Conformance Constraint Discovery: Measuring Trust in Data-Driven Systems

Anna Fariha*[†] University of Massachusetts Amherst afariha@cs.umass.edu Ashish Tiwari[†] Arjun Radhakrishna Sumit Gulwani Microsoft {astiwar,arradha,sumitg}@microsoft.com

Alexandra Meliou University of Massachusetts Amherst ameli@cs.umass.edu

ABSTRACT

The reliability of inferences made by data-driven systems hinges on the data's continued conformance to the systems' initial settings and assumptions. When serving data (on which we want to apply inference) deviates from the profile of the initial training data, the outcome of inference becomes unreliable. We introduce conformance constraints, a new data profiling primitive tailored towards quantifying the degree of non-conformance, which can effectively characterize if inference over that tuple is untrustworthy. Conformance constraints are constraints over certain arithmetic expressions (called projections) involving the numerical attributes of a dataset, which existing data profiling primitives such as functional dependencies and denial constraints cannot model. Our key finding is that projections that incur low variance on a dataset construct effective conformance constraints. This principle yields the surprising result that lowvariance components of a principal component analysis, which are usually discarded for dimensionality reduction, generate stronger conformance constraints than the high-variance components. Based on this result, we provide a highly scalable and efficient techniquelinear in data size and cubic in the number of attributes-for discovering conformance constraints for a dataset. To measure the degree of a tuple's non-conformance with respect to a dataset, we propose a quantitative semantics that captures how much a tuple violates the conformance constraints of that dataset. We demonstrate the value of conformance constraints on two applications: trusted machine learning and data drift. We empirically show that conformance constraints offer mechanisms to (1) reliably detect tuples on which the inference of a machine-learned model should not be trusted, and (2) quantify data drift more accurately than the state of the art.

ACM Reference Format:

Anna Fariha, Ashish Tiwari, Arjun Radhakrishna, Sumit Gulwani, and Alexandra Meliou. 2021. Conformance Constraint Discovery: Measuring Trust in Data-Driven Systems. In *Proceedings of the 2021 International Conference on*

SIGMOD '21, June 20-25, 2021, Virtual Event, China

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-8343-1/21/06...\$15.00 https://doi.org/10.1145/3448016.3452795 Management of Data (SIGMOD '21), June 20-25, 2021, Virtual Event, China. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3448016.3452795

1 INTRODUCTION

The core of modern data-driven systems typically comprises of models learned from large datasets, and they are usually optimized to target particular data and workloads. While these data-driven systems have seen wide adoption and success, their reliability and proper functioning hinge on the data's continued conformance to the systems' initial settings and assumptions. When the serving data (on which the system operates) deviates from the profile of the initial data (on which the system was trained), system performance degrades and system behavior becomes unreliable. Therefore, a mechanism to assess the trustworthiness of a system's inferences is paramount, especially for systems that perform safety-critical or high-impact operations.

A machine-learned (ML) model typically works best if the serving dataset follows the profile of the dataset the model was trained on; when it doesn't, the model's inference can be unreliable. One can profile a dataset in many ways, such as by modeling the data distribution of the dataset [2], or by finding the (implicit) *constraints* that the dataset satisfies [49]. Distribution-oriented approaches learn data likelihood (e.g., joint or conditional distribution) from the training data, and can be used to check if the serving data is unlikely. An unlikely tuple does not necessarily imply that the model would fail for it. The problem with the distribution-oriented approaches is that they tend to overfit, and, thus, are overly conservative towards unseen tuples, leading them to report many such false positives.

We argue that certain constraints offer a more effective and robust mechanism to quantify trust of a model's inference on a serving tuple. The reason is that learning systems implicitly exploit such constraints during model training, and build models that assume that the constraints will continue to hold for serving data. For example, when there exist high correlations among attributes in the training data, learning systems will likely reduce the weights assigned to redundant attributes that can be deduced from others, or eliminate them altogether through dimensionality reduction. If the serving data preserves the same correlations, such operations are inconsequential; otherwise, we may observe model failure.

In this paper, we characterize datasets with a new data-profiling primitive, *conformance constraints*, and we present a mechanism to identify *strong* conformance constraints, whose violation indicates unreliable inference. Conformance constraints specify constraints over *arithmetic relationships* involving multiple numerical attributes

^{*}Part of the research was done while the author was an intern at Microsoft. †Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

	Departure Date	Departure Time [DT]	Arrival Time [AT]	Duration (min) [DUR]
t_1	May 2	14:30	18:20	230
t_2	July 22	09:05	12:15	195
t_3	June 6	10:20	12:20	115
t_4	May 19	11:10	13:05	117
t_5	April 7	22:30	06:10	458

Figure 1: Sample of the airlines dataset (details are in Section 6.1), showing departure, arrival, and duration only. The dataset does not report arrival date, but an arrival time earlier than departure time (e.g., last row), indicates an overnight flight. All times are in 24 hour format and in the same time zone. There is some noise in the values.

of a dataset. We argue that a tuple's conformance to the conformance constraints is more critical for accurate inference than its conformance to the training data distribution. This is because any violation of conformance constraints is likely to result in a catastrophic failure of a learned model that is built upon the assumption that the conformance constraints will always hold. Thus, we can use a tuple's deviation from the conformance constraints as a proxy for the trust on a learned model's inference for that tuple. We proceed to describe a real-world example of conformance constraints, drawn from our case-study evaluation on *trusted machine learning* (TML).

EXAMPLE 1. We used a dataset with flight information that includes data on departure and arrival times, flight duration, etc. (Fig. 1) to train a linear regression model to predict flight delays. The model was trained on a subset of the data that happened to include only daytime flights (such as the first four tuples). In an empirical evaluation of the regression accuracy, we found that the mean absolute error of the regression output more than quadruples for overnight flights (such as the last tuple t₅), compared to daytime flights. The reason is that tuples representing overnight flights deviate from the profile of the training data that only contained daytime flights. Specifically, daytime flights satisfy the conformance constraint that "arrival time is later than departure time and their difference is very close to the flight duration", which does not hold for overnight flights. Note that this constraint is just based on the covariates (predictors) and does not involve the target attribute delay. Critically, although this conformance constraint is unaware of the regression task, it was still a good proxy of the regressor's performance. In contrast, approaches that model data likelihood may report long daytime flights as unlikely, since all flights in the training data $(t_1 - t_4)$ were also short flights, resulting in false alarms, as the model works very well for most daytime flights, regardless of the duration (i.e., for both short and long daytime flights).

Example 1 demonstrates that when training data has *coincidental* relationships (e.g., the one between AT, DT, and DUR for daytime flights), then ML models may *implicitly* assume them as *invariants*. Conformance constraints can capture such data invariants and flag non-conforming tuples (overnight flights) during serving.

Conformance constraints. Conformance constraints complement the existing data profiling literature, as the existing constraint models, such as functional dependencies and denial constraints, cannot model arithmetic relationships. For example, the conformance constraint of Example 1 is: $-\epsilon_1 \leq AT - DT - DUR \leq \epsilon_2$, where ϵ_1 and ϵ_2 are small values. Conformance constraints can capture complex linear dependencies across attributes within a *noisy* dataset. For example, if the flight departure and arrival data reported the hours

and the minutes across separate attributes, the constraint would be on a different arithmetic expression: $(60 \cdot \operatorname{arrHour} + \operatorname{arrMin}) - (60 \cdot \operatorname{depHour} + \operatorname{depMin}) - \operatorname{duration}$.

The core component of a conformance constraint is the arithmetic expression, called *projection*, which is obtained by a linear combination of the numerical attributes. There is an unbounded number of projections that we can use to form arbitrary conformance constraints. For example, for the projection AT, we can find a broad range [ϵ_3 , ϵ_4], such that all training tuples in Example 1 satisfy the conformance constraint $\epsilon_3 \leq AT \leq \epsilon_4$. However, this constraint is too inclusive and a learned model is unlikely to exploit such a weak constraint. In contrast, the projection AT – DT – DUR leads to a stronger conformance constraint with a narrow range as its bounds, which is selectively permissible, and, thus, more effective.

Challenges and solution sketch. The principal challenge is to discover an *effective* set of conformance constraints that are likely to affect a model's inference implicitly. We first characterize "good" projections (that construct effective constraints) and then propose a method to discover them. We establish through theoretical analysis two important results: (1) A projection is good over a dataset if it is almost constant (i.e., has low variance) over all tuples in that dataset. (2) A set of projections, collectively, is good if the projections have small pair-wise correlations. We show that low variance components of a principal component analysis (PCA) on a dataset yield such a set of projections. Note that this is different from—and, in fact, completely opposite of—the traditional approaches (e.g., [51]) that perform multidimensional analysis based on the high-variance principal components, after reducing dimensionality using PCA.

Scope. Fig. 2 summarizes prior work on related problems, but our scope differs significantly. Specifically, we can detect if a serving tuple is non-conforming with respect to the training dataset only based on its predictor attributes, and require no knowledge of the ground truth. This setting is essential in many practical applications when we observe extreme verification latency [62], where ground truths for serving tuples are not immediately available. For example, consider a self-driving car that is using a trained controller to generate actions based on readings of velocity, relative positions of obstacles, and their velocities. In this case, we need to determine, only based on the sensor readings (predictors), when the driver should be alerted to take over vehicle control. Furthermore, we do not assume access to the model, i.e., model's predictions on a given tuple. This setting is necessary for (1) safety-critical applications, where the goal is to quickly alert the user, without waiting for the availability of the prediction, (2) auditing and privacy-preserving applications where the prediction cannot be shared, and (3) when we are unaware of the detailed functionality of the system due to privacy concerns or lack of jurisdiction. We focus on identifying tuple-level non-conformance as opposed to dataset-level non-conformance that usually requires observing entire data's distribution. However, our tuple-level approach trivially extends (by aggregation) to the entire dataset.

Contrast with prior art. We now discuss where conformance constraints fit with respect to the existing literature (Fig. 2).

Data profiling techniques. Conformance constraints fall under the umbrella of data profiling, which refers to the task of extracting technical metadata about a given dataset [1]. A key task in data profiling

	Legend	constraints			viol	ation	setting			technique				TML				
HP: Hyper Parameter FD: Functional Dependency DC: Denial Constraint \oslash : Does not require \bot : Not applicable \star : Supports via extension !: Partially		parametric	arithmetic	approximate	conditional	notion of weight	interpretable	continuous	tuple-wise	noisy data	numerical attr.	categorial attr.	⊘thresholds	⊘distance metric	⊘HP tuning	scalable	task agnostic	⊘access to model
	Conformance Constraints	\checkmark	1	✓	✓	✓	*	\checkmark	✓	✓	✓	✓	\checkmark	✓	✓	\checkmark	✓	\checkmark
	FD [48]						✓			_		\checkmark	\checkmark	\checkmark	\checkmark		÷	4
	Approximate FD [42]			\checkmark			V			×	-	~		~	√ _	-	essed in prior wor	
50	Conditional ED [10]	Ι.	-	~	/		V			~	~	✓ ✓		7	1	1		
i.	Pattern FD [50]	÷	-		V		× ./		✓ ./	× ./		✓ .∕		✓ ./	✓ ./			
rof	Soft FD [32]			$\overline{}$				\checkmark	v	V	$\overline{}$	V		V	V	\checkmark		
ta I	Relaxed FD [10]		\vdash	\checkmark		-	\checkmark			\checkmark	\checkmark	\checkmark			\checkmark	-		
Dai	FDX [75]						\checkmark			\checkmark		\checkmark				\checkmark		
<u> </u>	Differential Dependency [60]						\checkmark				\checkmark	\checkmark			\checkmark		Jdr.	
	DC [8, 12]	!		\checkmark	\checkmark	\checkmark	\checkmark		\checkmark	\checkmark	\checkmark	\checkmark		\checkmark	\checkmark		t a(
	Approximate DC [44, 49]	!		\checkmark	\checkmark	\checkmark	\checkmark		\checkmark	\checkmark	\checkmark	\checkmark		\checkmark	\checkmark		1	1
	Statistical Constraint [74]			\checkmark			\checkmark	\checkmark		\checkmark	\checkmark	\checkmark		\checkmark	\checkmark	\leq		
	Ordinary Least Square	\checkmark	\checkmark	\checkmark			*	\checkmark	\checkmark	\checkmark	\checkmark	*	\checkmark	\checkmark	\checkmark	\leq		
50	Total Least Square	\checkmark	\checkmark	\checkmark			*	\checkmark	\checkmark	\checkmark	\checkmark	*	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
.Ĩ	Auto-encoder [15]			_	L			\checkmark	\checkmark	\checkmark	\checkmark	\checkmark				\checkmark	\checkmark	\checkmark
arr	Schelter et al. [56] ⁺			_	L			\checkmark			\checkmark	\checkmark		\checkmark	\checkmark	\checkmark		
цЧ	Jiang et al. [34]			_	L			\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark			\checkmark		
	Hendrycks et al. [26]			_	L				\checkmark	\checkmark	\checkmark	\checkmark	V.	\checkmark	\checkmark	\checkmark		
Model's Prediction Probability				_	L			\checkmark	\checkmark			Vä	arie	es				

Requires additional information

Figure 2: Conformance constraints complement existing data profiling primitives and provide an efficient mechanism to quantify trust in prediction, with minimal assumption on the setting.

is to learn relationships among attributes. Functional dependencies (FD) [48] and their variants only capture if a relationship exists between two sets of attributes, but do not provide a closed-form (parametric) expression of the relationship. Using the FD "{AT, DT} \rightarrow {DUR}" to model the constraint of Example 1 suffers from several limitations. First, since the data is noisy, no exact FD can be learned. Metric FDs [40] allow small variations in the data, but hinge on appropriate distance metrics and thresholds. For example, if time is split across two attributes (hour and minute), the distance metric is non-trivial: it needs to encode that $\langle hour = 4, min = 59 \rangle$ and $\langle hour = 5, min = 1 \rangle$ are similar, while $\langle hour = 4, min = 1 \rangle$ and $\langle hour = 5, min = 59 \rangle$ are not. In contrast, conformance constraints can model the composite attribute (60 · hour + minute) by automatically discovering the coefficients 60 and 1.

Denial constraints (DC) [8, 12, 44, 49] encapsulate a number of different data-profiling primitives such as FDs and their variants (e.g, [19]). Exact DCs can adjust to noisy data by adding predicates until the constraint becomes exact over the entire dataset, but this can make the constraint extremely large and complex, which might even fail to provide the desired generalization. For example, a finite DC—whose language is limited to universally quantified first-order logic—cannot model the constraint of Example 1, which involves an arithmetic expression (addition and multiplication with a constant). Expressing conformance constraints requires a richer language that includes linear arithmetic expressions. Pattern functional dependencies (PFD) [50] move towards addressing this limitation of DCs, but they focus on text attributes: they are regex-based and treat digits as characters. However, modeling arithmetic relationships of numerical attributes requires interpreting digits as numbers.

To adjust for noise, FDs and DCs either relax the notion of constraint violation or allow a user-defined fraction of tuples to violate the (strict) constraint [10, 30, 32, 40, 42, 44, 49]. Some approaches [32, 74, 75] use statistical techniques to model other types of data profiles such as correlations and conditional dependencies. However, they require additional parameters such as noise and violation thresholds and distance metrics. In contrast, conformance constraints do not require any parameter from the user and work on noisy datasets.

Existing data profiling techniques are not designed to learn what ML models exploit and are sensitive to noise in the numerical attributes. Moreover, data constraint discovery algorithms typically search over an exponential set of candidates, and hence, are not scalable: their complexity grows exponentially with the number of attributes or quadratically with data size. In contrast, our technique for deriving conformance constraints is highly scalable (linear in data size) and efficient (cubic in the number of attributes). It does not explicitly explore the candidate space, as PCA—which lies at the core of our technique—performs the search *implicitly* by iteratively refining weaker constraints to stronger ones.

Learning techniques. While *ordinary least square* finds the lowestvariance projection, it minimizes observational error on only the target attribute, and, thus, does not apply to our setting. *Total least square* offers a partial solution as it takes observational errors on all predictor attributes into account; but, it finds only one projection the lowest variance one—that fits the data tuples best. But there may exist other projections with slightly higher variances and we consider them all. As we show empirically in Section 6.2, constraints derived from multiple projections, collectively, capture various aspects of the data, and result in an effective data profile targeted towards certain tasks such as data-drift quantification [21].

Contributions. We make the following contributions:

- We ground the motivation of our work with two case studies on trusted machine learning (TML) and data drift. (Section 2)
- We introduce and formalize conformance constraints, a new data profiling primitive that specifies constraints over arithmetic relationships among numerical attributes of a dataset. We describe a *conformance language* to express conformance constraints, and a *quantitative semantics* to quantify how much a tuple violates the conformance constraints. In applications of constraint violations, some violations may be more or less critical than others. To capture that, we consider a notion of constraint importance, and weigh violations against constraints accordingly. (Section 3)
- We formally establish that strong conformance constraints are constructed from projections with small variance and small mutual correlation on the given dataset. Beyond simple linear constraints (e.g., the one in Example 1), we derive *disjunctive* constraints, which are disjunctions of linear constraints. We achieve this by dividing the dataset into disjoint partitions, and learning linear constraints for each partition. We provide an efficient, scalable, and highly parallelizable algorithm for computing a set of linear conformance constraints and disjunctions over them. We also analyze its runtime and memory complexity. (Section 4)
- We formalize the notion of *unsafe* tuples in the context of trusted machine learning and provide a mechanism to detect unsafe tuples using conformance constraints. (Section 5)
- We empirically analyze the effectiveness of conformance constraints in two case-study applications—TML and data-drift quantification. We show that conformance constraints can reliably predict the trustworthiness of linear models and quantify data drift precisely, outperforming the state of the art. (Section 6)

2 CASE STUDIES

Like other data-profiling primitives, conformance constraints have general applicability in understanding and describing datasets. However, their true power lies in quantifying the degree of a tuple's non-conformance with respect to a reference dataset. Within the scope of this paper, we focus on two case studies to motivate our work. We provide an extensive evaluation over these applications in Section 6.

Trusted machine learning (TML) refers to the problem of quantifying trust in the inference made by a machine-learned model on a new serving tuple [34, 52, 55, 65, 71]. When a model is trained using a dataset, the conformance constraints for that dataset specify a safety envelope [65] that characterizes the tuples for which the model is expected to make trustworthy predictions. If a serving tuple falls outside the safety envelope (violates the conformance constraints), then the model is likely to produce an untrustworthy inference. Intuitively, the higher the violation, the lower the trust. Some classifiers produce a confidence measure along with the class prediction, typically by applying a softmax function to the raw numeric prediction values. However, such confidence measures are not well-calibrated [25, 34], and, therefore, cannot be reliably used as a measure of trust in the prediction. Additionally, a similar mechanism is not available for inferences made by regression models.

In the context of TML, we formalize the notion of *unsafe tuples*, on which the prediction may be untrustworthy. We establish that conformance constraints provide a sound and complete procedure for detecting unsafe tuples, which indicates that the search for conformance constraints should be guided by the class of models considered by the corresponding learning system (Section 5).

Data drift [6, 24, 43, 51] specifies a significant change in a dataset with respect to a reference dataset, which typically requires that systems be updated and models retrained. To quantify how much a dataset D' drifted from a reference dataset D, our three-step approach is: (1) compute conformance constraints for D, (2) evaluate the constraints on all tuples in D' and compute their violations (degrees of non-conformance), and (3) finally, aggregate the tuple-level violations to get a dataset-level violation. If all tuples in D' satisfy the constraints, then we have no evidence of drift. Otherwise, the aggregated violation serves as the drift quantity.

While we focus on these two applications here, we mention other applications of conformance constraints in our technical report [21].

3 CONFORMANCE CONSTRAINTS

In this section, we first provide the general definition of conformance constraints. Then we propose a language for representing them. Finally, we define quantitative semantics over conformance constraints, which allows us to quantify their violation.

Basic notation. We use $\mathcal{R}(A_1, A_2, ..., A_m)$ to denote a relation schema where A_i denotes the i^{th} attribute of \mathcal{R} . We use Dom_i to denote the domain of attribute A_i . Then the set $\text{Dom}^m = \text{Dom}_1 \times \cdots \times \text{Dom}_m$ specifies the domain of all possible tuples. We use $t \in \text{Dom}^m$ to denote a tuple in the schema \mathcal{R} . A dataset $D \subseteq \text{Dom}^m$ is a specific instance of the schema \mathcal{R} . For ease of notation, we assume some order of tuples in D and we use $t_i \in D$ to refer to the i^{th} tuple and $t_i.A_j \in \text{Dom}_j$ to denote the value of the j^{th} attribute of t_i .

Conformance constraint. A conformance constraint Φ characterizes a set of allowable or conforming tuples and is expressed through a *conformance language* (Section 3.1). We write $\Phi(t)$ and $\neg \Phi(t)$ to denote that *t* satisfies and violates Φ , respectively.

DEFINITION 2 (CONFORMANCE CONSTRAINT). A conformance constraint for a dataset $D \subseteq \text{Dom}^m$ is a formula $\Phi : \text{Dom}^m \mapsto \{\text{True}, \text{False}\}$ such that $|\{t \in D \mid \neg \Phi(t)\}| \ll |D|$.

The set $\{t \in D \mid \neg \Phi(t)\}$ denotes atypical tuples in *D* that do not satisfy the conformance constraint Φ . In our work, we do not need to know the set of atypical tuples, nor do we need to purge the atypical tuples from the dataset. Our techniques derive constraints in ways that ensure there are very few atypical tuples (Section 4).

3.1 Conformance Language

Projection. A central concept in our conformance language is *projection*. Intuitively, a projection is a derived attribute that specifies a "lens" through which we look at the tuples. More formally, a projection is a function $F : \mathbf{Dom}^m \mapsto \mathbb{R}$ that maps a tuple $t \in \mathbf{Dom}^m$ to a real number $F(t) \in \mathbb{R}$. In our language for conformance constraints, we only consider projections that correspond to linear combinations of the numerical attributes of a dataset. Specifically, to define a projection, we need a set of numerical coefficients for all attributes of the dataset and the projection is defined as a sum over the attributes, weighted by their corresponding coefficients. We extend a projection F to a dataset D by defining F(D) to be the sequence of reals obtained by applying F on each tuple in D individually.

Grammar. Our language for conformance constraints consists of formulas Φ generated by the following grammar:

The language consists of (1) bounded constraints $lb \leq F(\vec{A}) \leq$ ub where *F* is a projection on **Dom**^{*m*}, \vec{A} is the tuple of formal parameters $(A_1, A_2, ..., A_m)$, and lb, ub $\in \mathbb{R}$ are reals; (2) equality constraints A = c where *A* is an attribute and *c* is a constant in *A*'s domain; and (3) operators (\triangleright , \land , and \lor ,) that connect the constraints. Intuitively, \triangleright is a switch operator that specifies which constraint ϕ applies based on the value of the attribute A, \land denotes conjunction, and \lor denotes disjunction. Formulas generated by ϕ and Ψ are called *simple constraints* and *compound constraints*, respectively. Note that a formula generated by ψ_A only allows equality constraints on a single attribute, namely *A*, among all the disjuncts.

EXAMPLE 3. Consider the dataset D consisting of the first four tuples $\{t_1, t_2, t_3, t_4\}$ of Fig. 1. A simple constraint for D is:

$$\phi_1: -5 \leq \mathsf{AT} - \mathsf{DT} - \mathsf{DUR} \leq 5$$

Here, the projection $F(\vec{A}) = AT - DT - DUR$, with attribute coefficients (1, -1, -1), 1b = -5, and ub = 5. A compound constraint is:

$\psi_2 : M = ``May" \triangleright -2 \le$	AT - DT - DUR	≤ 0
\lor M = "June" \triangleright 0 \leq	AT - DT - DUR	≤ 5
\vee M = "fuly" \triangleright -5 \leq	AT – DT – DUR	≤ 0

For ease of exposition, we assume that all times are converted to minutes (e.g., $06: 10 = 6 \times 60 + 10 = 370$) and M denotes the departure month, extracted from Departure Date.

Note that arithmetic expressions that specify linear combination of numerical attributes (highlighted above) are disallowed in denial constraints [12] which only allow raw attributes and constants (more details are in our technical report [21]).

3.2 **Quantitative Semantics**

Conformance constraints have a natural Boolean semantics: a tuple either satisfies a constraint or it does not. However, Boolean semantics is of limited use in practice, because it does not quantify the degree of constraint violation. We interpret conformance constraints using a quantitative semantics, which quantifies violations, and reacts to noise more gracefully than Boolean semantics.

The quantitative semantics $[[\Phi]](t)$ is a measure of the violation of Φ on a tuple *t*—with a value of 0 indicating no violation and a value greater than 0 indicating some violation. In Boolean semantics, if $\Phi(t)$ is True, then $[[\Phi]](t)$ will be 0; and if $\Phi(t)$ is False, then $[[\Phi]](t)$ will be 1. Formally, $[[\Phi]]$ is a mapping from **Dom**^{*m*} to [0, 1].

Quantitative semantics of simple constraints. We build upon ϵ -insensitive loss [70] to define the quantitative semantics of simple constraints, where the bounds 1b and ub define the ϵ -insensitive zone:¹

$$\begin{split} & \llbracket [\mathsf{lb} \le F(A) \le \mathsf{ub}] \rrbracket(t) \coloneqq \eta(\alpha \cdot \max(0, F(t) - \mathsf{ub}, \mathsf{lb} - F(t))) \\ & \llbracket \land (\phi_1, \dots, \phi_K) \rrbracket(t) \coloneqq \sum_k^K \gamma_k \cdot \llbracket \phi_k \rrbracket(t) \end{split}$$

Below, we describe the parameters of the quantitative semantics, and provide further details on them in our technical report [21].

Scaling factor $\alpha \in \mathbb{R}^+$.

Projections are unconstrained functions and different projections can map the same tuple to vastly different values. We use a scaling factor α to standardize the values computed by a projection *F*, and to bring the values of different projections to the same comparable scale. The scaling factor is automatically computed as the inverse of the standard deviation, $\alpha = \frac{1}{\sigma(F(D))}$. We set α to a large positive number when $\sigma(F(D)) = 0$.

Normalization function $\eta(.) : \mathbb{R} \mapsto [0, 1]$ **.**

The normalization function maps values in the range $[0, \infty)$ to the range [0, 1). While any monotone mapping from $\mathbb{R}^{\geq 0}$ to [0, 1) can be used, we pick $\eta(z) = 1 - e^{-z}$.

Importance factors $\gamma_k \in \mathbb{R}^+$, $\sum_k^K \gamma_k = 1$. The weights γ_k control the contribution of each bounded-projection constraint in a conjunctive formula. This allows for prioritizing constraints that are more significant than others. In our work, we derive the importance factor of a constraint automatically, based on its projection's standard deviation over D.

Quantitative semantics of compound constraints. Compound constraints are first simplified into simple constraints, and they get their meaning from the simplified form. We define a function $simp(\psi, t)$ that takes a compound constraint ψ and a tuple *t* and returns a

simple constraint. It is defined recursively as follows:

$$simp(\lor((A = c_1) \triangleright \phi_1, (A = c_2) \triangleright \phi_2, \ldots), t) := \phi_k \text{ if } t.A = c_k$$
$$simp(\land(\psi_{A_1}, \psi_{A_2}, \ldots), t) := \land(simp(\psi_{A_1}, t), simp(\psi_{A_2}, t), \ldots)$$

If the condition in the definition above does not hold for any c_k , then simp (ψ, t) is undefined and simp $(\wedge(\dots, \psi, \dots), t)$ is also undefined. If $simp(\psi, t)$ is undefined, then $[[\psi]](t) := 1$. When $simp(\psi, t)$ is defined, the quantitative semantics of ψ is given by:

$$[[\psi]](t) := [[simp(\psi, t)]](t)$$

Since compound constraints simplify to simple constraints, we mostly focus on simple constraints. Even there, we pay special attention to bounded-projection constraints (ϕ) of the form $1b \leq F(\vec{A}) \leq ub$, which lie at the core of simple constraints.

EXAMPLE 4. Consider the constraint ϕ_1 from Example 3. For $t \in D$, $[[\phi_1]](t) = 0$ since ϕ_1 is satisfied by all tuples in D. The standard deviation of the projection F over D, $\sigma(F(D)) = \sigma(\{0, -5, 5, -2\}) = 3.6$. *Now consider the last tuple* $t_5 \notin D$. $F(t_5) = (370 - 1350) - 458 =$ -1438, which is way below the lower bound -5 of ϕ_1 . Now we compute how much t_5 violates $\phi_1: [[\phi_1]](t_5) = [[-5 \le F(\vec{A}) \le 5]](t_5) = \eta(\alpha \cdot$ $\max(0, -1438 - 5, -5 + 1438)) = 1 - e^{-\frac{1433}{3.6}} \approx 1$. Intuitively, this implies that t_5 strongly violates ϕ_1 .

4 **CONFORMANCE CONSTRAINT SYNTHESIS**

In this section, we describe our techniques for deriving conformance constraints. We start with the synthesis of simple constraints (the ϕ constraints in our language specification), followed by compound constraints (the Ψ constraints in our language specification). Finally, we analyze the time and memory complexity of our algorithm.

4.1 Simple Conformance Constraints

Synthesizing simple conformance constraints involves (a) discovering the projections, and (b) discovering the lower and upper bounds for each projection. We start by discussing (b), followed by the principle to identify effective projections, based on which we solve (a).

4.1.1 Synthesizing Bounds for Projections. Fix a projection F and consider the bounded-projection constraint ϕ : $lb \leq F(A) \leq ub$. Given a dataset *D*, a trivial choice for the bounds is: lb = min(F(D))and ub = max(F(D)). However, this choice is very sensitive to noise: adding a single atypical tuple to D can produce very different constraints. Instead, we use a more robust choice as follows:

$$lb = \mu(F(D)) - C \cdot \sigma(F(D)), \quad ub = \mu(F(D)) + C \cdot \sigma(F(D))$$

Here, $\mu(F(D))$ and $\sigma(F(D))$ denote the mean and standard deviation of the values in F(D), respectively, and C is some positive constant. With these bounds, $[[\phi]](t) = 0$ implies that F(t) is within $C \times \sigma(F(D))$ from the mean $\mu(F(D))$. In our experiments, we set C = 4, which ensures that in expectation, very few tuples in D will violate the constraint for many distributions of the values in F(D). Specifically, if F(D) follows a normal distribution, then 99.99% of the population is expected to lie within 4 standard deviations from mean. Note that we make no assumption on the original data distribution of each attribute.

¹For a target value y, predicted value \hat{y} , and a parameter ϵ , the ϵ -insensitive loss is 0 if $|y - \hat{y}| < \epsilon$ and $|y - \hat{y}| - \epsilon$ otherwise.



Figure 3: Clear and shaded regions depict conformance and nonconformance zones, respectively. (a) Correlated projections X and Yyield conformance constraints forming a large conformance zone, (b) Uncorrelated (orthogonal) projections X - Y and X + Y yield conformance constraints forming a smaller conformance zone.

Setting the bounds 1b and ub as $C \cdot \sigma(F(D))$ -away from the mean, and the scaling factor α as $\frac{1}{\sigma(F(D))}$, guarantees the following property for our quantitative semantics:

LEMMA 5. Let D be a dataset and let ϕ_k be $|\mathbf{b}_k| \leq F_k(\vec{A}) \leq \mathbf{u}_k$ for k = 1, 2. Then, for any tuple t, if $\frac{|F_1(t) - \mu(F_1(D))|}{\sigma(F_1(D))} \geq \frac{|F_2(t) - \mu(F_2(D))|}{\sigma(F_2(D))}$, then $[[\phi_1]](t) \geq [[\phi_2]](t)$.

This means that larger deviation from the mean (proportionally to the standard deviation) results in higher degree of violation under our semantics. The proof follows from the fact that the normalization function $\eta(.)$ is monotonically increasing, and hence, $[[\phi_k]](t)$ is a monotonically non-decreasing function of $\frac{|F_k(t)-\mu(F_k(D))|}{\sigma(F_k(D))}$.

4.1.2 Principle for Synthesizing Projections. We start by investigating what makes a constraint more effective than others. An effective constraint (1) should not overfit the data, but rather generalize by capturing the properties of the data, and (2) should not underfit the data, because it would be too permissive and fail to identify deviations effectively. Our flexible bounds (Section 4.1.1) serve to avoid overfitting. In this section, we focus on identifying the principles that help us avoid underfitting. We first describe the key technical ideas for characterizing effective projections through example and then proceed to formalization.

EXAMPLE 6. Let D be a dataset of three tuples {(1,1.1),(2,1.7),(3,3.2)} with two attributes X and Y. Consider two arbitrary projections: X and Y. For X: $\mu(X(D)) = 2$ and $\sigma(X(D)) = 0.8$. So, bounds for its conformance constraint are: $lb = 2-4\times0.8 = -1.2$ and $ub = 2+4\times0.8 = 5.2$. This gives us the conformance constraint: $-1.2 \le X \le 5.2$. Similarly, for Y, we get the conformance constraint: $-1.6 \le Y \le 5.6$. Fig. 3(a) shows the conformance zone (clear region) defined by these two conformance constraints. The shaded region depicts non-conformance zone. The conformance zone is large and too permissive: it allows many atypical tuples with respect to D, such as (0, 4) and (4, 0).

A natural question arises: are there other projections that can better characterize conformance with respect to the tuples in *D*? The answer is yes and next we show another pair of projections that shrink the conformance zone significantly.

EXAMPLE 7. In Fig. 3(b), the clear region is defined by the conformance constraints $-0.8 \le X - Y \le 0.8$ and $-2.8 \le X + Y \le 10.8$, over projections X - Y and X + Y, respectively. The region is indeed much smaller than the one in Fig. 3(a) and allows fewer atypical tuples. How can we derive projection X - Y from the projections X and Y, given D? Note that X and Y are highly correlated in D. In Lemma 11, we show that two highly correlated projections can be linearly combined to construct another projection with lower standard deviation that generates a *stronger* constraint. We proceed to formalize *stronger constraint*—which defines whether a constraint is more effective than another in quantifying violation—and *incongruous tuples*—which help us estimate the subset of the data domain for which a constraint is stronger than the others.

DEFINITION 8 (STRONGER CONSTRAINT). A conformance constraint ϕ_1 is stronger than another conformance constraint ϕ_2 on a subset $H \subseteq \text{Dom}^m$ if $\forall t \in H$, $[[\phi_1]](t) \ge [[\phi_2]](t)$.

Given a dataset $D \subseteq \mathbf{Dom}^m$ and a projection F, for any tuple t, let $\Delta F(t) = F(t) - \mu(F(D))$. For projections F_1 and F_2 , the correlation coefficient ρ_{F_1,F_2} (over D) is defined as $\frac{\frac{|D|}{\sum t \in D} \Delta F_1(t) \Delta F_2(t)}{\sigma(F_1(D))\sigma(F_2(D))}$.

DEFINITION 9 (INCONGRUOUS TUPLE). A tuple t is incongruous w.r.t. a projection pair $\langle F_1, F_2 \rangle$ on D if: $\Delta F_1(t) \cdot \Delta F_2(t) \cdot \rho_{F_1,F_2} < 0$.

Informally, an incongruous tuple for a pair of projections does not follow the general trend of correlation between the projection pair. For example, if F_1 and F_2 are positively correlated ($\rho_{F_1,F_2} > 0$), an incongruous tuple *t* deviates in opposite ways from the mean of each projection ($\Delta F_1(t) \cdot \Delta F_2(t) < 0$). Our goal is to find projections that yield a conformance zone with very few incongruous tuples.

EXAMPLE 10. In Example 6, X and Y are positively correlated with $\rho_{X,Y} \approx 1$. The tuple t = (0, 4) is incongruous w.r.t. $\langle X, Y \rangle$, because $X(t) = 0 < \mu(X(D)) = 2$, whereas $Y(t) = 4 > \mu(Y(D)) = 2$. Intuitively, the incongruous tuples do not behave like the tuples in D when viewed through the projections X and Y. Note that the narrow conformance zone of Fig. 3(b) no longer contains the incongruous tuple (0, 4). In fact, the conformance zone defined by the conformance constraints derived from projections X - Y and X + Y are free from a vast majority of the incongruous tuples.

We proceed to state Lemma 11, which informally says that any two highly correlated projections can be linearly combined to construct a new projection to obtain a stronger constraint. We write ϕ_F to denote the conformance constraint $1b \leq F(\vec{A}) \leq ub$, synthesized from *F*. (All proofs are in our technical report [21].)

LEMMA 11. Let D be a dataset and F_1, F_2 be two projections on D s.t. $|\rho_{F_1,F_2}| \ge \frac{1}{2}$. Then, $\exists \beta_1, \beta_2 \in \mathbb{R}$ s.t. $\beta_1^2 + \beta_2^2 = 1$ and for the new projection $F = \beta_1 F_1 + \beta_2 F_2$:

- (1) $\sigma(F(D)) < \sigma(F_1(D))$ and $\sigma(F(D)) < \sigma(F_2(D))$, and
- (2) φ_F is stronger than both φ_{F1} and φ_{F2} on the set of tuples that are incongruous w.r.t. ⟨F1, F2⟩.

We now extend the result to multiple projections in Theorem 12.

THEOREM 12 (LOW STANDARD DEVIATION CONSTRAINTS). Given a dataset D, let $\mathcal{F}=\{F_1,\ldots,F_K\}$ denote a set of projections on D s.t. $\exists F_i, F_j \in \mathcal{F}$ with $|\rho_{F_i,F_j}| \geq \frac{1}{2}$. Then, there exist a nonempty subset $I \subseteq \{1,\ldots,K\}$ and a projection $F = \sum_{k \in I} \beta_k F_k$, where $\beta_k \in \mathbb{R}$ s.t. (1) $\forall k \in I, \sigma(F(D)) < \sigma(F_k(D))$,

(2) $\forall k \in I, \phi_F$ is stronger than ϕ_{F_k} on the subset H, where $H=\{t \mid \forall k \in I(\beta_k \Delta F_k(t) \ge 0) \lor \forall k \in I(\beta_k \Delta F_k(t) \le 0)\}, and$ (3) $\forall k \notin I, |\rho_{F,F_k}| < \frac{1}{2}.$

Algorithm 1: Procedure to	generate linear	projections
---------------------------	-----------------	-------------

Inputs : A dataset $D \subset \mathbf{Dom}^m$
Output : A set $\{(F_1, \gamma_1), \ldots, (F_K, \gamma_K)\}$ of projections and
importance factors
1 $D_N \leftarrow D$ after dropping non-numerical attributes
$2 D'_N \leftarrow [\vec{1}; D_N]$
3 $\{\vec{w}_1, \ldots, \vec{w}_K\} \leftarrow \text{ eigenvectors of } {D'_N}^T D'_N$
4 foreach $1 \le k \le K$ do
5 $\vec{w}_k' \leftarrow \vec{w}_k$ with first element removed
$ 6 \qquad F_k \leftarrow \lambda \vec{A} : \frac{\vec{A}^T \vec{\mathbf{w}}_k'}{\ \vec{\mathbf{w}}_k' \ } $
7 $\gamma_k \leftarrow \frac{1}{\log(2+\sigma(F_k(D_N)))}$
8 return $\{(F_1, \frac{\gamma_1}{Z}), \ldots, (F_K, \frac{\gamma_K}{Z})\}$, where $Z = \sum_k \gamma_k$

The theorem establishes that to detect violations for tuples in H: (1) projections with low standard deviations define stronger constraints (and, thus, are preferable), and (2) a set of constraints with highly correlated projections is suboptimal (as they can be linearly combined to generate stronger constraints). Note that H is a conservative estimate for the set of tuples where ϕ_F is stronger than each ϕ_{F_E} ; there exist tuples outside H for which ϕ_F is stronger.

Bounded projections vs. convex polytope. Bounded projections (Example 7) relate to the computation of convex polytopes [68]. For example, one can compute a convex hull-the minimal convex polytope that includes all the training tuples-and then any tuple falling outside it is considered non-conforming. However, a convex hull overfits to the training tuples and is extremely sensitive to outliers. For example, consider a training dataset over attributes *X* and *Y*: {(1, 10), (2, 20), (3, 30)}. A convex hull in this case would be a line segment-starting at (1, 10) and ending at (3, 30)-and the tuple (5,50) will fall outside it. Unlike convex hull-whose goal is to find the smallest possible "inclusion zone" that includes all training tuples-our goal is to find a "conformance zone" that reflects the *trend* of the training tuples. This is inspired from the fact that ML models aim to generalize to tuples outside training set; thus, conformance constraints also need to capture trends and avoid overfitting. Our definition of good conformance constraints (low variance and low mutual correlation) balances overfitting and overgeneralization. Therefore, beyond the minimal bounding hyper-box over the training tuples, we take into consideration the distribution of the interaction among attributes (trends). For the above example, conformance constraints will model the interaction trend: Y = 10X, allowing the tuple (5, 50).

4.1.3 PCA-inspired Projection Derivation. Theorem 12 sets the requirements for good projections (see also [43, 46, 69] that make similar observations in different ways). It indicates that we can start with any arbitrary projections and then iteratively improve them. However, we can get the desired set of best projections in one shot using an algorithm inspired by principal component analysis (PCA). PCA relies on computing eigenvectors. There exist different algorithms for computing eigenvectors (from the infinite space of possible vectors). The general mechanism involves applying numerical approaches to iteratively converge to the eigenvectors (up to a desired precision) as no analytical solution exists in general. Algorithm 1 returns projections that correspond to the principal components of a slightly modified version of the given dataset: **Line 1** Drop all non-numerical attributes from *D* to get the numeric dataset D_N . This is necessary because PCA only applies to numerical values. Instead of dropping, one can also consider embedding techniques to convert non-numerical attributes to numerical ones. **Line 2** Add a new column to D_N that consists of the constant 1, to obtain the modified dataset $D'_N := [\vec{1}; D_N]$, where $\vec{1}$ denotes the column vector with 1 everywhere. We do this transformation to capture the additive constant within principal components, which ensures that the approach works even for unnormalized data.

Line 3 Compute *K* eigenvectors of the square matrix $D'_N{}^T D'_N$, where *K* denotes the number of columns in D'_N . These eigenvectors provide coefficients to construct projections.

Lines 5–6 Remove the first element (coefficient for the newly added constant column) of all eigenvectors and normalize them to generate projections. Note that we no longer need the constant element of the eigenvectors since we can appropriately adjust the bounds, 1b and ub, for each projection by evaluating it on D_N . **Line 7** Compute importance factor for each projection. Since pro-

jections with smaller standard deviations are more discerning (i.e., stronger), we assign each projection an importance factor (γ) that is inversely proportional to its standard deviation over D_N .

Line 8 Return the linear projections with corresponding normalized importance factors.

We now claim that the projections returned by Algorithm 1 include the projection with minimum standard deviation and the correlation between any two projections is 0. This indicates that we cannot further improve the projections, and, thus they are optimal.

THEOREM 13 (CORRECTNESS OF ALGORITHM 1). Given a numerical dataset D over the schema \mathcal{R} , let $\mathcal{F} = \{F_1, F_2, \ldots, F_K\}$ be the set of linear projections returned by Algorithm 1. Let $\sigma^* = \min_k^K \sigma(F_k(D))$. If $\mu(A_k(D)) = 0$ for all attribute A_k in \mathcal{R} , then,² (1) $\sigma^* \leq \sigma(F(D)) \forall F = \vec{A}^T \vec{w}$ where $||\vec{w}|| \geq 1$, and (2) $\forall F_j, F_k \in \mathcal{F}$ s.t. $F_j \neq F_k, \rho_{F_j, F_k} = 0$.

Using projections F_1, \ldots, F_K , and importance factors $\gamma_1, \ldots, \gamma_K$, returned by Algorithm 1, we generate the simple (conjunctive) constraint with *K* conjuncts: $\bigwedge_k \mathbf{lb}_k \leq F_k(\vec{A}) \leq \mathbf{ub}_k$. We compute the bounds \mathbf{lb}_k and \mathbf{ub}_k following Section 4.1.1 and use the importance factor γ_k for the k^{th} conjunct in the quantitative semantics.

EXAMPLE 14. Algorithm 1 finds the projection of the conformance constraint of Example 1, but in a different form. The actual airlines dataset has an attribute distance (DIS) that represents miles travelled by a flight. In our experiments, we found the following conformance constraint³ over the dataset of daytime flights:

$$0.7 \times \text{AT} - 0.7 \times \text{DT} - 0.14 \times \text{DUR} - 0.07 \times \text{DIS} \approx 0$$
(1)

This constraint is not quite interpretable by itself, but it is in fact a linear combination of two expected and interpretable constraints: ⁴

$$AT - DT - DUR \approx 0 \tag{2}$$

$$DUR - 0.12 \times DIS \approx 0$$
 (3)

²When the condition $\forall A_k \ \mu(A_k(D)) = 0$ does not hold, slightly modified variants of the claim hold. However, by normalizing *D* (i.e., by subtracting attribute mean $\mu(A_k(D))$ from each $A_k(D)$), it is always possible to satisfy the condition.

³For ease of exposition, we use $F(\vec{A}) \approx 0$ to denote $\epsilon_1 \leq F(\vec{A}) \leq \epsilon_2$, where $\epsilon_i \approx 0$. ⁴We developed a tool [20] to explain causes of non-conformance. [21]

Here, (2) is the one mentioned in Example 1 and (3) follows from the fact that average aircraft speed is about 500 mph implying that it requires 0.12 minutes per mile. $0.7 \times (2) + 0.56 \times (3)$ yields:

$$0.7 \times (\text{AT} - \text{DT} - \text{DUR}) + 0.56 \times \text{DUR} - 0.56 \times 0.12 \times \text{DIS} \approx 0$$
$$\Rightarrow 0.7 \times \text{AT} - 0.7 \times \text{DT} - 0.14 \times \text{DUR} - 0.07 \times \text{DIS} \approx 0$$

=

Which is exactly the conformance constraint (1). Algorithm 1 found the optimal projection of (1), which is a linear combination of the projections of (2) and (3). The reason is: there is a correlation between the projections of (2) and (3) over the dataset (Theorem 12). One possible explanation of this correlation is: whenever there is an error in the reported duration of a tuple, it violates both (2) and (3). Due to this natural correlation, Algorithm 1 returned the optimal projection of (1), that "covers" both projections of (2) or (3).

4.2 Compound Conformance Constraints

The quality of our PCA-based simple linear constraints relies on how many low variance linear projections we are able to find on the given dataset. For many datasets, it is possible we find very few, or even none, such linear projections. In these cases, it is fruitful to search for compound constraints; we first focus on *disjunctive constraints* (defined by ψ_A in our language grammar).

The PCA-based approach fails in cases where there exist different piecewise linear trends within the data, as it will result into low-quality constraints, with very high variances. In such cases, partitioning the dataset and then learning constraints separately on each partition will result in significant improvement of the learned constraints. A disjunctive constraint is a compound constraint of the form $\bigvee_k ((A = c_k) \triangleright \phi_k)$, where each ϕ_k is a constraint for a specific partition of D. Finding disjunctive constraints involves horizontally partitioning the dataset D into smaller disjoint datasets D_1, D_2, \ldots, D_L . Our strategy for partitioning D is to use categorical attributes with a small domain in D; in our implementation, we use those attributes A_i for which $|\{t.A_i | t \in D\}| \leq 50$. If A_i is such an attribute with values v_1, v_2, \ldots, v_L , we partition D into L disjoint datasets D_1, D_2, \dots, D_L , where $D_l = \{t \in D | t.A_j = v_l\}$. Let $\phi_1, \phi_2, \ldots, \phi_L$ be the *L* simple conformance constraints we learn for D_1, D_2, \ldots, D_L using Algorithm 1, respectively. We compute the following disjunctive conformance constraint for D:

 $((A_j = v_1) \triangleright \phi_1) \lor ((A_j = v_2) \triangleright \phi_2) \lor \dots \lor ((A_j = v_L) \triangleright \phi_L)$

We repeat this process and partition D across multiple attributes and generate a compound disjunctive constraint for each attribute. Then we generate the final compound conjunctive conformance constraint (Ψ) for D, which is the conjunction of all these disjunctive constraints. Intuitively, this final conformance constraint forms a set of *overlapping* hyper-boxes around the data tuples.

4.3 Theoretical Analysis

4.3.1 Runtime Complexity. Computing simple constraints involves two computational steps: (1) computing $X^T X$, where X is an $n \times m$ matrix with *n* tuples and *m* attributes, which takes $O(nm^2)$ time, and (2) computing the eigenvalues and eigenvectors of an $m \times m$ positive definite matrix, which has complexity $O(m^3)$ [47]. Once we obtain the linear projections using the above two steps, we need to compute the mean and variance of these projections on the original dataset, which takes $O(nm^2)$ time. In summary, the overall procedure is cubic in the number of attributes and linear in the number of tuples. For computing disjunctive constraints, we greedily pick attributes that take at most L (typically small) distinct values, and then run the above procedure for simple constraints at most L times. This adds just a constant factor overhead per attribute.

4.3.2 Memory Complexity. The procedure can be implemented in $O(m^2)$ space. The key observation is that $X^T X$ can be computed as $\sum_{i=1}^{n} t_i t_i^T$, where t_i is the i^{th} tuple in the dataset. Thus, $X^T X$ can be computed incrementally by loading only one tuple at a time into memory, computing $t_i t_i^T$, and then adding that to a running sum, which can be stored in $O(m^2)$ space. Note that instead of such an incremental computation, this can also be done in an embarrassingly parallel way where we horizontally partition the data (row-wise) and each partition is computed in parallel.

4.3.3 Implication, Redundancy, and Minimality. Definition 8 gives us the notion of implication on conformance constraints: for a dataset *D*, satisfying ϕ_1 that is stronger than ϕ_2 implies that *D* would satisfy ϕ_2 as well. Lemma 11 and Theorem 12 associate *re*dundancy with correlation: correlated projections can be combined to construct a new projection that makes the correlated projections redundant. Theorem 13 shows that our PCA-based procedure finds a non-redundant (orthogonal and uncorrelated) set of projections. For disjunctive constraints, it is possible to observe redundancy across partitions. However, our quantitative semantics ensures that redundancy does not affect the violation score. Another notion relevant to data profiles (e.g., FDs) is minimality. In this work, we do not focus on finding the minimal set of conformance constraints. Towards achieving minimality for conformance constraints, a future direction is to explore techniques for optimal data partitioning. However, our approach computes only *m* conformance constraints for each partition. Further, for a single tuple, only $m_N \cdot m_C$ conformance constraints are applicable, where m_N and m_C are the number of numerical and categorical attributes in D (i.e., $m = m_N + m_C$). The quantity $m_N \cdot m_C$ is upper-bounded by $\frac{m^2}{4}$.

5 TRUSTED MACHINE LEARNING

In this section, we provide a theoretical justification of why conformance constraints are effective in identifying tuples for which learned models are likely to make incorrect predictions. To that end, we define *unsafe* tuples and show that an "ideal" conformance constraint provides a sound and complete mechanism to detect unsafe tuples. In Section 4, we showed that low-variance projections construct strong conformance constraints. We now make a similar argument, but in a slightly different way: we show that projections with zero variance give us equality constraints that are useful for trusted machine learning. We start with an example to provide the intuition.

EXAMPLE 15. Consider the airlines dataset D and assume that all tuples in D satisfy the equality constraint $\phi := AT - DT - DUR = 0$ (i.e., lb = ub = 0). Note that for equality constraint, the corresponding projection has zero variance—the lowest possible variance. Now, suppose that the task is to learn some function f(AT, DT, DUR). If the above constraint holds for D, then the ML model can instead learn the function g(AT, DT, DUR) = f(DT + DUR, DT, DUR). g will perform just as well as f on D: in fact, it will produce the same output as f on D. If a new serving tuple t satisfies ϕ , then g(t) = f(t), and the prediction

will be correct. However, if t does not satisfy ϕ , then g(t) will likely be significantly different from f(t). Hence, violation of the conformance constraint is a strong indicator of performance degradation of the learned prediction model. Note that f need not be a linear function: as long as g is also in the class of models that the learning procedure is searching over, the above argument holds.

We proceed to formally define unsafe tuples. We use [D; Y] to denote the *annotated dataset* obtained by appending the target attribute *Y* to a dataset *D*, and coDom to denote *Y*'s domain.

DEFINITION 16 (UNSAFE TUPLE). Given a class C of functions with signature $\mathbf{Dom}^m \mapsto \mathsf{coDom}$, and an annotated dataset $[D; Y] \subset (\mathbf{Dom}^m \times \mathsf{coDom})$, a tuple $t \in \mathbf{Dom}^m$ is unsafe w.r.t. C and [D; Y], $if \exists f, g \in C$ s.t. f(D) = g(D) = Y but $f(t) \neq g(t)$.

Intuitively, *t* is unsafe if there exist two different predictor functions *f* and *g* that agree on all tuples in *D*, but disagree on *t*. Since, we can never be sure whether the model learned *f* or *g*, we should be cautious about the prediction on *t*. Example 15 suggests that *t* can be unsafe when all tuples in *D* satisfy the equality conformance constraint $f(\vec{A}) - g(\vec{A}) = 0$ but *t* does not. Hence, we can use the following approach for trusted machine learning:

(1) Learn conformance constraints Φ for the dataset *D*.

(2) Declare *t* as unsafe if *t* does not satisfy Φ .

The above approach is sound and complete for characterizing unsafe tuples, thanks to the following proposition.

PROPOSITION 17. There exists a conformance constraint Φ for D s.t. the following statement is true: " $\neg \Phi(t)$ iff t is unsafe w.r.t. C and [D; Y] for all $t \in \mathbf{Dom}^m$ ".

The required conformance constraint Φ is: $\forall f, g \in C : f(D) = g(D) = Y \Rightarrow f(\vec{A}) - g(\vec{A}) = 0$. Intuitively, when all possible pairs of functions that agree on *D* also agree on *t*, only then the prediction on *t* can be trusted. (More discussion is in our technical report [21].)

5.1 Applicability

Generalization to noisy settings. While our analysis and formalization for using conformance constraints for TML focused on the noise-free setting, the intuition generalizes to noisy data. Specifically, suppose that f and g are two possible functions a model may learn over D; then, we expect that the difference f - g will have small variance over D, and, thus, would be a good conformance constraint. In turn, the violation of this constraint would mean that f and g diverge on a tuple t (making t unsafe); since we are oblivious of the function the model learned, prediction on t is untrustworthy.

False positives. Conformance constraints are designed to work in a model-agnostic setting. Although this setting is of great practical importance, designing a perfect mechanism for quantifying trust in ML model predictions, while remaining completely model-agnostic, is challenging. It raises the concern of *false positives*: conformance constraints may incorrectly flag tuples for which the model's prediction is in fact correct. This may happen when the model ignores the trend that conformance constraints learn. Since we are oblivious of the prediction task and the model, it is preferable that conformance constraints behave rather *conservatively* so that the users can be cautious about potentially unsafe tuples. Moreover, if a model ignores some attributes (or their interactions) during training, it is still necessary to learn conformance constraints over them. Particularly, in case of concept drift [66], the ground truth may start depending on those attributes, and by learning conformance constraints over all attributes, we can better detect potential model failures.

False negatives. Another concern involving conformance constraints is of *false negatives*: linear conformance constraints may miss nonlinear constraints, and, thus, fail to identify some unsafe tuples. However, the linear dependencies modeled in conformance constraints persist even after sophisticated (nonlinear) attribute transformations. Therefore, violation of conformance constraints is a strong indicator of potential failure of a possibly nonlinear model.

Modeling nonlinear constraints. While linear conformance constraints are the most common ones, we note that our framework can be easily extended to support nonlinear conformance constraints using *kernel functions* [57]—which offer an efficient, scalable, and powerful mechanism to learn nonlinear decision boundaries for support vector machines (also known as *kernel trick*). Briefly, instead of explicitly augmenting the dataset with transformed nonlinear attributes—which grows exponentially with the desired degree of polynomials—kernel functions enable *implicit* search for nonlinear models. The same idea also applies for PCA called kernel-PCA [5, 34]. While we limit our evaluation to only linear kernel, polynomial kernels—e.g., radial basis function (RBF) [38]—can be plugged into our framework to model nonlinear conformance constraints.

In general, our conformance language is not guaranteed to model all possible functions that an ML model can potentially learn, and, thus, is not guaranteed to find the best conformance constraint. However, our empirical evaluation on real-world datasets shows that our language models conformance constraints effectively.

6 EXPERIMENTAL EVALUATION

We now present experimental evaluation to demonstrate the effectiveness of conformance constraints over our two case-study applications (Section 2): trusted machine learning and data drift. Our experiments target the following research questions:

- How effective are conformance constraints for trusted machine learning? Is there a relationship between constraint violation score and the ML model's prediction accuracy? (Section 6.1)
- Can conformance constraints be used to quantify data drift? How do they compare to other state-of-the-art drift-detection techniques? (Section 6.2)

Efficiency. In all our experiments, our algorithms for deriving conformance constraints were extremely fast, and took only a few seconds even for datasets with 6 million rows. The number of attributes were reasonably small (~40), which is true for most practical applications. As our theoretical analysis showed (Section 4.3), our approach is linear in the number of data rows and cubic in the number of attributes. Since the runtime performance of our techniques is straightforward, we opted to not include further discussion of efficiency here and instead focus this empirical analysis on the techniques' effectiveness.

Implementation: CCSYNTH. We created an implementation of conformance constraints and our method for synthesizing them, CCSYNTH, in Python 3 [11]. Experiments were run on a Windows 10 machine (3.60 GHz processor and 16GB RAM).

	Train	Serving					
	114111	Daytime	Overnight	Mixed			
Average violation MAE	0.02% 18.95	0.02% 18.89	27.68% 80.54	8.87% 38.60			

Figure 4: Average constraint violation (in percentage) and MAE (for linear regression) of four data splits on the airlines dataset. The constraints were learned on Train, excluding the target attribute, delay.



Figure 5: Constraint violation strongly correlates with the absolute error of delay prediction of a linear regression model.

Datasets

Airlines [4] contains data about flights and has 14 attributes —year, month, day, day of week, departure time, arrival time, carrier, flight number, elapsed time, origin, destination, distance, diverted, and arrival delay. We used a subset of the data containing all flight information for year 2008. In this dataset, most of the attributes follow uniform distribution (e.g., month, day, arrival and departure time, etc.); elapsed time and distance follow skewed distribution with higher concentration towards small values (implying that shorter flights are more common); arrival delay follows a slightly skewed gaussian distribution implying most flights are on-time, few arrive late and even fewer arrive early. The training and serving datasets contain 5.4M and 0.4M rows, respectively.

Human Activity Recognition (HAR) [63] is a real-world dataset about physical activities for 15 individuals, 8 males and 7 females, with varying fitness levels and BMIs. We use data from two sensors accelerometer and gyroscope—attached to 6 body locations—head, shin, thigh, upper arm, waist, and chest. We consider 5 activities lying down, running, sitting, standing, and walking. The dataset contains 36 numerical attributes (2 sensors × 6 body-locations × 3 co-ordinates) and 2 categorical attributes—activity-type and person-ID. We pre-processed the dataset to aggregate the measurements over a small time window, resulting in 10,000 tuples per person and activity, for a total of 750,000 tuples.

Extreme Verification Latency (EVL) [62] is a widely used benchmark to evaluate drift-detection algorithms in non-stationary environments under extreme verification latency. It contains 16 synthetic datasets with incremental and gradual concept drifts over time. The number of attributes of these datasets vary from 2 to 6 and each of them has one categorical attribute.

6.1 Trusted Machine Learning

We now demonstrate the applicability of conformance constraints in the TML problem. We show that, serving tuples that violate the training data's conformance constraints are unsafe, and therefore, an ML model is more likely to perform poorly on those tuples.

Airlines. We design a regression task of predicting the arrival delay and train a linear regression model for the task. Our goal

is to observe whether the mean absolute error of the predictions (positively) correlates to the constraint violation for the serving tuples. In a process analogous to the one described in Example 1, our training dataset (Train) comprises of a subset of daytime flights flights that have arrival time later than the departure time (in 24 hour format). We design three serving sets: (1) Daytime: similar to Train, but another subset, (2) Overnight: flights that have arrival time earlier than the departure time (the dataset does not explicitly report the date of arrival), and (3) Mixed: a mixture of Daytime and Overnight. A few sample tuples of this dataset are in Fig. 1.

Our experiment involves the following steps: (1) CCSYNTH computes conformance constraints Φ over Train, while *ignoring* the target attribute delay. (2) We compute average constraint violation for all four datasets—Train, Daytime, Overnight, and Mixed—against Φ (first row of Fig. 4). (3) We train a linear regression model over Train—including delay—that learns to predict arrival delay. (4) We compute mean absolute error (MAE) of the prediction accuracy of the regressor over the four datasets (second row of Fig. 4). We find that constraint violation is a very good proxy for prediction error, as they vary in a similar manner across the four datasets. The reason is that the model implicitly assumes that the constraints (e.g., $AT - DT - DUR \approx 0$) derived by CCSYNTH will always hold, and, thus, deteriorates when the assumption no longer holds.

To observe the rate of false positives and false negatives, we investigate the relationship between constraint violation and prediction error at tuple-level granularity. We sample 1000 tuples from Mixed and organize them by decreasing order of violations (Fig. 5). For all the tuples (on the left) that incur high constraint violations, the regression model incurs high error for them as well. This implies that CCSYNTH reports no false positives. There are some false negatives (right part of the graph), where violation is low, but the prediction error is high. Nevertheless, such false negatives are very few.

HAR. We design a supervised classification task to identify persons from their activity data that contains 36 numerical attributes. We construct train_x with data for sedentary activities (lying, standing, and sitting), and train_y with the corresponding person-IDs. We learn conformance constraints on train_x, and train a Logistic Regression (LR) classifier using the annotated dataset [train_x; train_y]. During serving, we mix mobile activities (walking and running) with held-out data for sedentary activities and observe how the classification's mean accuracy-drop (i.e., how much the mean prediction accuracy decreases compared to the mean prediction accuracy over the training data) relates to average constraint violation. To avoid artifacts due to sampling bias, we repeat this experiment 10 times for different subsets of the data by randomly sampling 5000 data points for each of training and serving. Fig. 6(a) depicts our findings: classification degradation has a clear positive correlation with violation (pcc = 0.99 with p-value = 0).

Noise sensitivity. Intuitively, noise weakens conformance constraints by increasing variance in the training data, which results in reduced violations of the serving data. However, this is desirable: as more noise makes machine-learned models less likely to overfit, and, thus, more robust. In our experiment for observing noise sensitivity of conformance constraints, we use mobile activity data as the serving set and start with sedentary data as the training set. Then we gradually introduce noise in the training set by mixing mobile



Figure 6: (a) As a higher fraction of mobile activity data is mixed with sedentary activity data, conformance constraints are violated more, and the classifier's mean accuracy-drop increases. (b) As more noise is added during training, conformance constraints get weaker, leading to less violation and decreased accuracy-drop. (c) CCSVNTH detects the gradual local drift on the HAR dataset as more people start changing their activities. In contrast, weighted-PCA (W-PCA) fails to detect drift in absence of a strong global drift.



Figure 7: Inter-person constraint violation heat map. Each person has a very low self-violation.

activity data. As Fig. 6(b) shows, when more noise is added to the training data, conformance constraints start getting weaker; this leads to reduction in violations. However, the classifier also gains robustness with more noise, which is evident from gradual decrease in accuracy-drop (i.e., increase in accuracy). Therefore, even under the presence of noise, the positive correlation between classification degradation and violation persists (pcc = 0.82 with p-value = 0.002).

Key takeaway: CCSYNTH derives conformance constraints whose violation is a strong proxy of model prediction accuracy. Their correlation persists even in the presence of noise.

6.2 Data Drift

We now present results of using conformance constraints for driftdetection; specifically, for *quantifying* drift in data. Given a baseline dataset *D*, and a new dataset *D'*, we measure the drift as average violation of tuples in *D'* on conformance constraints learned for *D*. **HAR.** We perform two drift-quantification experiments on HAR: *Gradual drift*. For observing how CCSYNTH detects gradual drift, we introduce drift in an organic way. The initial training dataset contains data of exactly one activity for each person. This is a realistic scenario as one can think of it as taking a snapshot of what a group of people are doing during a reasonably small time window. We introduce gradual drift to the initial dataset by altering the activity of one person at a time. To control the amount of drift, we use a parameter K. When K = 1, the first person switches their activity, i.e., we replace the tuples corresponding to the first person performing activity A with new tuples that correspond to the same person performing another activity B. When K = 2, the second person switches their activity in a similar fashion, and so on. As we increase K from 1 to 15, we expect a gradual increase in the drift magnitude compared to the initial training data. When K = 15, all persons have switched their activities from the initial setting, and we expect to observe maximum drift. We repeat this experiment 10 times, and display the average constraint violation in Fig. 6(c): the drift magnitude (violation) indeed increases as more people alter their activities.

The baseline weighted-PCA approach (W-PCA) fails to model local constraints (who is doing what), and learns some weaker global constraints indicating that "a group of people are performing some activities". Thus, it fails to detect the gradual local drift. CCSYNTH can detect drift when individuals switch activities, as it learns *disjunctive* constraints that encode who is doing what.

Inter-person drift. The goal of this experiment is to observe how effectively conformance constraints can model the representation of an entity and whether such learned representations can be used to accurately quantify drift between two entities. We use half of each person's data to learn the constraints, and compute violation on the held-out data. CCSYNTH learns disjunctive constraints for each person over all activities, and then we use the violation w.r.t. the learned constraints to measure how much the other persons drift. While computing drift between two persons, we compute activity-wise constraint violation scores and then average them out. In Fig. 7, the violation score at row p1 and column p2 denotes how much p2 drifts from p1. As one would expect, we observe a very low self-drift across the diagonal. Interestingly, our result also shows that some people are more different from others, which appears to have some correlation with (the hidden ground truth) fitness and BMI values. This asserts that the constraints we learn for each person are an accurate abstraction of that person's activities, as people do not deviate too much from their usual activity patterns.

EVL. We now compare CCSYNTH against other state-of-the-art drift detection approaches on the EVL benchmark.

Baselines. We use two drift-detection baselines as described below:

(1) PCA-SPLL [43] similar to us, also argues that principal components with lower variance are more sensitive to a general drift, and uses those for dimensionality reduction. It then models multivariate distribution over the reduced dimensions and applies semiparametric log-likelihood (SPLL) to detect drift between two multivariate distributions. However, PCA-SPLL discards all high-variance principal components and does not model disjunctive constraints.

(2) CD (Change Detection) [51] is another PCA-based approach for drift detection in data streams. But unlike PCA-SPLL, it ignores low-variance principal components. CD projects the data onto top k high-variance principal components, which results into multiple univariate distributions. We compare against two variants of CD: CD-Area, which uses the intersection area under the curves of two density functions as a divergence metric, and CD-MKL, which uses Maximum KL-divergence as a symmetric divergence metric, to compute divergence between the univariate distributions.



Figure 8: In the EVL benchmark, CCSYNTH quantifies drift correctly for all cases, outperforming other approaches. PCA-SPLL fails to detect drift in a few cases by discarding all principal components; CD-MKL and CD-Area are too sensitive to small drift and detect spurious drifts.

Fig. 8 depicts how CCSYNTH compares against CD-MKL, CD-Area, and PCA-SPLL, on 16 datasets in the EVL benchmark. For PCA-SPLL, we retain principal components that contribute to a cumulative explained variance below 25%. Beyond drift detection, which just detects if drift is above some threshold, we focus on drift quantification. A tuple (x, y) in the plots denotes that drift magnitude for dataset at x^{th} time window, w.r.t. the dataset at the first time window, is y. Since different approaches report drift magnitudes in different scales, we normalize the drift values within [0, 1]. Additionally, since different datasets have different number of time windows, for the ease of exposition, we normalize the time window indices. Below we state our key findings from this experiment:

CCSYNTH's drift quantification matches the ground truth. In all of the datasets in the EVL benchmark, CCSYNTH is able to correctly quantify the drift, which matches the ground truth [17] exceptionally well. In contrast, as CD focuses on detecting the drift point, it is ill-equipped to precisely quantify the drift, which is demonstrated in several cases (e.g., 2CHT), where CD fails to distinguish the deviation in drift magnitudes. In contrast, both PCA-SPLL and CCSYNTH correctly quantify the drift. Since CD only retains highvariance principal components, it is more susceptible to noise and considers noise in the dataset as significant drift, which leads to incorrect drift quantification. In contrast, PCA-SPLL and CCSYNTH ignore the noise and only capture the general notion of drift.

CCSYNTH models local drift. When the dataset contains instances from multiple classes, the drift may be just local, and not global (e.g., 4CR dataset [21]). In such cases, PCA-SPLL fails to detect drift (4CR, 4CRE-V2, and FG-2C-2D). In contrast, CCSYNTH learns disjunctive constraints and quantifies local drifts accurately.

Key takeaways: CCSYNTH can effectively detect data drift, both global and local, is robust across drift patterns, and significantly outperforms the state-of-the-art methods.

7 RELATED WORK

There is extensive literature on data-profiling [1] primitives that model relationships among data attributes, such as functional dependencies (FD) [48, 75] and their variants [10, 19, 30, 32, 40, 42, 50], differential dependencies [60], denial constraints [8, 12, 44, 49], statistical constraints [74], etc. However, none of them focus on learning approximate arithmetic relationships that involve multiple numerical attributes in a noisy setting, which is the focus of our work. Some FD variants [10, 30, 32, 40, 42] consider noisy setting, but they require noise parameters to be explicitly specified by the user. In contrast, we do not require any explicit noise parameter.

The issue of trust, resilience, and interpretability of artificial intelligence (AI) systems has been a theme of increasing interest recently [33, 35, 55, 71], particularly for safety-critical data-driven AI systems [65, 72]. A standard way to decide whether to trust a classifier or not, is to use the classifier-produced confidence score. However, this is not always effective since the classifier's confidence scores are not well-calibrated [34]. While some recent techniques [15, 26, 34, 56] aim at validating the inferences made by machine-learned models on unseen tuples, they usually require knowledge of the inference task, access to the model, and/or expected cases of data shift, which we do not. Furthermore, they usually require costly hyper-parameter tuning and do not generate closed-form data profiles like conformance constraints (Fig. 2). Prior work on data drift, change detection, and covariate shift [3, 7, 9, 13, 14, 16, 18, 22, 23, 28, 29, 31, 37, 39, 54, 58, 59, 61, 67, 73] relies on modeling data distribution. However, data distribution does not capture constraints, which is the primary focus of our work.

Few works [15, 26, 45] use autoencoder's [27, 53] input reconstruction error to determine if a new data point is out-of-distribution. Our approach is similar to outlier-detection [41] and one-classclassification [64]. However, conformance constraints differ from these approaches as they perform under the additional requirement to generalize the data in a way that is exploited by a given class of ML models. In general, there is a clear gap between representation learning (that models data likelihood) [2, 27, 36, 53] and the (constraint-oriented) data-profiling techniques to address the problem of trusted AI and our aim is to bridge this gap.

8 SUMMARY AND FUTURE DIRECTIONS

We introduced conformance constraints, and the notion of unsafe tuples for trusted machine learning. We presented an efficient and scalable approach for synthesizing conformance constraints and empirically demonstrated their effectiveness in tagging unsafe tuples and quantify data drift. We studied two use-cases from a large pool of potential applications using linear conformance constraints. In future, we want to explore more powerful nonlinear conformance constraints using autoencoders. Moreover, we plan to explore approaches to learn conformance constraints in a decision-tree-like structure where categorical attributes will guide the splitting conditions and leaves will contain simple conformance constraints.

Acknowledgements: This work was partially supported by the NSF grants IIS-1453543 and CCF-1763423, and by Oracle Labs.

REFERENCES

- Ziawasch Abedjan, Lukasz Golab, and Felix Naumann. 2015. Profiling relational data: a survey. VLDB J. 24, 4 (2015), 557–581.
- [2] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas J. Guibas. 2018. Learning Representations and Generative Models for 3D Point Clouds. In *ICML*, Vol. 80. 40–49.
- [3] Charu C. Aggarwal. 2003. A Framework for Diagnosing Changes in Evolving Data Streams. In SIGMOD. 575–586.
- [4] Airlines Dataset 2009. http://kt.ijs.si/elena_ikonomovska/data.html.
- [5] Carlos Alzate and Johan AK Suykens. 2008. Kernel component analysis using an epsilon-insensitive robust loss function. *IEEE Transactions on Neural Networks* 19, 9 (2008), 1583–1598.
- [6] Jean Paul Barddal, Heitor Murilo Gomes, Fabrício Enembreck, and Bernhard Pfahringer. 2017. A survey on feature drift adaptation: Definition, benchmark, challenges and future directions. *Journal of Systems and Software* 127 (2017), 278–294.
- [7] Albert Bifet and Ricard Gavaldà. 2007. Learning from Time-Changing Data with Adaptive Windowing. In SDM. 443–448.
- [8] Tobias Bleifuß, Sebastian Kruse, and Felix Naumann. 2017. Efficient Denial Constraint Discovery with Hydra. PVLDB 11, 3 (2017), 311–323.
- [9] Li Bu, Cesare Alippi, and Dongbin Zhao. 2018. A pdf-Free Change Detection Test Based on Density Difference Estimation. *IEEE Trans. Neural Netw. Learning Syst.* 29, 2 (2018), 324–334.
- [10] Loredana Caruccio, Vincenzo Deufemia, and Giuseppe Polese. 2016. On the discovery of relaxed functional dependencies. In Proceedings of the 20th International Database Engineering & Applications Symposium. 53–61.
- [11] CCSynth Source Code 2021. https://github.com/microsoft/prose/tree/main/misc/ CCSynth.
- [12] Xu Chu, Ihab F. Ilyas, and Paolo Papotti. 2013. Discovering Denial Constraints. PVLDB 6, 13 (2013), 1498–1509.
- [13] Tamraparni Dasu, Shankar Krishnan, Suresh Venkatasubramanian, and Ke Yi. 2006. An information-theoretic approach to detecting changes in multidimensional data streams. In Symp. on the Interface of Statistics, Computing Science, and Applications.
- [14] Rodrigo Fernandes de Mello, Yule Vaz, Carlos Henrique Grossi Ferreira, and Albert Bifet. 2019. On learning guarantees to unsupervised concept drift detection on data streams. *Expert Syst. Appl.* 117 (2019), 90–102.
- [15] Taylor Denouden, Rick Salay, Krzysztof Czarnecki, Vahdat Abdelzad, Buu Phan, and Sachin Vernekar. 2018. Improving Reconstruction Autoencoder Out-ofdistribution Detection with Mahalanobis Distance. *CoRR* abs/1812.02765 (2018). arXiv:1812.02765
- [16] Denis Moreira dos Reis, Peter A. Flach, Stan Matwin, and Gustavo E. A. P. A. Batista. 2016. Fast Unsupervised Online Drift Detection Using Incremental Kolmogorov-Smirnov Test. In SIGKDD. 1545–1554.
- [17] Extreme Verification Latency Benchmark Video (Nonstationary Environments -Archive): 2020. https://sites.google.com/site/nonstationaryarchive/home.
- [18] William J. Faithfull, Juan José Rodríguez Diez, and Ludmila I. Kuncheva. 2019. Combining univariate approaches for ensemble change detection in multivariate data. *Information Fusion* 45 (2019), 202–214.
- [19] Wenfei Fan, Floris Geerts, Laks V. S. Lakshmanan, and Ming Xiong. 2009. Discovering Conditional Functional Dependencies. In *ICDE*. 1231–1234.
- [20] Anna Fariha, Ashish Tiwari, Arjun Radhakrishna, and Sumit Gulwani. 2020. ExTuNe: Explaining Tuple Non-conformance. In SIGMOD. 2741–2744.
- [21] Anna Fariha, Ashish Tiwari, Arjun Radhakrishna, Sumit Gulwani, and Alexandra Meliou. 2020. Conformance Constraint Discovery: Measuring Trust in Data-Driven Systems. *CoRR* abs/2003.01289 (2020).
- [22] Mohamed Medhat Gaber and Philip S Yu. 2006. Classification of changes in evolving data streams using online clustering result deviation. In Proc. Of International Workshop on Knowledge Discovery in Data Streams.
- [23] João Gama, Pedro Medas, Gladys Castillo, and Pedro Pereira Rodrigues. 2004. Learning with Drift Detection. In Advances in Artificial Intelligence - SBIA. 286– 295.
- [24] João Gama, Indre Zliobaite, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. ACM Comput. Surv. 46, 4 (2014), 44:1–44:37.
- [25] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In Proceedings of the 34th International Conference on Machine Learning-Volume 70. 1321–1330.
- [26] Dan Hendrycks and Kevin Gimpel. 2017. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. In ICLR.
- [27] Geoffrey E Hinton and Ruslan R Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *science* 313, 5786 (2006), 504–507.
- [28] Shen-Shyang Ho. 2005. A martingale framework for concept change detection in time-varying data streams. In *ICML*. 321–327.
- [29] Bryan Hooi and Christos Faloutsos. 2019. Branch and Border: Partition-Based Change Detection in Multivariate Time Series. In Proceedings of the 2019 SIAM International Conference on Data Mining, SDM. 504–512.

- [30] Yka Huhtala, Juha Kärkkäinen, Pasi Porkka, and Hannu Toivonen. 1999. TANE: An efficient algorithm for discovering functional and approximate dependencies. *The computer journal* 42, 2 (1999), 100–111.
- [31] Dino Ienco, Albert Bifet, Bernhard Pfahringer, and Pascal Poncelet. 2014. Change detection in categorical evolving data streams. In SAC. 792–797.
- [32] Ihab F. Ilyas, Volker Markl, Peter J. Haas, Paul Brown, and Ashraf Aboulnaga. 2004. CORDS: Automatic Discovery of Correlations and Soft Functional Dependencies. In SIGMOD. 647–658.
- [33] Susmit Jha. 2019. Trust, Resilience and Interpretability of AI Models. In Numerical Software Verification - 12th International Workshop, NSV@CAV. 3–25.
- [34] Heinrich Jiang, Been Kim, Melody Y. Guan, and Maya R. Gupta. 2018. To Trust Or Not To Trust A Classifier. In *NeurIPS*. 5546–5557.
- [35] Daniel Kang, Deepti Raghavan, Peter Bailis, and Matei Zaharia. 2020. Model Assertions for Monitoring and Improving ML Models. In MLSys.
- [36] Theofanis Karaletsos, Serge J. Belongie, and Gunnar Rätsch. 2016. When crowds hold privileges: Bayesian unsupervised representation learning with oracle constraints. In *ICLR*.
- [37] Yoshinobu Kawahara and Masashi Sugiyama. 2009. Change-Point Detection in Time-Series Data by Direct Density-Ratio Estimation. In SDM. 389–400.
- [38] S Sathiya Keerthi and Chih-Jen Lin. 2003. Asymptotic behaviors of support vector machines with Gaussian kernel. *Neural computation* 15, 7 (2003), 1667–1689.
- [39] Daniel Kifer, Shai Ben-David, and Johannes Gehrke. 2004. Detecting Change in Data Streams. In *PVLDB*, 180–191.
- [40] Nick Koudas, Avishek Saha, Divesh Srivastava, and Suresh Venkatasubramanian. 2009. Metric functional dependencies. In *ICDE*. 1275–1278.
- [41] Hans-Peter Kriegel, Peer Kröger, Erich Schubert, and Arthur Zimek. 2012. Outlier Detection in Arbitrarily Oriented Subspaces. In ICDM. 379–388.
- [42] Sebastian Kruse and Felix Naumann. 2018. Efficient discovery of approximate dependencies. PVLDB 11, 7 (2018), 759–772.
- [43] Ludmila I. Kuncheva and William J. Faithfull. 2014. PCA Feature Extraction for Change Detection in Multidimensional Unlabeled Data. *IEEE Trans. Neural Netw. Learning Syst.* 25, 1 (2014), 69–80.
- [44] Ester Livshits, Alireza Heidari, Ihab F. Ilyas, and Benny Kimelfeld. 2020. Approximate Denial Constraints. PVLDB 13, 10 (2020), 1682–1695.
- [45] Haochuan Lu, Huanlin Xu, Nana Liu, Yangfan Zhou, and Xin Wang. 2019. Data Sanity Check for Deep Learning Systems via Learnt Assertions. CoRR abs/1909.03835 (2019). arXiv:1909.03835
- [46] Tveten Martin and Ingrid K. Glad. 2019. Online Detection of Sparse Changes in High-Dimensional Data Streams Using Tailored Projections. arXiv preprint arXiv:1908.02029 (2019).
- [47] Victor Y. Pan and Zhao Q. Chen. 1999. The Complexity of the Matrix Eigenproblem. In ACM Symposium on Theory of Computing. 507–516.
- [48] Thorsten Papenbrock, Jens Ehrlich, Jannik Marten, Tommy Neubert, Jan-Peer Rudolph, Martin Schönberg, Jakob Zwiener, and Felix Naumann. 2015. Functional dependency discovery: An experimental evaluation of seven algorithms. *PVLDB* 8, 10 (2015), 1082–1093.
- [49] Eduardo HM Pena, Eduardo C de Almeida, and Felix Naumann. 2019. Discovery of approximate (and exact) denial constraints. *PVLDB* 13, 3 (2019), 266–278.
- [50] Abdulhakim Qahtan, Nan Tang, Mourad Ouzzani, Yang Cao, and Michael Stonebraker. 2020. Pattern functional dependencies for data cleaning. *PVLDB* 13, 5 (2020), 684–697.
- [51] Abdulhakim Ali Qahtan, Basma Alharbi, Suojin Wang, and Xiangliang Zhang. 2015. A PCA-Based Change Detection Framework for Multidimensional Data Streams: Change Detection in Multidimensional Data Streams. In SIGKDD. 935– 944.
- [52] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In SIGKDD. 1135–1144.
- [53] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. 1985. Learning internal representations by error propagation. Technical Report. California Univ San Diego La Jolla Inst for Cognitive Science.
- [54] Leszek Rutkowski, Maciej Jaworski, Lena Pietruczuk, and Piotr Duda. 2015. A New Method for Data Stream Mining Based on the Misclassification Error. IEEE Trans. Neural Netw. Learning Syst. 26, 5 (2015), 1048–1059.
- [55] Suchi Saria and Adarsh Subbaswamy. 2019. Tutorial: Safe and Reliable Machine Learning. CoRR abs/1904.07204 (2019). arXiv:1904.07204
- [56] Sebastian Schelter, Tammo Rukat, and Felix Bießmann. 2020. Learning to Validate the Predictions of Black Box Classifiers on Unseen Data. In SIGMOD. 1289–1299.
- [57] Bernhard Schölkopf, Alexander J Smola, Francis Bach, et al. 2002. Learning with kernels: support vector machines, regularization, optimization, and beyond.
- [58] Tegjyot Singh Sethi and Mehmed M. Kantardzic. 2017. On the reliable detection of concept drift from streaming unlabeled data. *Expert Syst. Appl.* 82 (2017), 77–99.
- [59] Tegjyot Singh Sethi, Mehmed M. Kantardzic, and Elaheh Arabmakki. 2016. Monitoring Classification Blindspots to Detect Drifts from Unlabeled Data. In IEEE International Conference on Information Reuse and Integration, IRI. 142–151.
- [60] Shaoxu Song and Lei Chen. 2011. Differential dependencies: Reasoning and discovery. ACM Transactions on Database Systems (TODS) 36, 3 (2011), 1–41.
- [61] Xiuyao Song, Mingxi Wu, Christopher M. Jermaine, and Sanjay Ranka. 2007. Statistical change detection for multi-dimensional data. In SIGKDD. 667–676.

- [62] V. M. A. Souza, D. F. Silva, J. Gama, and G. E. A. P. A. Batista. 2015. Data Stream Classification Guided by Clustering on Nonstationary Environments and Extreme Verification Latency. In SDM. 873–881.
- [63] Timo Sztyler and Heiner Stuckenschmidt. 2016. On-body Localization of Wearable Devices: An Investigation of Position-Aware Activity Recognition. In *PerCom*. 1–9
- [64] David M. J. Tax and Klaus-Robert Müller. 2003. Feature Extraction for One-Class Classification. In ICANN/ICONIP. 342–349.
- [65] Ashish Tiwari, Bruno Dutertre, Dejan Jovanovic, Thomas de Candia, Patrick Lincoln, John M. Rushby, Dorsa Sadigh, and Sanjit A. Seshia. 2014. Safety envelope for security. In *HiCoNS*. 85–94.
- [66] Alexey Tsymbal. 2004. The problem of concept drift: definitions and related work. Computer Science Department, Trinity College Dublin 106, 2 (2004), 58.
- [67] Alexey Tsymbal, Mykola Pechenizkiy, Padraig Cunningham, and Seppo Puuronen. 2006. Handling Local Concept Drift with Dynamic Integration of Classifiers: Domain of Antibiotic Resistance in Nosocomial Infections. In IEEE International Symposium on Computer-Based Medical Systems (CBMS). 679–684.
- [68] Francesco Turchini, Lorenzo Seidenari, and Alberto Del Bimbo. 2017. Convex polytope ensembles for spatio-temporal anomaly detection. In *International Conference on Image Analysis and Processing*. 174–184.

- [69] Martin Tveten. 2019. Which principal components are most sensitive to distributional changes? arXiv preprint arXiv:1905.06318 (2019).
- [70] Vladimir Vapnik, Steven E Golowich, and Alex J Smola. 1997. Support vector method for function approximation, regression estimation and signal processing. In *NeurIPS*. 281–287.
- [71] Kush R. Varshney. 2019. Trustworthy machine learning and artificial intelligence. ACM Crossroads 25, 3 (2019), 26–29.
- [72] Kush R. Varshney and Homa Alemzadeh. 2017. On the Safety of Machine Learning: Cyber-Physical Systems, Decision Sciences, and Data Products. *Big Data* 5, 3 (2017), 246–255.
- [73] Heng Wang and Zubin Abraham. 2015. Concept Drift Detection for Imbalanced Stream Data. CoRR abs/1504.01044 (2015). arXiv:1504.01044
- [74] Jing Nathan Yan, Oliver Schulte, Mohan Zhang, Jiannan Wang, and Reynold Cheng. 2020. SCODED: Statistical Constraint Oriented Data Error Detection. In SIGMOD. 845–860.
- [75] Yunjia Zhang, Zhihan Guo, and Theodoros Rekatsinas. 2020. A Statistical Perspective on Discovering Functional Dependencies in Noisy Data. In SIGMOD. 861–876.