration, as well as, enable more reliable



i. SEER-CUTRACT: CUTRACT *certain* instances are similar to their SEER nearest prototypes, whilst CUTRACT *uncertain* instances are different to their nearest SEER prototypes (e.g. PSA).



ii. Adult: The *certain* and *uncertain* instances, represent two different demographics, aligning with the known dataset biases toward age. The *uncertain* instances specifically highlight a sub-group of black, females.



iii. Electricity: The *certain* instances are similar to the training set in features and time. The *uncertain* instances identified represent a later time period, wherein concept drift has likely occurred.

Figure 7. Insights of prototypes identified by Data-SUITE. Tables describe the average prototypes for *certain* and *uncertain* instances.

model deployment. 2. Data-SUITE opens up future avenues to advance the data-centric AI research agenda, taking it a step further by both explaining why instances might be classed as uncertain or inconsistent. Further, how this information could be leveraged either to correct the identified instances or improve the data collection process to improve overall data quality. Finally, the current formulation has focused on tabular data. For usage with high-dimensional data, we refer the reader to Appendix A.2.3 for proposals on possible modifications to Data-SUITE.

Acknowledgements

The authors are grateful to Fergus Imrie, Zhaozhi Qian, Yuchao Qin, Krzysztof Kacprzyk, Kamile Stankeviciute and the 4 anonymous ICML reviewers for their useful comments & feedback. Nabeel Seedat would like to acknowledge Hameeda Saif for her constant support and feedback. Nabeel Seedat is funded by the Cystic Fibrosis Trust, Jonathan Crabbé by Aviva and Mihaela van der Schaar by the Office of Naval Research (ONR), NSF 1722516.

References

- Alaa, A. and Van Der Schaar, M. Discriminative jackknife: Quantifying uncertainty in deep learning via higher-order influence functions. In *International Conference on Machine Learning*, pp. 165–174. PMLR, 2020.
- Algan, G. and Ulusoy, I. Image classification with deep learning in the presence of noisy labels: A survey. *Knowledge-Based Systems*, 215:106771, 2021.
- Asuncion, A. and Newman, D. Uci machine learning repository, 2007.
- Balasubramanian, V., Ho, S.-S., and Vovk, V. *Conformal prediction for reliable machine learning: theory, adaptations and applications*. Newnes, 2014.
- Bedford, T. and Cooke, R. M. Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial Intelligence*, 32(1):245–268, 2001.
- Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., and Kasneci, G. Deep neural networks and tabular data: A survey. *arXiv preprint arXiv:2110.01889*, 2021.
- Boström, H., Linusson, H., Löfström, T., and Johansson, U. Evaluation of a variance-based nonconformity measure for regression forests. In *Symposium on Conformal and Probabilistic Prediction with Applications*, pp. 75–89. Springer, 2016.
- Breunig, M. M., Kriegel, H.-P., Ng, R. T., and Sander, J. Lof: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp. 93–104, 2000.
- Chan, A., Alaa, A., Qian, Z., and Van Der Schaar, M. Unlabelled data improves bayesian uncertainty calibration under covariate shift. In *International Conference on Machine Learning*, pp. 1392–1402. PMLR, 2020.
- Duggan, M. A., Anderson, W. F., Altekruze, S., Penberthy, L., and Sherman, M. E. The surveillance, epidemiology and end results (seer) program and pathology: towards strengthening the critical relationship. *The American journal of surgical pathology*, 40(12):e94, 2016.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.
- Ghosh, S., Yao, J., and Doshi-Velez, F. Structured variational learning of bayesian neural networks with horseshoe priors. In *International Conference on Machine Learning*, pp. 1744–1753. PMLR, 2018.
- Gianfrancesco, M. A., Tamang, S., Yazdany, J., and Schmajuk, G. Potential biases in machine learning algorithms using electronic health record data. *JAMA internal medicine*, 178(11):1544–1547, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems 2014*, volume 27, 2014a.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014b.
- Graves, A. Practical variational inference for neural networks. *Advances in neural information processing systems*, 24, 2011.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of wasserstein gans. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 5769–5779, 2017.
- Gurumoorthy, K. S., Dhurandhar, A., Cecchi, G., and Aggarwal, C. Efficient data representation by selecting prototypes with importance weights. In *2019 IEEE International Conference on Data Mining (ICDM)*, pp. 260–269. IEEE, 2019.
- Haeusser, P., Frerix, T., Mordvintsev, A., and Cremers, D. Associative domain adaptation. In *Proceedings of the IEEE international conference on computer vision*, pp. 2765–2773, 2017.
- Harries, M. and Wales, N. S. Splice-2 comparative evaluation: Electricity pricing. 1999.
- Hsu, Y.-C., Shen, Y., Jin, H., and Kira, Z. Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10951–10960, 2020.
- Jain, A., Patel, H., Nagalapatti, L., Gupta, N., Mehta, S., Guttula, S., Mujumdar, S., Afzal, S., Sharma Mittal, R., and Munigala, V. Overview and importance of data quality for machine learning tasks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 3561–3562, 2020.
- Joe, H. *Dependence modeling with copulas*. CRC press, 2014.
- Johansson, U., Sönströd, C., and Linusson, H. Efficient conformal regressors using bagged neural nets. In *2015 Inter-*

- national Joint Conference on Neural Networks (IJCNN), pp. 1–8, 2015. doi: 10.1109/IJCNN.2015.7280763.
- Kandel, S., Paepcke, A., Hellerstein, J. M., and Heer, J. Enterprise data analysis and visualization: An interview study. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2917–2926, 2012.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations*, 2014.
- Kingma, D. P., Salimans, T., and Welling, M. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28:2575–2583, 2015.
- Koenker, R. and Hallock, K. F. Quantile regression. *Journal of economic perspectives*, 15(4):143–156, 2001.
- Kumagai, A. and Iwata, T. Unsupervised domain adaptation by matching distributions based on the maximum mean discrepancy via unilateral transformations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4106–4113, 2019.
- Kumar, A., Boehm, M., and Yang, J. Data management in machine learning: Challenges, techniques, and systems. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pp. 1717–1722, 2017.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Lecué, G., Lerasle, M., and Mathieu, T. Robust classification via mom minimization. *Machine Learning*, 109(8): 1635–1665, 2020.
- Lee, K., Lee, K., Lee, H., and Shin, J. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018.
- Lei, J., G’Sell, M., Rinaldo, A., Tibshirani, R. J., and Wasserman, L. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018.
- Leslie, D., Mazumder, A., Peppin, A., Wolters, M. K., and Hagerty, A. Does “ai” stand for augmenting inequality in the era of covid-19 healthcare? *BMJ*, 372, 2021.
- Li, Z., Zhao, Y., Botta, N., Ionescu, C., and Hu, X. Copod: copula-based outlier detection. In *2020 IEEE International Conference on Data Mining (ICDM)*, pp. 1118–1123. IEEE, 2020.
- Liu, F. T., Ting, K. M., and Zhou, Z.-H. Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1):1–39, 2012.
- Long, M., Cao, Y., Wang, J., and Jordan, M. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pp. 97–105. PMLR, 2015.
- Navratil, J., Arnold, M., and Elder, B. Uncertainty prediction for deep sequential regression using meta models. *arXiv preprint arXiv:2007.01350*, 2020.
- Ng, A., Aroyo, L., Coleman, C., Diamos, G., Reddi, V. J., Vanschoren, J., Wu, C.-J., and Zhou, S. Neurips data-centric ai workshop, 2021. URL <https://datacentricai.org/>.
- Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019.
- Park, C., Awadalla, A., Kohno, T., and Patel, S. Reliable and trustworthy machine learning for health using dataset shift detection. *Advances in Neural Information Processing Systems*, 34, 2021.
- Polyzotis, N., Roy, S., Whang, S. E., and Zinkevich, M. Data management challenges in production machine learning. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pp. 1723–1726, 2017.
- Prostate Cancer UK, C. Prostate cancer uk. URL <https://prostatecanceruk.org/>.
- Ren, M., Zeng, W., Yang, B., and Urtasun, R. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pp. 4334–4343. PMLR, 2018.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International conference on machine learning*, pp. 1530–1538. PMLR, 2015.
- Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P. K., and Aroyo, L. M. ”everyone wants to do the model work, not the data work”: Data cascades in high-stakes ai. 2021.
- Saria, S. and Subbaswamy, A. Tutorial: safe and reliable machine learning. *ACM Conference on Fairness, Accountability, and Transparency*, 2019.
- Seedat, N. and Kanan, C. Towards calibrated and scalable uncertainty representations for neural networks. *arXiv preprint arXiv:1911.00104*, 2019.
- Shafer, G. and Vovk, V. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.

- Sklar, A. Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8:229–231, 1959.
- Song, H., Kim, M., Park, D., Shin, Y., and Lee, J.-G. Learning from noisy labels with deep neural networks: A survey. *arXiv preprint arXiv:2007.08199*, 2020.
- Srivastava, A., Valkov, L., Russell, C., Gutmann, M. U., and Sutton, C. Veegan: Reducing mode collapse in gans using implicit variational learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 3310–3320, 2017.
- Suciu, D., Olteanu, D., Ré, C., and Koch, C. Probabilistic databases. *Synthesis lectures on data management*, 3(2): 1–180, 2011.
- Varshney, K. R. On mismatched detection and safe, trustworthy machine learning. In *2020 54th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–4. IEEE, 2020.
- Vovk, V. Transductive conformal predictors. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pp. 348–360. Springer, 2013.
- Vovk, V., Gammerman, A., and Shafer, G. Conformal prediction. *Algorithmic learning in a random world*, pp. 17–51, 2005.
- Wasserman, L. *All of statistics: a concise course in statistical inference*, volume 26. Springer, 2004.
- Williams, C. K. and Rasmussen, C. E. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- Yan, H., Ding, Y., Li, P., Wang, Q., Xu, Y., and Zuo, W. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 945–954, 2017. doi: 10.1109/CVPR.2017.107.
- Yang, J., Zhou, K., Li, Y., and Liu, Z. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021.
- Yoon, J., Zhang, Y., Jordon, J., and van der Schaar, M. Vime: Extending the success of self-and semi-supervised learning to tabular domain. *Advances in Neural Information Processing Systems*, 33, 2020.
- Zhang, L., Goldstein, M., and Ranganath, R. Understanding failures in out-of-distribution detection with deep generative models. In *International Conference on Machine Learning*, pp. 12427–12436. PMLR, 2021.
- Zhao, Y., Hu, X., Cheng, C., Wang, C., Wan, C., Wang, W., Yang, J., Bai, H., Li, Z., Xiao, C., et al. Suod: Accelerating large-scale unsupervised heterogeneous outlier detection. *Proceedings of Machine Learning and Systems*, 3, 2021.
- Zliobaite, I. How good is the electricity benchmark for evaluating concept drift adaptation. *arXiv preprint arXiv:1301.3524*, 2013.

A. Data-SUITE details & related work

A.1. Extended Related Work

We present a comparison of our framework Data-SUITE, and contrast it to the related work of uncertainty quantification and learning with noisy labels. Table 4, highlights both similarities and differences across multiple dimensions. We highlight 3 key features which distinguish Data-SUITE: (1) Data-centric uncertainty is a novel paradigm compared to the predominant model-centric approaches, (2) Our method offers increased flexibility, as it is used independent of task-specific predictive models. Any conclusions that we draw from Data-SUITE are not model-specific. (3) Our method provides theoretical guarantees concerning the validity of coverage.

Table 4. Comparison of related work

	Data-centric uncertainty	Model-centric uncertainty	Task Model independent	No noise assumptions	Coverage guarantees
Data-SUITE (Ours)	✓	×	✓	✓	✓
Uncertainty Quantification	×	✓	×	✓	×
Noisy labels	×	✓	×	×	×

A.2. Data-SUITE Details

We present a block diagram of our framework Data-SUITE in Figure 8. We next have in-depth discussions on both the generator and conformal predictor. We outline the motivations as well as technical details not covered in the main paper.

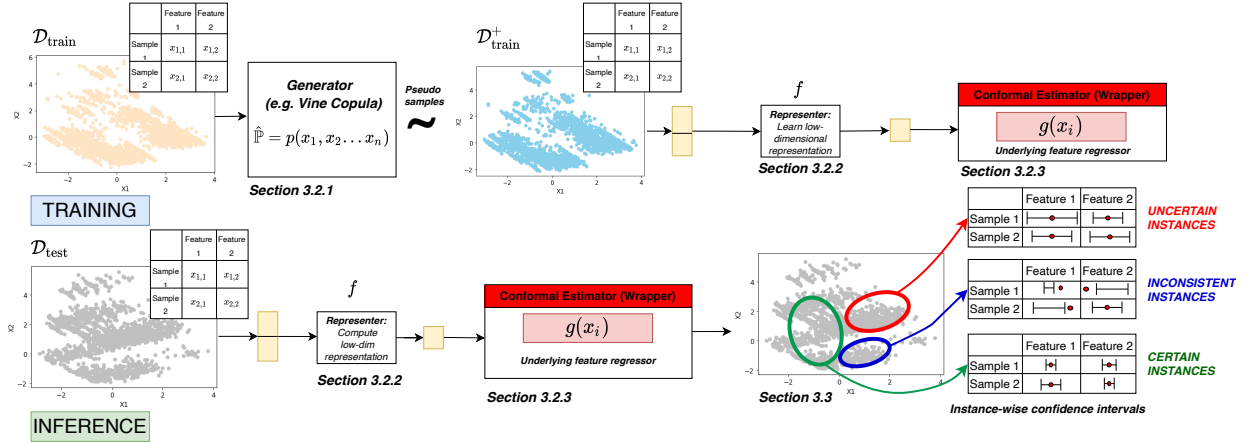


Figure 8. Outline of our framework **Data-SUITE**

A.2.1. GENERATOR: COPULAS

Motivation. Recall in our formulation, we have a set of training instances $\mathcal{D}_{\text{train}}$ and we wish to learn the dependency between features \mathcal{X} . Hence, we model the multivariate joint distribution of $\mathcal{D}_{\text{train}}$ to capture the dependence structure between random variables. A significant challenge is modeling distributions in high dimensions. Parametric approaches such as Kernel Density prediction (KDE), often using Gaussian distributions, are largely inflexible.

On the other hand, nonparametric versions are typically infeasible with complex data distributions and the curse of dimensionality. Additionally, while Variational Autoencoders (VAEs) (Kingma & Welling, 2014) and Generative Adversarial Networks (GANs) (Goodfellow et al., 2014a) can model learn the joint distributions, they both have limitations for our setting. VAEs make strong distributional assumptions (Rezende & Mohamed, 2015), while GANs involve training multiple models, which leads to associated difficulties (Srivastava et al., 2017; Gulrajani et al., 2017) (both from training computational burden and inherent to GANs - e.g., mode collapse).

An attractive approach, particularly for tabular data, is Copulas; which flexibly couple the marginal distributions of the different random variables into a joint distribution. One important reason lies in the following theorem:

Property A.1 (Sklar’s theorem). A d -dimensional random vector $\mathbf{X} = (X_1, \dots, X_d)$ with joint distribution F and marginal distributions F_1, \dots, F_d can be expressed as $f(X_1, \dots, X_d) = C_\theta\{F_1(X_1)\dots F_d(X_d)\}$, where $C_\theta : [0, 1]^d \rightarrow [0, 1]$ is a Copula. This is attractive in high dimensions, as it separates the learning of univariate marginal distributions from learning of the coupled multivariate dependence structure.

Parametric copulas have limited expressivity; hence we use pair copula constructions (vine copulas) (Bedford & Cooke, 2001), which are hierarchical and express a multivariate copula as a cascade of bivariate copulas.

Copula Details. To learn the copula, we factorize the d -dimensional copula density into a product of $d(d-1)/2$, bivariate conditional densities by means of the chain rule.

The graphical model which has edges being each bivariate-copula that encodes the (conditional) dependence between variables. The graphical model additionally consists of levels (where there are as many levels as features of the dataset). Each node will be a variable, and edges are coupling between variables based on bivariate copulas. As each level is constructed, the number of nodes decreases per level. The product over all pair-copula densities is then taken to define the joint copula.

Once we have learned the copula density, we sample the copula to obtain an augmented dataset of pseudo/synthetic samples. The copula samples are then easily transformed back to the natural data scale using the inverse probability integral transform.

Specifically, assume we have $\mathbf{U} = (U_1, U_2 \dots U_d)$ random variable $U(0, 1)$. We can then use the Copula C_θ to define variables $\mathbf{S} = (S_1, S_2 \dots S_d)$, where $S_1 = C^{-1}(U_1), S_2 = C^{-1}(U_2|U_1) \dots S_d = C^{-1}(U_d|U_1, U_2 \dots U_{d-1})$. This means that \mathbf{S} is the inverse Rosenblatt transform of \mathbf{U} and hence, $S \sim C$, which allows us to simulate synthetic/pseudo samples. For more on Copulas in general, we refer the reader to (Joe, 2014).

Complexity. As we go through each tree, there are a decreasing number of pair copulas. i.e. ($T_1 = d - 1, T_2 = d - 2 \dots T_{d-1} = 1$). Hence, the complexity of this algorithm is $O(n) \times d \times \text{truncation level}$, which for all purposes is $O(n)$.

A.2.2. CONFORMAL PREDICTOR

Motivation. Conformal prediction allows us to transform any underlying point predictor into a valid interval predictor. We will not discuss the generalized framework of conformal prediction (*Transductive Conformal prediction*), which requires model training to be redone for every data point. This has a large computational burden for modern datasets with many datapoints. For more, see (Shafer & Vovk, 2008; Balasubramanian et al., 2014).

We instead only discuss *Inductive Conformal prediction*, which is used in Data-SUITE. The inductive method splits the two processes needed: 1. the training of the underlying model and 2. computing the conformal estimates.

Conformal Prediction Details. Practically, we split the training set ($|\mathcal{D}_{\text{train}}^+| = n$) into two disjoint sets, namely the proper training set and calibration set: $\mathcal{D}_{\text{train}}^+ = \mathcal{D}_{\text{train2}}^+ \sqcup \mathcal{D}_{\text{cal}}^+$, where $|\mathcal{D}_{\text{train2}}^+| = m$ and $|\mathcal{D}_{\text{cal}}^+| = n - m$.

We use the proper training set to create our prediction rules for the feature-wise regressor (g). The calibration set is used for “conformalization”, i.e. for computing the non-conformity scores and p-values.

The non-conformity score μ_i of each example is a function which computes the disagreement (i.e. non-conformance) between the prediction and the true value. Note, we only compute non-conformity scores on the calibration set.

To obtain our intervals, we need to determine the critical value ϵ based on the non-conformity scores. We first sort them in descending order. The critical value ϵ is then the $\lceil (|\mathcal{D}_{\text{cal}}^+| + 1)(1 - \alpha) \rceil$ -th smallest residual (Vovk, 2013). Consequently, for any confidence level $(1 - \alpha)$, we can use the critical value to find $p(x) > \alpha$, which corresponds to the maximum and minimum values with p-values larger than α . As a consequence, we can obtain the maximum and minimum values of our predictive intervals.

The whole process is detailed in Algorithm 1. Normalization to obtain adaptive instance-wise intervals can be done using a normalization function sigma_i , as described in the main paper.

Algorithm 1 General Inductive Conformal prediction

function **CE** ($\alpha, \mathcal{D}_{\text{train2}}^+, \mathcal{D}_{\text{cal}}^+, \mu$)

Input : Significance α , nonconformity measure μ , proper training set, $\mathcal{D}_{\text{train2}}^+$ and calibration set, $\mathcal{D}_{\text{cal}}^+$

Output : Interval estimate

```

1 Train the underlying model  $g$  on  $\mathcal{D}_{\text{train2}}^+$ 
2 Compute non-conformity scores on the calibration set  $\mathcal{D}_{\text{cal}}^+$ 
3  $P = \{ \}$ ; foreach  $(\mathbf{x}, x_i) \in \mathcal{D}_{\text{cal}}^+$  do
4   |  $\mu_i(x) = |x_i - g_i \circ f(\mathbf{x})|$ ;
5   | Add  $\mu_i(x)$  to  $P$ 
6 end
7 CONFORMALIZATION.
8 Sort  $P$  in descending order to obtain scores  $S$ 
9 Determine the critical value of  $\epsilon \leftarrow [(|\mathcal{D}_{\text{cal}}^+| + 1)(1 - \alpha)]$ -th smallest residual in  $S$ .
10 Construct the interval predictor for each new value:
11 procedure  $\Gamma^\alpha(\mathbf{x} : \mathcal{X})$ 
12   | foreach  $\mathbf{x} \in \mathcal{D}_{\text{test}}$  do
13   |   | Apply  $x_m = g_i \circ f(\mathbf{x})$ 
14   | end
15   | return  $[l_i(x), r_i(x)] = [x_m - \epsilon, x_m + \epsilon]$ 

```

Remarks on theoretical guarantees. A motivation for conformal prediction in high-stakes settings is the theoretical guarantees on CI coverage validity (See Property A.2).

i.e., at a confidence level $(1 - \alpha)$ of 95% ($\alpha = 0.05$), the true value will be within the CIs in at least 95% of the cases. The framework is non-parametric and only makes the exchangeability assumption (which we detail next) and the guarantees of validity of coverage hold for any choice of dataset, underlying model or nonconformity measure – which makes it a versatile option.

Property A.2 (Validity). *Under the exchangeability assumption, the conformal predictor will return an interval, $\mathbb{P}(Y \in [l_i(x), r_i(x)]) \geq 1 - \alpha$, i.e, error $\leq \alpha$.*

The validity of the CI holds if the data is independent and identically distributed (IID). In practice, the weaker assumption of exchangeability (see Assumption A.1) also guarantees validity (Lei et al., 2018). This means that we are not required to impose any additional requirements for the validity of the CI, since the aforementioned assumptions on the underlying data are typically made for any ML model.

We also highlight that the validity is maintained even in the case of normalization, as long as the proper training set and calibration set are disjoint.

Assumption A.1 (Exchangeability). *In a dataset of n observations, the data points do not follow any particular order, i.e., all n permutations are equiprobable. Exchangeability is weaker than IID observations; however, IID observations satisfy exchangeability.*

A.2.3. DATA-SUITE BEYOND TABULAR

Currently, the focus of Data-SUITE has primarily engaged with tabular data being the most ubiquitous data type across industries (Borisov et al., 2021). However, there is of course value in application to other high dimensional modalities such as images or text. We provide some possible ideas of how Data-SUITE could be adapted for these other modalities.

Generator. As discussed in the main text, a copula might not be ideal for very high-dimensional (large d_X) data in domains such as computer vision or genomics. In those cases, copula modeling can be replaced by domain-specific augmentation techniques. For example, Generative Adversarial Networks (GANs).

Representer. PCA while the workhorse of tabular data, is typically not suitable for modalities such as images. Methods such as autoencoders could easily replace PCA in this instance.

Conformal Predictor. The primary challenge lies in the conformal predictor. Data-SUITE builds feature-wise regressors – which fits the tabular feature-wise setting. Two alternatives are proposed for other modalities. 1. Instead of feature-wise regressors, to construct instance-wise regressors or 2. consider the components of an intermediate representation vector as the “features”.

B. Benchmarks & Experimental Details

B.1. Benchmarks & Implementations

B.1.1. DATA-SUITE

Implementation details. Data-SUITE adopts a pipeline-based approach to constructing CI’s, leveraging copula modeling, representation learning and conformal estimation.

We break down each of these below:

Copula Modeling: We use the Copula to estimate a multivariate distribution with univariate marginal distributions. We make use vine copulas (Bedford & Cooke, 2001). for this task.

Specifically, we use Direct-Vines (D-Vines), which impose constraints on the edges such that we only learn vines of the structure given in Figure 9 below. When fitting each tree, we select the “best” base copula (Gaussian, Frank, Clayton or Gumbel) based on likelihood.

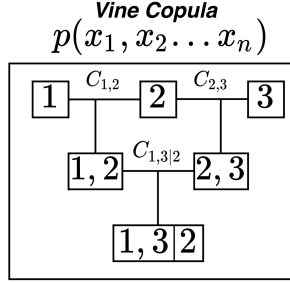


Figure 9. D-Vine structure

Finally, when we sample from the copula, we specifically sample $n_{samples} = |\mathcal{D}_{train}|$.

Representation learning: For the representer, we specifically make use of Principal Component Analysis (PCA). That said, we can simply replace this block with any alternative such as an autoencoder.

We pre-process all data prior to the representer, by standardizing the data, such that each feature has zero mean and unit variance (see Eq. 9). Note if the features were indeed categorical, we perform one-hot encoding of the features.

$$z = \frac{(x - \mu)}{\sigma} \quad (9)$$

When applying PCA, to learn the latent representation, we halve the dimensionality of \mathcal{D}_{train} , i.e. $d_X/2$.

Conformal Prediction: The two most important design decisions for conformal estimation is selecting the underlying feature-wise regressor (g_i) and the non-conformity score (μ_i).

We selected g_i , as in conventional machine learning by grid-search over different possible underlying models, evaluated on a validation set. We evaluate an MLP, KNN, Decision Trees and Random Forests. We ultimately selected a Decision Tree to serve as the base model for all $n = |d_X|$ feature-wise regressors. We use the following parameters max depth=None, min samples split=2, min samples leaf=5.

Interestingly, a simpler base model proved to be more effective and outperformed more parameterized models, which would require more compute. This is an additional advantage of Data-SUITE. We motivate that the simpler model might be directly related to the fact that we build feature-wise regressors, hence the mapping function is easier to approximate.

Finally, we use the absolute error non-conformity score. In practice, we can use any non-conformity score, however, intuitively for our application where we want to approximate the true value as closely as possible, the absolute error makes the most sense. Our non-conformity score used is given by Eq. 10 below.

$$\mu_i(x) = |x_i - (g_i \circ f)(x)| \quad (10)$$

B.1.2. BAYESIAN NEURAL NETWORK (BNN)

Bayesian modelling, when applied to neural networks, involves a likelihood function $p(Y | X, \theta)$, where the parameters θ are estimated by a neural network. In contrast to conventional neural networks which have point estimates for the parameters,

BNNs aim to learn distributions over parameters. However, modern neural networks have many parameters and weights. This makes exact inference largely intractable.

Hence, approximate Bayesian methods such as variational inference are often used instead. By this we mean we do not compute the posterior in the conventional manner. Rather, the variational approximation means we replace the posterior $p(\theta|\mathcal{D})$ with a more tractable variational distribution $q(\theta; \lambda)$. The Kullback-Leibler (KL) divergence between the true distribution and the variational distribution is then minimized as the loss function, which equates to maximizing the evidence lower bound (ELBO), see Equation 11. We note, however that the reparameterization trick (Kingma et al., 2015) is needed to make back propagation possible.

$$\mathcal{L}_{\text{VI}}(\mathcal{D}; \lambda) := D_{\text{KL}}(q(\theta; \lambda) || p(\theta)) - \mathbb{E}[\log p(\mathcal{D} | \theta)] \quad (11)$$

Implementation details. We train a 5-layer MLP model. A Gaussian prior is placed over the weights and we optimize the KL divergence during training. Our implementation is based on (Ghosh et al., 2018) and we use the implementation from ⁵.

B.1.3. DEEP ENSEMBLES (ENSEMBLE)

Deep Ensembles (Lakshminarayanan et al., 2017) is widely regarded as the state-of-the-art non-Bayesian uncertainty estimation method. The rationale is that training multiple randomly initialized models allows more robust predictions. Uncertainty can be computed as the variance of the different model predictions. We highlight some important features of Deep Ensembles: (1) optimization based on a *proper scoring rule* such that the loss has a unique minimum - which encourages the model to approximate the true probability distribution. However, the proper nature of the scoring rule introduces a distributional assumption. (2) Deep Ensembles uses adversarial perturbations based on the Fast Gradient Sign Method (Goodfellow et al., 2014b) and given by Equation 12.

$$\mathbf{x}' := \mathbf{x} + \boldsymbol{\eta} \odot \text{sgn}(\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, y; \theta)), \quad (12)$$

where \odot denotes element-wise multiplication.

This modifies the loss function used for gradient descent:

$$\mathcal{L}_{\text{tot}}(\mathbf{X}, \mathbf{y}; \theta) := \mathcal{L}(\mathbf{X}, \mathbf{y}; \theta) + \mathcal{L}(\mathbf{X}', \mathbf{y}; \theta), \quad (13)$$

To construct prediction intervals with an uncertainty estimate, an assumption of a conditionally normal distribution is assumed and the intervals, with uncertainty estimates are given as per Eq. 14.

$$\Gamma(\mathbf{x}^*) := \left[\mathbb{E}[y^* | \mathbf{x}^*] - \sqrt{\text{var}[y^* | \mathbf{x}^*]}, \mathbb{E}[y^* | \mathbf{x}^*] + \sqrt{\text{var}[y^* | \mathbf{x}^*]} \right]. \quad (14)$$

Implementation details. We have 5 models in the ensemble, all randomly initialized. Each model is a 3-layer MLP, which we train for 10 epochs. The learning rate was empirically determined based on a validation set. To compute the uncertainty estimates at test time we obtain predictions from each model in the ensemble. The prediction interval is computed as per Eq. 14.

B.1.4. GAUSSIAN PROCESS (GP)

Gaussian process (GP) models (?) are fully characterized by: the mean and covariance functions μ and Σ . The inference step can be performed exactly, as marginalizing the multivariate normal distributions can be written in closed form. See Equation 15, whereby conditioning on the dataset \mathcal{D} , we can compute the posterior.

$$(y^* | \mathbf{x}^*, \mathcal{D}) \sim \mathcal{N}\left(\mu^* + (\Sigma^*)^t \Sigma^{-1}(\mathbf{y} - \boldsymbol{\mu}), \Sigma^{**} - (\Sigma^*)^t \Sigma^{-1} \Sigma^*\right). \quad (15)$$

⁵<https://github.com/IBM/UQ360>

We note the two assumptions made by GPs: (1) the data is conditionally normally distributed and (2) the covariance selection is correct. That said, violations to these assumptions can severely impact performance. In addition, GPs have an issue of computational complexity due to matrix inversion of $O(n^3)$. For high-dimensional data, stochastic approximations are often used instead.

Implementation details. The GP is fit with a radial basis function (RBF) kernel and we make use of the Scikit-learn⁶ implementation of GPs which is based on (Williams & Rasmussen, 2006).

B.1.5. MONTE-CARLO DROPOUT (MCD)

Dropout layers are typically used as a regularizer at training time . (?). As shown by (Gal & Ghahramani, 2016) dropout networks can be used at test time (Monte-Carlo Dropout (MCD)) and are a variational approximation to deep Gaussian processes. Note that this induces the assumption of normality similar to GPs.

A key component of MCD, is that dropout at test time effectively allows us to obtain an ensemble of different models without having to retrain the model itself. For this, we do multiple stochastic forward passes at test time.

When making predictions, we compute a conditional mean. This is approximated by Monte-Carlo integration, i.e. mean of multiple forward passes. The prediction interval to characterize the uncertainty is given by Eq. 16:

$$\Gamma(\mathbf{x}^*) := \left[\mathbb{E}[y^* | \mathbf{x}^*] - \sqrt{\text{var}[y^* | \mathbf{x}^*]}, \mathbb{E}[y^* | \mathbf{x}^*] + \sqrt{\text{var}[y^* | \mathbf{x}^*]} \right]. \quad (16)$$

Implementation details. We train a 3-layer MLP with dropout ($p = 0.1$) for 10 epochs. The learning rate was empirically determined based on a validation set. To compute uncertainty estimates at test time, we perform 20 Monte-Carlo Samples (forward passes at test time). The prediction interval is computed as per Eq. 16.

B.1.6. QUANTILE REGRESSION (QR)

Quantile regressors estimate the conditional quantiles of a distribution (rather than conditional mean). Consequently, there is a natural measure of the underlying distribution’s spread.

Typically, the MSE loss is replaced by the pinball loss, otherwise known as quantile loss, given by Equation 17, which balances the number of points above and below the quantiles (Koenker & Hallock, 2001). Neural networks are easily applied and optimize the loss function with the appropriate number of output heads.

$$\mathcal{L}_{\text{pinball}}(\mathcal{T}) := \sum_{(\mathbf{x}, y) \in \mathcal{T}} \max \left((1 - \alpha)(\hat{q}_\alpha(\mathbf{x}) - y), \alpha(y - \hat{q}_\alpha(\mathbf{x})) \right). \quad (17)$$

Implementation details. The base model for predicting the quantiles is a Gradient Boosting Regressor, with 10 estimators, max depth=5, minimum samples per leaf=5, minimum samples for split=10. We use the implementation from⁷.

⁶<https://scikit-learn.org>

⁷<https://github.com/IBM/UQ360>

B.2. Synthetic Experiment Details

B.2.1. DATASET & EXPERIMENT CONFIGURATIONS

The synthetic data $\mathbf{X} = [X_1, X_2, X_3]$ is drawn IID from a multivariate Gaussian distribution, parameterized by mean vector μ and a positive definite covariance matrix Σ .

$$\mu = \begin{bmatrix} 5.0 \\ 0.0 \\ 10.0 \end{bmatrix} \quad \Sigma = \begin{bmatrix} 3.40 & -2.75 & -2.00 \\ -2.75 & 5.50 & 1.50 \\ -2.00 & 1.50 & 1.25 \end{bmatrix}$$

We sample $n = 1000$ for both $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$ respectively. We encode inconsistency and uncertainty into the features of the test set $\mathcal{D}_{\text{test}}$ using a multivariate additive model $\hat{X} = X + Z$, where $Z \in R^{n \times m}$, is the perturbation matrix.

We conduct three experiments with different configurations (D_a, D_b, D_c), see Table 5

Table 5. Different configurations of the synthetic data

	Noise distribution	Perturbation Proportions	Perturbation variance
D_a	{Normal}	{0.1, 0.25, 0.5, 0.75}	{2}
D_b	{Normal}	{0.5}	{1, 2, 3}
D_c	{Beta, Gamma, Normal, Weibull}	{0.5}	{2}

B.2.2. DOWNSTREAM MODEL

In Section 4.2, we have a downstream task wherein we compute the MSE for different models. The base model which we train using $\mathcal{D}_{\text{train}}$ is a linear regression model.

B.3. Real-Data Experiment Details

B.3.1. DATASETS

SEER Dataset The SEER dataset consists of 240,486 patients enrolled in the American SEER program (Duggan et al., 2016). The dataset consists of features used to characterize **prostate cancer**: including age, PSA (severity score), Gleason score, clinical stage, treatments etc. A summary of the covariate features can be found in Table 6. The classification task is to predict patient mortality, which is binary label $\in \{0, 1\}$.

The dataset is highly imbalanced, where 94% of patients survive. Hence, we extract a balanced subset of of 20,000 patients (i.e. 10,000 with label=0 and 10,000 with label=1).

Table 6. Summary of features for the SEER Dataset (Duggan et al., 2016)

Feature	Range
Age	37 – 95
PSA	0 – 98
Comorbidities	0, 1, 2, ≥ 3
Treatment	Hormone Therapy (PHT), Radical Therapy - RDx (RT-RDx), Radical Therapy -Sx (RT-Sx), CM
Grade	1, 2, 3, 4, 5
Stage	1, 2, 3, 4
Primary Gleason	1, 2, 3, 4, 5
Secondary Gleason	1, 2, 3, 4, 5

CUTRACT Dataset The CUTRACT dataset is a private dataset consisting of 10,086 patients enrolled in the British Prostate Cancer UK program (Prostate Cancer UK). Similar, to the SEER dataset, it consists of the same features to characterize prostate cancer. Additionally, it has the same task to predict mortality. A summary of the covariate features can be found in Table 7.

Once again, the dataset is highly imbalanced, hence we then choose extract a balanced subset of of 2,000 patients (i.e. 1000 with label=0 and 1000 with label=1).

Table 7. Summary of features for the CUTRACT Dataset (Prostate Cancer UK)

Feature	Range
Age	44 – 95
PSA	1 – 100
Comorbidities	0, 1, 2, ≥ 3
Treatment	Hormone Therapy (PHT), Radical Therapy - RDx (RT-RDx), Radical Therapy -Sx (RT-Sx), CM
Grade	1, 2, 3, 4, 5
Stage	1, 2, 3, 4
Primary Gleason	1, 2, 3, 4, 5
Secondary Gleason	1, 2, 3, 4, 5

ADULT Dataset The ADULT dataset (Asuncion & Newman, 2007) has 32,561 instances with a total of 13 attributes capturing demographic (age, gender, race), personal (marital status) and financial (income) features amongst others. The classification task predicts whether a person earns over \$50K or not. We encode the features (e.g. race, sex, gender etc) and a summary can be found in Table 8.

There is a known bias between gender and income in the dataset. We perform a train-test split such that $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$ have approximately equal sizes, with 15,378 and 14,784 samples respectively. In particular, to highlight the data exploration use-case.

Note that there is an imbalance across certain features. However, these are amongst the sensitive attributes. Thus, we do not want to balance the datasets based on this, as we wish to show both in the data exploration and model deployment experiments that we can identify these biases in the datasets. Balancing might eliminate these biases.

Table 8. Summary of features for the ADULT Dataset (Asuncion & Newman, 2007)

Feature	Range
Age	17 – 90
education-num	1 – 16
marital-status	0, 1
relationship	0, 1, 2, 3, 4
race	0, 1, 2, 3, 4
sex	0, 1
capital-gain	0, 1
capital-loss	0, 1
hours-per-week	1 – 99
country	0, 1
employment-type	0, 1, 2, 3
salary	0, 1

ELECTRICITY Dataset. The Electricity dataset (Harries & Wales, 1999), represents energy pricing in Australia, over the period of May-1996 to December 1998, with recordings every 30 minutes giving 45312 samples. The dataset records the energy prices and demand for New South Wales and Victoria, and the amount of power transferred between the two states. The goal is to predict whether the transfer price increases or decreases.

The covariates outlined in Table 9 are normalized to the interval $[0, 1]$

We temporally partition the dataset where $\mathcal{D}_{\text{train}}$ is Mid-1996 to early-1997 and $\mathcal{D}_{\text{test}}$ is early-1997 to 1998. We note that the dataset has been characterized as having concept shift for some features over the test period (however without an explicit timepoint or label). This could be due to behavioral changes or consumption pattern changes (Zliobaite, 2013). Hence, $\mathcal{D}_{\text{test}}$ consists of data which is both congruous and incongruous with $\mathcal{D}_{\text{train}}$.

Table 9. Summary of features for the ELECTRICITY Dataset (Harries & Wales, 1999)

Feature	Range
data	0 – 1
period	0 – 1
nswprice	0 – 1
nswdemand	0 – 1
vicprice	0 – 1
vicdemand	0 – 1
transfer	0 – 1
class	0 – 1

B.3.2. DOWNSTREAM MODEL

In Section 4.3, we have a downstream task for each of the three datasets. For all the datasets, our base model which we train using $\mathcal{D}_{\text{train}}$ is a Random Forest Classification model with 100 estimators in the ensemble and splits are based on the 'Gini' criterion.

C. Additional Experiments

This appendix presents additional experiments, validating further properties of Data-SUITE, conducting further comparisons or deep-dives into the regions identified by Data-SUITE and what insights can be garnered.

C.1. Data-SUITE Ablation

Data-SUITE adopts a pipeline-based approach. Hence, to better understand the effect of each component, we perform an ablation study of different constituent components. The constituent components are then compared to the complete pipeline, which we denote as Data-SUITE (ALL).

The two components that we explicitly test are:

- Data-SUITE (CONF): Test the conformal predictor without the representer. The conformal prediction process and instance stratification are as defined in the main paper.
- Data-SUITE (COP): Test the copula by itself. For this ablation, we fit the copula on $\mathcal{D}_{\text{train}}$. Then, to compute intervals per feature, we condition on the remaining features and sample 100 estimates from the copula. For example, for feature x_1 , we condition on $x_2 \dots x_n$. The uncertainty estimates are then the variance of these samples, which then is used to stratify the samples as before.

Firstly, we compare the coverage for each of the constituent components as per Table 10 for different configurations of the synthetic experiment. We see that Data-SUITE (All) is the only method to maintain coverage guarantees across all configurations. That said, we see a significant divergence between Data-SUITE (CONF) and Data-SUITE (COP), which suggests the conformal estimator is the most important component. The performance gap to Data-SUITE (All) is then likely the representer.

Table 10. Coverage of constituent components of Data-SUITE

PERTURBATION	Proportion (D_a)			Variance (D_b)	
	.1	.25	.5	1	2
Data-SUITE (All)	.97	.96	.96	.95	.96
Data-SUITE (CONF)	.89	.88	.90	.87	.90
Data-SUITE (COP)	.16	.15	.16	.17	.16

Secondly, we compare the downstream MSE for each of the constituent components as per Table 11 for different configurations of the synthetic experiment. Recall that a lower MSE (i.e. closer to Train Data (BASELINE)) is desired. Data-SUITE (All) outperforms the constituent components and is less sensitive to perturbations. Interestingly, despite the higher coverage for Data-SUITE (CONF) vs Data-SUITE (COP), when evaluating for a downstream task, Data-SUITE (COP) in fact produces results which are less sensitive to perturbations.

Table 11. Downstream MSE for ablations of constituent components of Data-SUITE

PERTURBATION	Proportion (D_a)			Variance (D_b)	
	.1	.25	.5	1	2
<i>Train Data (BASELINE)</i>	<i>.067</i>	<i>.059</i>	<i>.068</i>	<i>.065</i>	<i>.068</i>
Test Data	.222	.513	.889	.275	.889
Data-SUITE (All)	.069	.122	.197	.104	.197
Data-SUITE (CONF)	.125	.396	.846	.293	.846
Data-SUITE (COP)	.220	.277	.451	.236	.451

Takeaway: Both results show reduced performance for the constituent components of Data-SUITE, with each component having different individual sensitivities. This highlights the necessity of the interconnected Data-SUITE (All) framework.

C.2. Regions identified by Data-SUITE as NOT OOD

In this experiment, we address the question whether the regions identified by Data-SUITE as uncertain or inconsistent are in fact OOD. We benchmark four widely used methods (with different detection mechanisms), which have been applied in the literature for OOD and outlier detection:

- Mahalanobis distance (Lee et al., 2018)
- SUOD: Accelerating Large-Scale Unsupervised Heterogeneous Outlier Detection (Zhao et al., 2021)
- COPOD: Copula-Based Outlier Detection (Li et al., 2020)
- Isolation Forest (Liu et al., 2012)

We note for SUOD, much like the original paper, we make use of an ensemble of base estimators namely: Local outlier factor (LOF) (Breunig et al., 2000), COPOD, Isolation Forest.

For each dataset, we have the instance IDs identified by Data-SUITE for various proportions. We then apply each of the aforementioned methods and compute the overlap between the predicted OOD/Outlier instances and our identified uncertain and inconsistent instances.

The results for the overlap of uncertain instances can be found in Fig. 10. We see minimal overlap across methods ranging between 2-18%.

We additionally, evaluate the confidence scores (for those methods which provide confidence scores as outputs). The goal is to see if an instance is predicted as OOD/outlier, then with what confidence does the detection method ascribe to the instance. The results in Fig. 11 suggest that the methods were often unconfident, with average confidence scores ranging between 5-50%. This suggests the identified uncertain instances are unlikely OOD.

A similar question might be asked for the inconsistent instances. The results are similar to the uncertain instances as shown in Fig. 12. This result, similar to uncertain instances, suggests the identified inconsistent instances are unlikely OOD.

Takeaway: The Data-SUITE identified *uncertain* and *inconsistent* instances are unlikely OOD. The reason is the limited overlap of predicted OOD and both *uncertain* and *inconsistent* instances, coupled with unconfident (low probability) predictions for OOD.

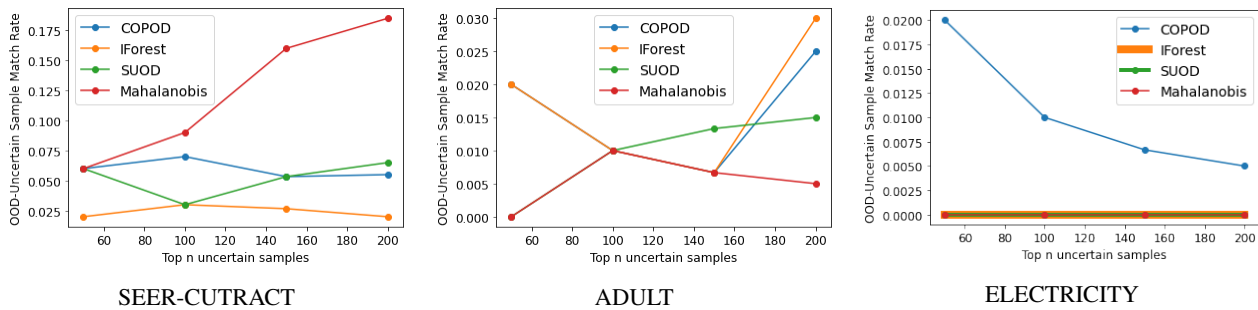


Figure 10. OOD-Uncertain Instances Match Rate or overlap

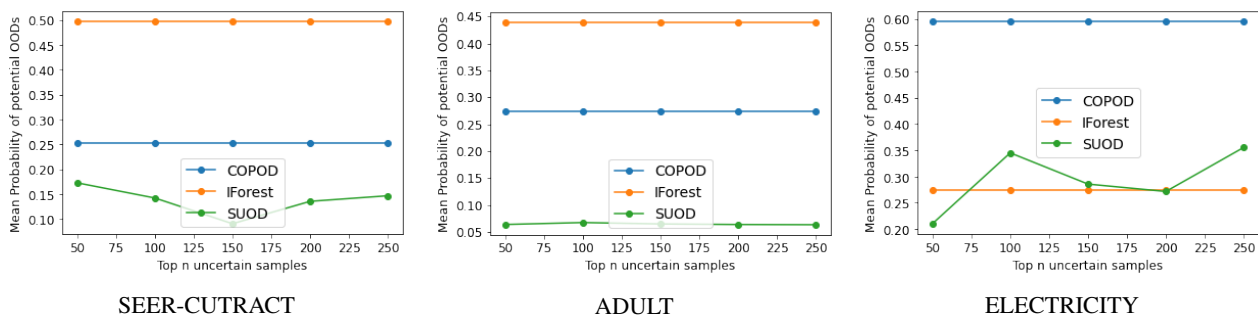


Figure 11. Mean probability of predicted OOD, i.e. confidence

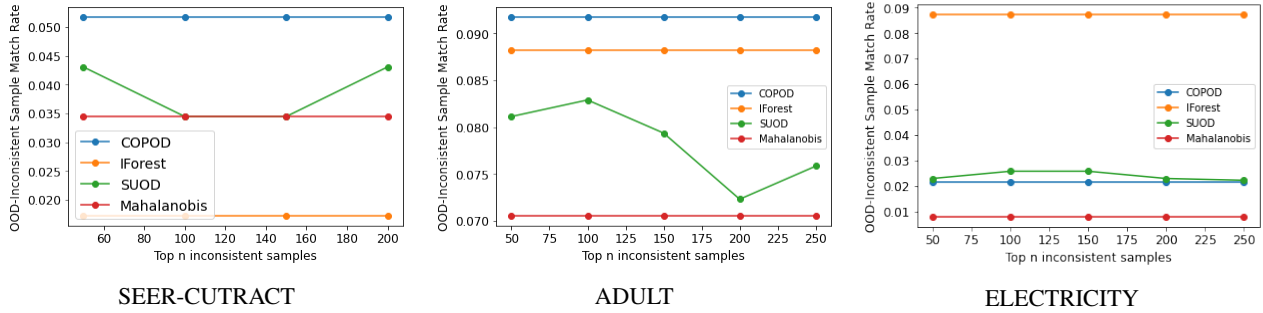


Figure 12. OOD-Inconsistent Instances Match Rate or overlap

C.3. EDA & digestible prototype insights

As discussed in the main paper, we now conduct a more detailed analysis of the regions identified by Data-SUITE, as well as, the digestible average prototypes. This illustrates the full potential of the detail that practitioners can get from Data-SUITE and how they can practically use Data-SUITE to garner insights about the data. We present the diagram again for easy reference (see Fig 14).

SEER-CUTRACT. Data-SUITE identifies distinct *certain* (Green) and *uncertain* (Red) regions of ID data as shown in Fig. 14 (i). Beyond average prototypes for $\mathcal{D}_{\text{test}}$ regions i.e., CUTRACT (UK), we also find the nearest neighbor SEER (USA) prototypes for each instance in the identified regions. Comparing these prototypes assists us to tease out the differences between the two geographic sites. We note three specific insights which can assist end-users. (1) *certain* instances represent less severe patients than the *uncertain* instances (see PSA values). (2) *certain* instances: CUTRACT (UK) and SEER (USA) prototypes are similar in all-feature values. (3) *uncertain* instances: CUTRACT (UK) and SEER (USA) prototypes have differences in certain feature values (PSA, more comorbidities, different treatment, staging score).

We conduct an extended deep-dive below, highlighting features such as PSA, that show no difference on a population level, or for *certain* instances. However, when teasing out *uncertain* instances and their prototypes, we see this difference. This example illustrates the full potential of how Data-SUITE can be used by practitioners to uncover differences across sites (to benefit clinical practice), whilst also identifying patients where model performance would either be reliable or substandard (independent of the model).

We highlight this in Figure 13, which shows no PSA difference between the USA and UK on a population level, or for certain instances, but when teasing out *uncertain* instances and their prototypes we see this difference.

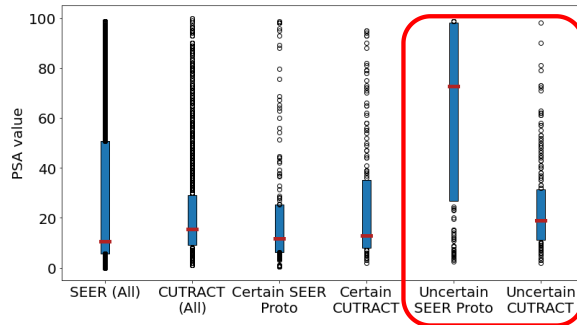
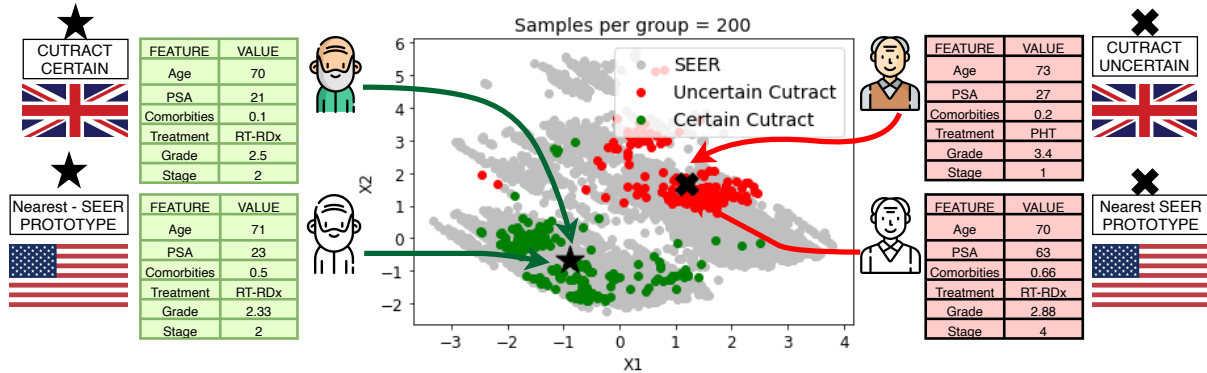


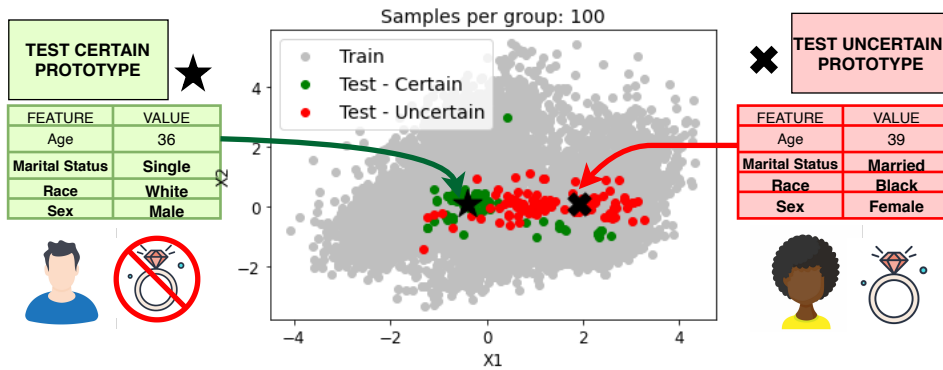
Figure 13. PSA Deep-Dive, which highlights how the difference is not evident at population level or for certain instances. Only the uncertain instances highlight the heterogeneity. However, it does highlight common support based on the whiskers of the box plot.

We extend this analysis to further illustrate the full potential of the detail that practitioners can get from Data-SUITE. We do this by comparing the Earth Movers Distance (EMD) - see Fig. 15 which is a common metric to flag drift. We see no PSA difference between the USA and UK on a population level, or for certain instances. However, when teasing out *uncertain* instances and their prototypes, we see this difference.

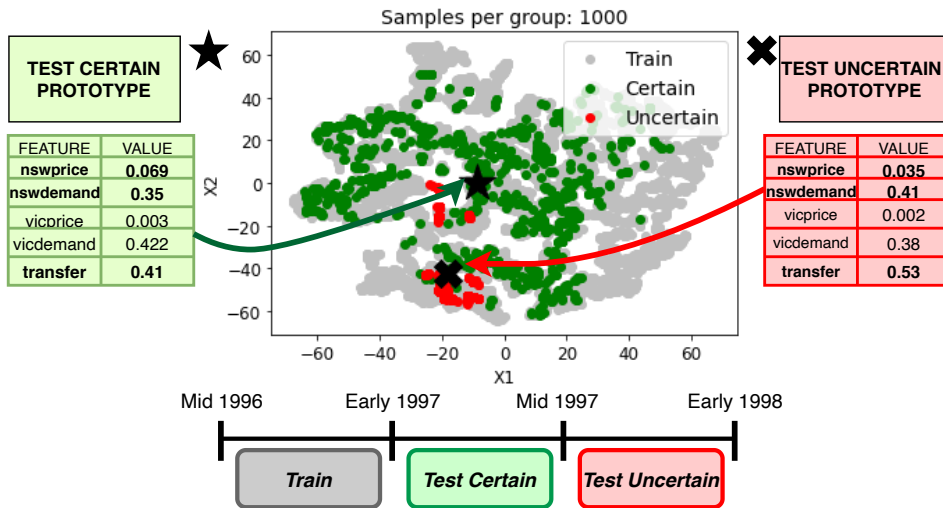
Takeaway: We highlight that while not evident at population level, the heterogeneous groups can be teased out by using Data-SUITE. This fully demonstrates the added capability of what Data-SUITE offers to practitioners.



i. SEER-CUTRACT: CUTRACT *certain* instances are similar to their SEER nearest prototypes, whilst CUTRACT *uncertain* instances are different to their nearest SEER prototypes (e.g. PSA).

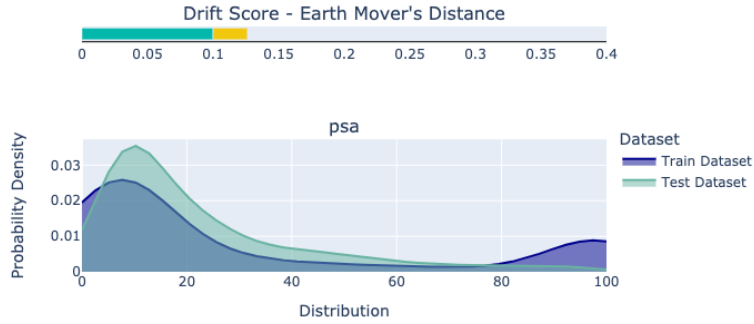


ii. Adult: The *certain* and *uncertain* instances, represent two different demographics, aligning with the known dataset biases toward females. The *uncertain* instances specifically highlight a sub-group of black, females.

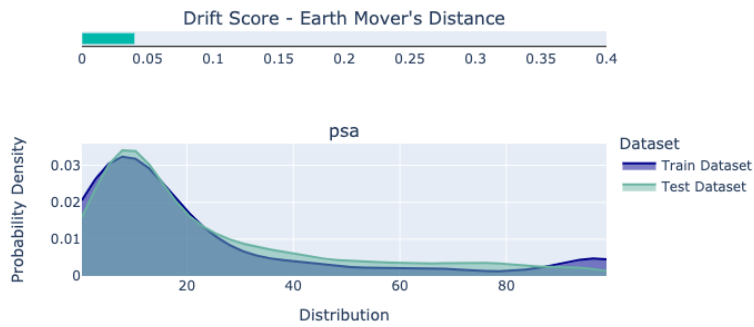


iii. Electricity: The *certain* instances are similar to the training set in features and time. The *uncertain* instances identified, represent a later time period, wherein concept drift has likely occurred.

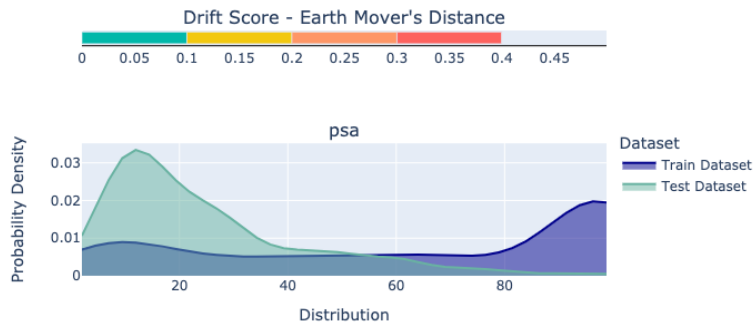
Figure 14. Insights of prototypes identified by Data-SUITE. Tables describe the average prototypes for *certain* and *uncertain* instances.



Earth Movers Distance on the FULL Dataset (Population): Train (SEER) & Test (CUTRACT) - No Difference



Earth Movers Distance: *CERTAIN*: Train (SEER Nearest Neighbors) & Test (CUTRACT) - No Difference



Earth Movers Distance: *UNCERTAIN*: Train (SEER Nearest Neighbors) & Test (CUTRACT)- Clear Difference

Figure 15. Earth Movers Distance Deep-Dive into PSA, which highlights how the difference is not evident at the population level or in certain instances. Only the uncertain instances highlight the heterogeneity

Adult. The dataset has a known bias between gender and income, which Data-SUITE successfully identifies. However, an added dimension between the most *certain* and *uncertain* instances as shown in Fig. 7 (ii) is that marital status and race are also relevant biases.

We capture this finding with prototypes, (i) *certain* prototype: younger, single, white, males (ii) *uncertain* prototype: older, married, black, females.

These prototypes can inform end-users such as data scientists and stakeholders of the dangers of naively building models on this dataset without first considering these issues.

Electricity. The dataset has concept drift over time, which Data-SUITE identifies as the *uncertain* instances, which have increased transfer of energy between NSW and Victoria, increased NSW energy price and decreased demand relative to the *certain* instances.

These prototypes can inform end-users of this change in consumption habits, despite the data still lying in-distribution.

Additionally, we note as per Fig. 7 (iii), the *certain* instances capture the time-frame closer to the training data (Start 1997-Mid 1997), whilst the *uncertain* instances are later in time (Mid 1997-Early 1998), hence more likely affected by the concept drift.

C.4. Assessment of different CI regressors g_i

In our formulation, we select a CI regressor (g_i). We wish to assess if the choice of CI regressor can impact the performance on a fixed downstream task predictive model.

We hence carry out an experiment where we assess different $g_i \in \{RF, MLP, SVRandTree\}$, based on the Mean Performance Improvement (MPI) -Eq. 8 on a downstream task.

Takeaway. We find the variance of MPI across CI regressor types ranges between 0.04%-3% on our 3 datasets. The low variability in MPI is desirable, as it indicates minimal impact of the choice of CI regressor on Data-SUITE’s performance.

C.5. Rank for a diverse set of downstream models

We now present additional results for the experiment with a diverse set of downstream models, as outlined in Sec. 4.4. See Figure 16

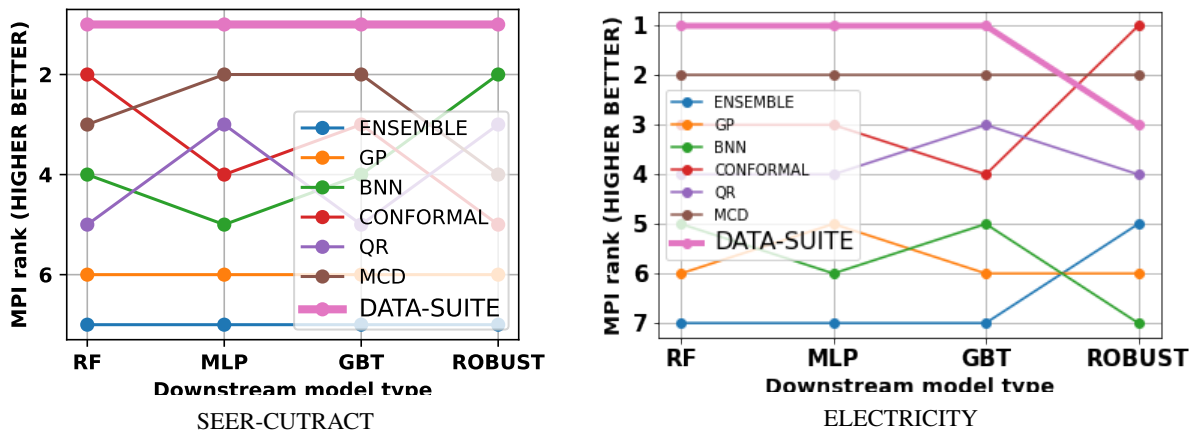


Figure 16. Rank assessment for on diverse downstream models

C.6. Inconsistent instances λ sensitivity

As described in Section 3.3, an instance x is *inconsistent* if the fraction of inconsistent features is above a predetermined threshold $\lambda \in [0, 1]$: $\nu(x) > \lambda$. In our implementation, we use $\lambda = 0.5$, i.e. if more than 50% of features are inconsistent, then the sample is considered inconsistent.

For completeness, we conduct an analysis of the sensitivity to the values of $\lambda \in [0, 1]$. We show the accuracy score as a function of λ , as well as, the number of instances that would be classified as inconsistent for that value of λ . The results are shown in Fig. 17. The flat, steady performance for small λ and then drop off approximately around $\lambda = 0.5$ across all datasets, suggests that our chosen value of $\lambda = 0.5$ is indeed a sensible choice.

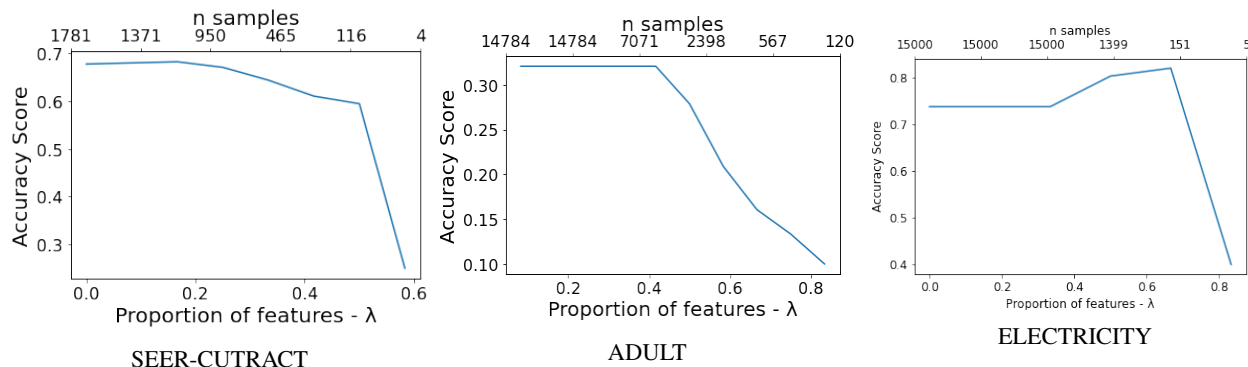


Figure 17. λ sweep to flag inconsistent instances

Takeaway: Based on the sweep across all three datasets, we see that $\lambda = 0.5$ is a sensible choice in practice.

C.7. Computation time comparison

All experiments were run on CPU on a MacBook Pro with an Intel Core i5 and 16GB RAM. Besides the actual performance values, practitioners often are interested in the computation time associated with different methods. This is especially true if the algorithms are applied to very large datasets with many instances.

We hence conduct a comparison of the computation time needed by each method to both train and to construct the predictive intervals for all instances in $\mathcal{D}_{\text{test}}$. We present the results in Table 12, with the recorded computation time provided to the nearest minute.

Table 12. Comparison of computation time across methods to the nearest minute

	SEER-CUTTRACT	Adult	Electricity
Data-SUITE	4	1	1
BNN	2	1	1
CONFORMAL	2	1	2
ENSEMBLE	28	3	3
GP	22	15	24
MCD	19	13	19
QR	1	1	1

Takeaway: Data-SUITE does not have the fastest computation time (due to the inter-connected framework). However, the computation times are neither too dissimilar and nor prohibitive for practitioners to use.

C.8. Comparison to domain adaptation

In the paradigm of domain adaptation, a popular metric is the maximum mean discrepancy (MMD) (Gretton et al., 2012). For example, (Kumagai & Iwata, 2019; Long et al., 2015; Yan et al., 2017; Haeusser et al., 2017) have used MMD as the metric to compare distributions and subsequently, to optimize to minimize the distributional difference of representations.

The MMD metric is a distance-based metric that compares the mean embeddings of two probability distributions source S and target T , in a reproducing kernel Hilbert space H_k . This is given by Eq. 18

$$MMD(S, T) = \|\mu_k(S) - \mu_k(T)\|_{H_k} \quad (18)$$

We can compute unbiased estimates of samples from the two distributions after we apply the kernel trick (in this case we use a Radial Basis Function Kernel).

In the domain adaptation literature, the goal is to minimize the latent feature divergence by optimizing the MMD. Hence, MMD is a fundamental component of domain adaptation and is used to identify instances with discrepancies between source and target domains.

As a comparison to our approach (Data-SUITE), we cast the problem as a potential domain adaptation problem and apply MMD, where the source distribution ($S \sim \mathcal{D}_{\text{train}}$) and target distribution ($T \sim \mathcal{D}_{\text{test}}$). We then use the computed MMD to stratify instances, where lower MMD means *certain* and large MMD means *uncertain*.

The mean performance improvement (MPI) is then computed for instances stratified based on MMD, similar to the experiment in Section 4.3. Table 13 contrasts the results using MMD vs Data-SUITE to stratify instances.

Takeaway. The approach of Data-SUITE achieves greater average performance improvement based on MPI across all three datasets, when compared to framing the problem in the context of domain adaptation (using the MMD metric to identify and stratify instances).

Table 13. Comparison of average performance improvement for Data-SUITE vs MMD (Domain Adaptation)

	Data-SUITE	MMD
SEER-CUTTRACT	0.11	0.088
Adult	0.64	-0.06
Electricity	0.26	-0.17

C.9. Comparison to prototypes

We compare the samples identified by Data-SUITE with prototypes. Specifically, we compare ProtoDash (Gurumoorthy et al., 2019).

Takeaway. When explaining $\mathcal{D}_{\text{test}}$ using prototypes, the prototypes uniformly cover the manifold, i.e. we cannot easily identify clusters of “certain” or “uncertain” as in Data-SUITE.

Alternatively, explaining $\mathcal{D}_{\text{test}}$ with prototypes from $\mathcal{D}_{\text{train}}$, the prototypes match the “certain” instances identified by Data-SUITE (DS), where the Pearson $r \approx 0.8$ for DS(certain) vs prototypes. On the contrary, it doesn’t match the uncertain samples where DS(uncertain) vs prototypes ≈ 0.55 .

Thus, the prototypes either are “uniform samples” or “certain samples” with no distinct “uncertain samples”. We conclude that Data-Suite has an advantage of identifying BOTH “certain” and “uncertain” samples, relevant for reliable model deployment (Desiderata 2)