

Mosaic: A Sample-Based Database System for Open World Query Processing

Laurel Orr, Samuel Ainsworth, Walter Cai,
Kevin Jamieson, Magda Balazinska, Dan Suciu
Paul G Allen School of Computer Science, University of Washington
{ljorr1, skainswo, walter, jamieson, magda, suciu}@cs.washington.edu

ABSTRACT

Data scientists have relied on samples to analyze populations of interest for decades. Recently, with the increase in the number of public data repositories, sample data has become easier to access. It has not, however, become easier to analyze. This sample data is arbitrarily biased with an unknown sampling probability, meaning data scientists must manually debias the sample with custom techniques to avoid inaccurate results. In this vision paper, we propose MOSAIC, a database system that treats samples as first-class citizens and allows users to ask questions over populations represented by these samples. Answering queries over biased samples is non-trivial as there is no existing, standard technique to answer population queries when the sampling probability is unknown. In this paper, we show how our envisioned system solves this problem by having a unique sample-based data model with extensions to the SQL language. We propose how to perform population query answering using biased samples and give preliminary results for one of our novel query answering techniques.

1. INTRODUCTION

At its core, data science’s goal is to answer questions about populations of interest. As many population datasets are either inaccessible (e.g., private medical data) or do not exist (e.g., all people living today), data scientists must use samples to approximately answer questions over these populations. Historically, data scientists accessed samples mainly by designing their own sampling scheme and manually collecting the sample for processing.

However, in recent years, the number of public data repositories has increased, e.g., Data.World, and government data has become more available. These sources give data scientists access to hundreds of samples, but they do not necessarily indicate how the samples were collected. This is problematic because samples can be arbitrarily biased which, if not corrected for, can lead to inaccurate analytical results.

Typically, samples are debiased using knowledge of how the sample was collected (sampling mechanism), but as the

mechanism is not always available to data scientists, they are forced to develop alternative techniques to ensure accurate sample analysis. These alternative debiasing techniques vary across disciplines and are often customized to fit a specific type of data [50, 33, 18]. For example, in [33], the authors use an iterative reweighting algorithm to reweight a sample of individuals using known population aggregates. In [18], the authors learn probabilities from a sample and population aggregates and then use sampling to generate a representative population for later analysis.

In all of these cases, the scientists manually implement solutions hard-coded for a specific dataset in order to answer queries about some population of data. They do not use a database management system (DBMS) for query processing. DBMSs are built to efficiently answer questions about data and should be an ideal system for these data scientists, but they are not being widely used by them.

One reason DBMSs have not been more broadly adopted by the data science community is the closed world assumption taken by traditional DBMSs. This assumption states that any tuple not in the database does not exist. With sample analytics, this assumption is not valid. Sample analytics require the open world assumption where tuples not in the database could still exist.

In this vision paper, we propose a novel DBMS that automatically assumes an open world and allows scientists to query populations that do not fully exist in the database. This is a new problem as this influx of publicly available, arbitrarily biased data is recent, and we are not aware of other prototype systems that tackle open world data debiasing. In recent prior work [42], we explored a first, specific approach to this problem. In this paper, we paint a broader vision based on our experience with our prototype.

The challenge in building an open world database is that no general purpose solution exists to debias sample data without knowing the sampling mechanism. We will need to invent generalized techniques for debiasing. Further, an open world DBMS will need to provide an interface for users to ask questions over the debiased data.

In this paper, we describe our envisioned system, MOSAIC, named because it is built from a collection of samples (pieces of the population). MOSAIC builds on our prior work in [42] and extends it in three main ways:

1. It treats samples as first-class citizens, allowing users to define and ingest sample data.
2. It allows users to define and query populations of interest through the use of population metadata (e.g., population aggregate data).

3. It gives users the ability to choose how “open” it can be when answering queries (i.e. how much can it assume tuples are missing from the data).

To achieve each open world feature, we make the following technical contributions:

1. A sample-oriented data model with new keywords to the SQL query language so users can add, query, and modify samples and populations (Sec. 3).
2. Two open world query processing techniques utilizing samples and population metadata (Sec. 4, Sec. 5).

2. MOTIVATING EXAMPLE

As a motivating example, inspired by [49], take a data scientist who is using Yahoo! emails to estimate the number of migrants in European countries. As there are no existing debiasing systems, to correct for selection bias from Internet usage varying by country, age, and gender, the data scientist manually models the bias and uses maximum likelihood to fit her model to reported, ground-truth migrant statistics from Eurostat. After learning a corrected weight of each tuple, she asks various queries on the reweighted email sample.

To show the ease and versatility if she instead used MOSAIC, we show some example MOSAIC queries below but explain the semantics of these queries in Sec. 3. We assume Eurostat has information on the number of migrants per country and the number of migrants per email provider, and we leave out attribute declarations for space.

```

1 CREATE TEMPORARY TABLE Eurostat;
  ...Ingest Eurostat reports to Eurostat table
3 CREATE GLOBAL POPULATION EuropeMigrants;
4 CREATE METADATA EuropeMigrants_M1 AS
5   (SELECT country, reported_count
6    FROM Eurostat);
7 CREATE METADATA EuropeMigrants_M2 AS
8   (SELECT email, reported_count
9    FROM Eurostat);
10 CREATE SAMPLE YahooMigrants AS
11   (SELECT * FROM EuropeMigrants
12    WHERE email = Yahoo);
  ...Ingest Yahoo sample to YahooMigrants

15 SELECT SEMI-OPEN country, email, COUNT(*)
16   FROM EuropeMigrants
17   GROUP BY country, email;
  -- UK, Yahoo, 20000
  -- FR, Yahoo, 9000
  -- ...

22 SELECT OPEN country, email, COUNT(*)
23   FROM EuropeMigrants
24   GROUP BY country, email;
  -- UK, Yahoo, 20000
  -- UK, AOL, 20
  -- ...

```

At a high level, she creates a population of all European migrants (line 3) and adds the Eurostat report as metadata (lines 4-9). She then creates a sample relation representing migrants with Yahoo! email (lines 10-12) and then directly queries the population of all European migrants. MOSAIC automatically debiases the Yahoo! migrant sample using the Eurostat metadata.

The first query (lines 15-17) shows the number of migrants per country and email provider if MOSAIC just performs sample reweighting. As the sample has no information on

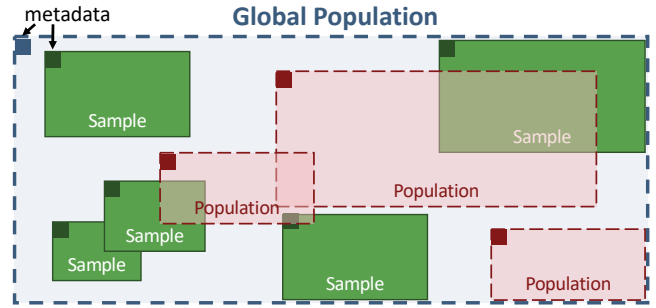


Figure 1: Data model of Mosaic.

non-Yahoo! email users, the only email provider is Yahoo!. The second query (lines 22-24) shows that MOSAIC can also use the metadata to generate entirely new, missing tuples, such as migrants who use AOL email.

3. DATA MODEL

We now describe the data model and main user interactions of MOSAIC. At a high level, MOSAIC defines two different specialized relations to help users analyze samples. In addition to these specialized relations, there is metadata associated with these relations to aid in analysis. Lastly, we define a new query construct to allow users to specify how much freedom MOSAIC has to reweight and create tuples when answering queries.

Before diving into the model, we need to define a key concept of sampling: the sampling mechanism. Given some reference global population being sampled, the sampling mechanism is the probability of an individual being included in the sample, denoted $Pr_S(t)$.

A standard approach for sample analysis is to reweight the sample, meaning each tuple gets assigned a weight indicating how many tuples in the population that sample tuple represents. If the sampling mechanism is known, the standard approach is to reweight the tuples by $1/Pr_S(t)$ [7].

3.1 Relations

There are three main types of relations in MOSAIC: population, sample, and auxiliary. We describe each relation and give our proposed syntax for declaring each of them.

(1) **(Global) Population Relation.** A population relation represents some set of tuples that could exist but are not fully known to MOSAIC. Populations are used for describing the set of tuples a scientist wants to query. When declaring a population, the user can use the GLOBAL keyword to enforce that this is the global population. The *global population (GP)* is the population that contains all other populations and samples and is used as a reference population when defining other populations and samples. Its use will become especially clear when doing sample reweighting in Sec. 4, and for now, we assume the user defines only one GP before defining other populations and samples.

A population is declared by

```

CREATE [GLOBAL] POPULATION <pop> (A1 type, ...)
  [AS (SELECT A1, ... FROM <gl_pop> WHERE <pred>)];

```

where there can be an arbitrary number of attributes. If the global keyword is not used, the population must be defined with a SELECT statement over a global population.

(2) **Sample Relation.** A sample relation represents a set

of tuples that do exist in the GP and that MOSAIC has access to. A sample can have an associated sampling mechanism defined over the GP, but it is not required.

A sample is declared by

```
CREATE SAMPLE <samp> (<A1 type, <A2 type, ...>) AS
(SELECT <A1, <A2, ...> FROM <gl_pop> [WHERE <pred>]
[USING MECHANISM <mechanism> PERCENT <perc>]);
```

where the mechanism is some probability of inclusion defined over the GP, and the sample percent represents the size of the sample.

Some example mechanisms are UNIFORM PERCENT 10 or STRATIFIED ON A_1 PERCENT 20. The former states that this sample is a 10 percent uniform sample, while the later states that this is a stratified sample on attribute A_1 where the sample contains 20 percent of the GP.

(3) **Auxiliary Relation.** Auxiliary relations are defined outside of any global population. They act the same as relations in traditional databases. In MOSAIC, they are primarily used for preprocessing, importing, and modifying data before ingesting it into MOSAIC’s specialized relations. As these relations are created using the same constructs as standard SQL, we leave them out.

3.2 Metadata

MOSAIC uses specialized metadata for query processing. For sample relations, the metadata is tuple weights, initialized to be one for every tuple (the user can modify their initial values if desired). These weights are updated by MOSAIC’s query engine (Sec. 4.1) for query processing.

For population relations, the metadata is ground-truth information about the population. For example, the metadata could be covariances of various attribute combinations. While we envision a system that incorporates a variety of metadata sources, in this paper, we focus on using aggregate values for one or two attributes; i.e., 1- or 2-dimensional histograms. These histograms (marginals) are commonly released by corporations or governments to provide statisticians with population information while preserving anonymity (e.g., Data.Gov yearly reports). There is nothing preventing our system, however, from supporting higher-dimensional histograms.

When MOSAIC answers queries over populations, it ensures these marginals are satisfied (Sec. 4).

The user specifies metadata via the command

```
CREATE METADATA <relation> AS
(SELECT <Ai, [Aj], COUNT(*)> FROM <aux_rel>
GROUP BY <Ai, [Aj]);
```

The user can update the initial sample weights via a similar command, but it is not required.

3.3 Queries

A core piece of MOSAIC’s model is its query processing over populations. The challenge with answering queries over a population is that MOSAIC does not have access to all the population tuples.

To address this, MOSAIC adds a *level of visibility* to queries. The visibility of a query represents how MOSAIC can use the underlying samples. Motivated by the debiasing techniques discussed in Sec. 1, MOSAIC supports three main ways samples are accessed and used in analysis. (1) MOSAIC can answer queries using the samples directly without any attempt at debiasing. This is equivalent to standard

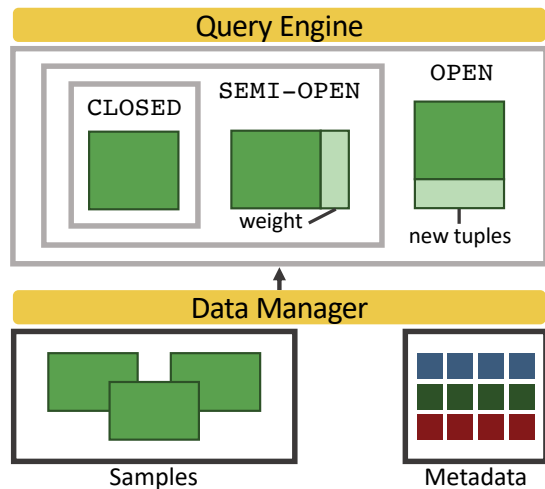


Figure 2: Proposed architecture of Mosaic.

data integration systems where the GP is treated as a global database, and the samples are local views over this database. (2) MOSAIC can reweight the samples to answer queries. Without the sampling mechanism, this is non-trivial. (3) MOSAIC can not only re-balance the samples but also infer and create missing tuples (e.g., inferring the number of UK migrants with AOL email). This allows MOSAIC to debias the sample and account for the fact that not all population tuples exist in the database.

When issuing a query, the user can decide which of the three visibility levels MOSAIC should use. The visibility level can be one of either CLOSED, SEMI-OPEN, or OPEN and appears after the SELECT. Closed means MOSAIC can only use samples (1), semi-open means it can reweight the samples (2), and open means it can reweight the samples and generate missing tuples (3).

The choice between CLOSED and SEMI-OPEN or OPEN represents if the user wants to make the open or closed world assumption. The choice between SEMI-OPEN and OPEN represents a trade-off between false-positive and false-negative tuples. In this setting, false-positives are tuples that MOSAIC returns as existing but are not (vice-versa for false-negatives). Specifically, if n is the number of tuples existing in the population but are not present in the sample, then SEMI-OPEN will have n false negative tuples but zero false positive. OPEN queries will potentially receive fewer false negatives but more false positives than SEMI-OPEN. The following table summarizes the trade-offs:

	False Negative	False Positive	Assumption
CLOSED	n	0	Closed
SEMI-OPEN	n	0	Open
OPEN	$\leq n$	≥ 0	Open

The choice of query visibility level depends on the user’s application and desired false positive and false negative rates.

We now describe how MOSAIC processes queries using metadata under the different levels of visibility.

4. QUERY EVALUATION

Fig. 2 shows the high level architecture of MOSAIC (we leave out auxiliary relations as they are standard SQL relations). The two types of stored data are the sample data

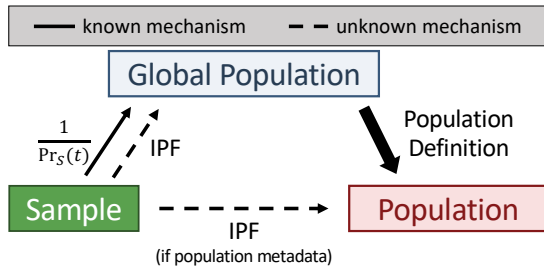


Figure 3: Semi-open world query evaluation.

and metadata. The user can interact with MOSAIC by either querying a population or updating the metadata.

To simplify the discussion, we make three assumptions:

1. The population attributes are contained in the sample attributes (i.e., we do not need to join samples).
2. When a population query gets issued, the query engine receives a single, optimal sample to use (this can be relaxed by unioning samples over shared attributes).
3. There is metadata, i.e., marginals, about the global population or the query population.

As discussed in Sec. 3, MOSAIC answers queries using one of three different visibilities. As closed queries are the same as answering queries using views in a LAV data integration system, we refer the reader to [30, 25] for query answering techniques and focus on semi-open and open queries. These two types of queries require novel system components, and in this section, we propose our solution that relies on marginals to answer arbitrary queries.

4.1 Semi-Open Queries and Reweighting

Recall that semi-open queries allow MOSAIC to perform sample reweighting. The two important subcases to consider are when the sampling mechanism is known versus unknown.

When the sampling mechanism is known, it is defined with respect to the global population. This is important as it guarantees that we can always approximately answer population queries. We use the known mechanism to reweight the sample by the inverse of its inclusion probability. Then we treat the population definition query as a view over this reweighted sample and run the user’s query over this view. Fig. 3 shows this process.

Research Question. The main research question for semi-open queries is how to answer queries when the sampling mechanism is unknown using marginals.

Proposed Solution. We addressed this problem partially in prior work when building THEMIS, an open world database system that automatically performs sample debiasing using population aggregates [42]. THEMIS merges two techniques to answer simple count queries and group by queries. One approach performs sample reweighting using a technique called Iterative Proportional Fitting (IPF) [27, 13]. The other approach builds a Bayesian network to represent the population probability distribution.

MOSAIC leverages the IPF technique presented in [42] to answer arbitrary queries over samples. Specifically, we reweight the sample so that the given marginals are satisfied.

If there is metadata over the query population, MOSAIC uses IPF to directly reweight the sample to represent

that population (bottom dashed line in Fig. 3) and answers queries over the reweighted sample. If there is only metadata over the global population (left dashed line in Fig. 3), MOSAIC reweights to represent the global population and then treats the query population as a view. Note that the accuracy will likely be lower when reweighting to fit global population (left dashed line) than reweighting to fit the query population directly (bottom dashed line) as biases that exist in the query population may not be captured when learning the global population.

4.2 Open Queries and Population Generation

Open queries over a population allow MOSAIC to use a sample and marginals to generate or infer missing tuples (i.e., answer queries with fewer false negatives). We use the term infer because some aggregate queries can be answered without needing to materialize missing tuples.

At a high level, if we can learn the population probability distribution, we can either answer queries over this distribution directly or over data sampled from this distribution that is representative of the population. For example, if we model the probability distribution as a Bayesian network, we can answer COUNT(*) queries using direct inference over the network. However, for top-k or group by queries, we need to materialize a sample.

Generating data to fit some joint distribution requires a *generative model* [23] as it allows sampling from the distribution. Explicitly defined generative models define a parametric specification for the likelihood which is then solved for and sampled from. Implicitly defined generative models provide some process for generating samples without specifying the likelihood. For example, Bayesian networks and variational autoencoders [16] are explicitly defined models while generative adversarial networks [24] are implicitly defined ones.

The benefit of using an explicitly defined model is the potential for aggregate queries to be answered without having to materialize data; e.g., the Bayesian network COUNT(*) example above. The downsides, however, are that the explicit distribution will make some, possibly incorrect, assumptions over the population. Bayesian networks, for example, assume the population satisfies the independent assumptions implied by the network. As we do not have access to the population data, we have no way of verifying if the model assumptions are accurate or not.

Therefore, in MOSAIC, we propose to take the alternative approach of using an implicitly defined generative model to learn the population as this requires no distributional assumptions. However, any generative model can be plugged in and used to answer open queries as long as it can be trained on sample data and marginals.

Research Question. The main challenge for learning a generative model is how to train one without access to the entire population. We only have a biased sample and 1-or 2-dimensional marginals, and we cannot just build a generative model over the sample as it is not representative of the population.

Proposed Solution. We propose to modify a Wasserstein generative adversarial network (WGAN) to use sample and marginal data rather than the entire population. Our motivation for modifying a WGAN is that the Wasserstein distance can be calculated very efficiently on 1-dimensional marginals (Sec. 5). We call our WGAN a marginal-

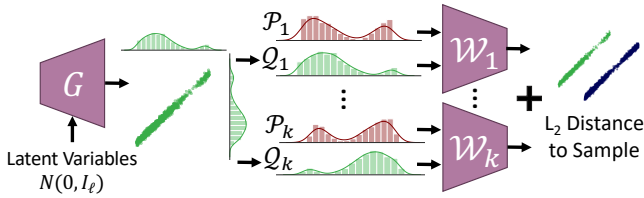


Figure 4: Marginal-constrained Wasserstein GAN.

constrained WGAN (M-WGAN) as it learns to sample population data that satisfies the marginals.

5. MARGINAL-CONSTRAINED WGANs

We now detail our proposed open query processing technique. We briefly review WGANs, describe our M-WGAN, and give preliminary results.

5.1 WGAN Overview

Generative adversarial networks (GANs) are comprised of two networks, a generator G and discriminator D , pitted against each other in a game theoretic sense. The discriminator attempts to distinguish between data from the true distribution and samples from the generator. Meanwhile, the generator takes as inputs random variables, typically Gaussian or uniform, from some latent space and is tuned to output samples that fool the discriminator. Upon successfully training both networks, the generator will output samples that closely match the true data distribution.

The Wasserstein or earth mover’s distance, denoted \mathcal{W} , quantifies the optimal transport plan of moving “mass” between two distributions [1]. A Wasserstein GAN is a framework for training G to generate data that minimizes the Wasserstein distance to the real distribution. Specifically, if \mathcal{P} represents the real distribution and \mathcal{Q} represents the generator distribution, WGANs solve

$$\min_G \mathcal{W}(\mathcal{P}, \mathcal{Q}).$$

Although the discriminator does not appear in this formulation, [5] explains how the discriminator is used to calculate an upper bound on \mathcal{W} .

The problem with using traditional WGANs is they are trained with complete knowledge of \mathcal{P} . We need to modify WGANs so they can use the sample and marginals.

5.2 Marginal-constrained WGANs

We make two alterations to the standard WGAN. The first is to modify the optimization function to handle multiple marginals and a sample, and the second is to compute the discriminator exactly. Our proposed network structure is shown in Fig. 4. Before discussing our solution, we need to define the assumptions of our WGAN model.

Manifold Hypothesis The population and sample live in low dimensional \mathbb{R}^r space but are embedded in higher-dimensional \mathbb{R}^d space with $r \ll d$ [39].

Sample Coverage For every point in our population, there is some point in our sample that is nearby.

The manifold hypothesis is a standard assumption made on high dimensional data and, in conjunction with the sample coverage assumption, means our sample informs us of the manifold our population data lives on. In particular, sample coverage requires our sample to give us the necessary

structural information without needing to be distributionally correct. A stratified sample, for example, meets this criteria.

To more easily introduce our modifications, we first assume we only have 1-dimensional marginals, and then we extend our model to 2-dimensional ones. A subtle point is that if the population marginals do not cover all d attributes (dimensions) of the sample, the model has no way of learning even the sample distribution of those attributes. Therefore, we add marginals from the sample into the set of population marginals for those uncovered attributes. We leave it as future work to determine the optimal set of 2-dimensional sample marginals to add in the 2-dimensional case.

1-dimensional marginals. Our first modification of the standard WGAN is to the error function. For notation, denote our sample as S , and assume we have an index set \mathcal{I}_1 of the 1-dimensional marginals. For each $i \in \mathcal{I}_1$, \mathcal{P}_i represents the distribution of attribute A_i of the population data (likewise for \mathcal{Q}_i). Our loss function becomes

$$\min_G \left[\sum_{i \in \mathcal{I}_1} \mathcal{W}(\mathcal{P}_i, \mathcal{Q}_i) + \lambda \mathbb{E}_{x \sim G} \left[\min_{y \in S} \|x - y\|_2 \right] \right]$$

where λ is a tuning parameter that trades off between fitting the population marginals and respecting the structure of the sample data. Note that G takes as input random Gaussian variables of latent dimension ℓ where ℓ is treated as another tuning parameter.

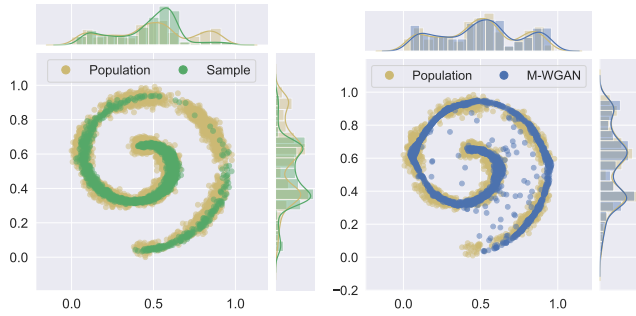
The second modification we make is to compute the Wasserstein distance exactly [48] instead of using the discriminator approach from [5]. This modification is crucial to the practicality of our approach. Not only is computing \mathcal{W} efficient for 1-dimensional data, but it makes the discriminator exact and avoids the need to train discriminator networks (one per marginal).

2-dimensional marginals. As we no longer have only 1-dimensional marginals, we cannot compute the exact Wasserstein distance efficiently as it requires solving a linear program for *every* step during training, which is prohibitively slow [21, 15]. However, by using the sliced Wasserstein distance [15], we can randomly project the marginals onto multiple one dimensional spaces and compute the Wasserstein distance exactly for each projection.

Specifically, in addition to \mathcal{I}_1 , let \mathcal{I}_2 be the index set of the 2-dimensional marginals such that for each $\{i, j\} \in \mathcal{I}_2$, $\mathcal{P}_{i,j}$ represents the *joint* distribution of the attribute pair A_i and A_j (likewise for \mathcal{Q}). Further, assume we have a set of p linear projections $\omega \in \Omega$ randomly generated and normalized to be on the unit sphere. Our loss function becomes

$$\min_G \left[\sum_{i \in \mathcal{I}_1} \mathcal{W}(\mathcal{P}_i, \mathcal{Q}_i) + \frac{1}{p} \sum_{\{i,j\} \in \mathcal{I}_2} \sum_{\omega \in \Omega} \mathcal{W}(\mathcal{P}_{i,j}\omega, \mathcal{Q}_{i,j}\omega) + \lambda \mathbb{E}_{x \sim G} \left[\min_{y \in S} \|x - y\|_2 \right] \right] \quad (1)$$

where $\mathcal{P}_{i,j}\omega$ is the 2-dimensional marginal over A_i and A_j projected by ω . Note that while this discussion has focused on 2-dimensional marginals, the projection technique is applicable for arbitrarily higher-dimensional marginals.



(a) Biased sample and population. (b) M-WGAN sample and population.

Figure 5: Population data with original biased sample (a) and M-WGAN generated sample (b).

5.3 Preliminary Results

We first evaluate just our M-WGAN¹ approach on a synthetic 2-dimensional population we can visualize. We then evaluate IPF (MOSAIC’s SEMI-OPEN evaluation technique) and M-WGAN on the flights data from [17] by running various aggregate queries. We compare against a uniformly reweighted sample which is the standard approximate query processing technique when there is no knowledge of how the sample was generated.

Synthetic Data. We generate a 2-dimensional spiral population following the experiments from [9] and generate a biased sample from this population with 10,000 rows (Fig. 5 (a)). After training our M-WGAN², we use it to generate a sample also with 10,000 rows (Fig. 5 (b)). Note that the generated data more closely matches the marginals while maintaining the spiral shape. Then, we uniformly reweight both the generated sample and the biased sample for query answering. We issue 100 random 2-dimensional range queries (i.e., counting the number of tuples lying in a bounding square box) for various box sizes. We repeat this for 10 different generated samples and report the average percent difference across the different samples.

Fig. 6 show box plots (X is average) of the average query error for the 100 queries where the whiskers show the 3rd and 97th percentiles. We repeat this as we increase the size of our bounding box to cover a greater fraction of the data. For example, a width coverage of 0.8 means the range queries for 80 percent of the data on one dimension and 80 percent of the data on the other dimension. We see that we always outperform the uniformly reweighted sample except when the range is very narrow (asking for 1/100th of the possible values). For this highly selective query, both methods achieve high query error as both are more likely to have false negative tuples.

Flights Data. We use the flights data from IDEBench [17], a real-world benchmark consisting of domestic flights in the United States (continuous attributes have

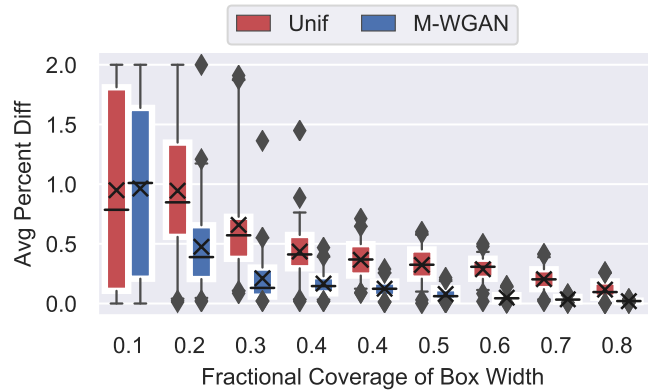


Figure 6: Average percent difference of uniform reweighting versus M-WGAN for 2 dimensional queries.

Flights	Abbrv	M-WGAN Dim
carrier	C	14
taxi_out	O	1
taxi_in	I	1
elapsed_time	E	1
distance	D	1

Table 1: Flights and IMDB attributes.

been rounded to whole numbers), to explore the strengths and weaknesses of our approach. In particular, we study how IPF and M-WGAN perform when we query numerical attributes only compared to GROUP BY queries with a single categorical group attribute. We use the attributes shown in Table 1 and filter the data to just include flights in 2015 or 2016 (426,411 rows total) for faster processing. For M-WGAN training, we one-hot encode the categorical variables and scale all attributes to be between 0 and 1. The attributes’ encoded dimensionality (equivalent to the number of distinct values) is also shown in Table 1.

For our sample, we generate a biased 5 percent sample (21,320 rows) of flights with an elapsed flight time of more than 200 minutes with a 95 percent bias, meaning 95 percent of the tuples have a long flight time. For our metadata, we generate population marginals from the four attribute pairs of (C, E), (O, E), (I, E), (D, E). As the numerical attributes are already whole numbers, we do not need to build histograms, and the marginals are just projections of the population data.

Inspired by the queries from [17], we run the eight aggregate queries shown in Table 2. To run the aggregate queries over a weighted sample, we simply modify the aggregate to be over a weight attribute (e.g. COUNT(*) becomes SUM(weight)). To reduce the variance of the M-WGAN approach, we generate 10 samples with the same number of rows as the original sample and uniformly reweight the generated sample to match the size of the population. We return the groups appearing in all 10 answers, averaging the aggregate value.

We add a softmax layer for the categorical variable and train our M-WGAN similarly as above but with the latent dimension ℓ being the same as the input dimensionality. As our goal is an accurate representation of the data, we did not feel it necessary to experiment with reducing ℓ . During

¹Code available at https://gitlab.cs.washington.edu/uwdb/project_mosaic_db.

²We use 3 ReLU FC layers with 100 nodes each. We use $\lambda = 0.04$, $\ell = 2$, a batch size of 500, and apply batch normalization after each layer. We use Pytorch’s Adam optimizer with the default settings and an initial learning rate of 0.001 that decreases by a factor of 10 if a plateau is reached during training.

Id	Query
1	SELECT AVG(D) FROM F WHERE E > 200
2	SELECT AVG(I) FROM F WHERE E < 200
3	SELECT AVG(E) FROM F WHERE D > 1000
4	SELECT AVG(O) FROM F WHERE D < 1000
5	SELECT C, AVG(D) FROM F WHERE E > 200 AND C IN ['WN', 'AA']
6	SELECT C, AVG(I) FROM F WHERE E < 200 AND C IN ['WN', 'AA']
7	SELECT C, AVG(E) FROM F WHERE D > 1000 AND C IN ['WN', 'AA']
8	SELECT C, AVG(O) FROM F WHERE D < 1000 AND C IN ['US', 'F9']

Table 2: The six SQL queries run in Fig. 7. We leave out the GROUP BY clause and replace Flights with F for space. ‘US’ and ‘F9’ stand for US Airways and Frontier (less popular carriers) while ‘AA’ and ‘WN’ stand for American and Southwest (more popular carriers).

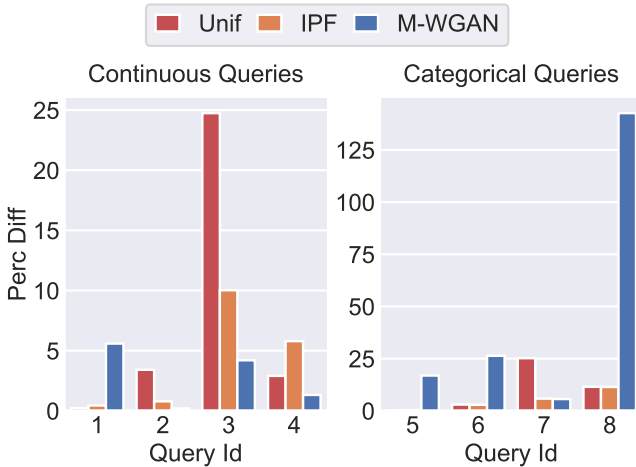


Figure 7: Average percent difference of uniform reweighting (default AQP) versus IPF versus M-WGAN for queries 1-4 (left) and queries 5-8 (right). The y axis is on a different scale between the left and right plots.

training, we leave the softmax output continuous and only force the output to be binary for data generation.

We choose the model parameters by a small hyperparameter grid search³, running the models for three epochs (each epoch is one pass over the population marginals). We select the model receiving the lowest average query error from running 200 random queries over the continuous attributes with the same template as queries 1-4 where the attributes and predicates are randomly generated. We then rerun the chosen model with four different random initializations for 80 epochs and choose the one receiving the lowest error on the same 200 queries. Our final parameters are $\lambda = 1^{-7}$, $p = 1000$, 5 layers with 50 nodes each, and a batch size of 500.

Fig. 7 (left) shows the average percent difference of the

³We search over the number of layers = 3, 5, 10, number of hidden nodes = 50, 200, and $\lambda = 0.000001, 0.0000001$. When the number of hidden nodes is 200 (50), we do not try 10 (3) layers.

first four aggregate queries using uniform reweighting, IPF, and M-WGAN. These are the queries over continuous attributes only. We see that all methods achieve an average error of less than 25 percent with M-WGAN having the lowest max error of 7.5 percent. Averaged across all four queries, M-WGAN gets the lowest query error compared to IPF and Unif.

Surprisingly, M-WGAN receives the highest error on the query with a predicate matching how the sample is biased. IPF and Unif get almost no error for this query because the sample accurately represents this predicate. This is because M-WGAN is slightly over-fitting to the population marginals and not generating as accurate a representation of the data satisfying this predicate as is in the sample.

IPF and Unif receive higher errors on query 3 because long distance flights are going to have longer flight times. As our sample is biased towards long flight times, the flight time is overestimated, even with a marginal over flight time and distance.

This trend of M-WGAN receiving the lowest error on average for numerical queries is supported by the fact that on the 200 random queries used for parameter selection, when both the true answer and M-WGAN answer are not-empty (161 queries total), *all* of our M-WGAN models achieve a lower query error than Unif. IPF also achieves a lower error than Unif. The not-empty filter is due to the fact that M-WGANs struggles to capture light hitters (rare values) and will often not generate these rare values. This is because the error reduction in capturing the tail of a distribution (i.e. light hitter values) is often not significant enough for the GAN to generate those values.

Fig. 7 (right) shows the average percent difference of the last four aggregate queries using uniform reweighting, IPF, and M-WGAN. Note that Unif and IPF get an close to zero error for query 5. For query 8, M-WGAN receives a large error as it does not generate any flights with the carrier ‘US’. This is a result of the carriers attribute being categorical and having a skewed distribution in the data. Our M-WGAN has to model an 18 dimensional space where the 14 one-hot encoded dimensions are sparse and not continuous (i.e., a one-hot encoded attribute is either a 0 or a 1). The M-WGAN takes a continuous input, the latent variable ℓ , and tries to generate data where 14 dimensions are binary and the distribution is skewed. This means there is not a significant reduction in error (Eq. 1) in generating the less popular carriers and potentially an increase in error from generating data falling between 0 and 1 for one of the binary dimensions, making it difficult to learn. We further discuss this problem in Sec. 7.

These experiments have revealed that while our M-WGAN approach is promising, there are still open questions around how to make it more accurate for heterogeneous, skewed data.

6. RELATED WORK

Our vision relates to five main areas of research: data integration/exchange, approximate query processing, distribution learning, sample selection bias correction, and social science sample analysis. To the best of our knowledge, our prior work [42] was the first system to attempt automatic sample debiasing. MOSAIC is the first database to truly treat samples as first-class citizens in a declarative manner.

The database community has extensively researched data

integration query processing (closed query processing in MOSAIC) [30, 25, 2, 28]. While data integration systems do consider the difference between the open and closed world assumption when answering queries over views, they do not reconstruct the global database. They are instead concerned with generating global schemas and giving equivalent or maximally contained answers. We go beyond data integration by providing the SEMI-OPEN and OPEN query constructs.

In a similar vein to MOSAIC’s goal of supporting data scientists is the work on open data integration and exchange [36] where the focus is on helping scientists to discover and integrate publicly available data sources without apriori knowledge of what data is available. This work centers on finding relevant relations on the web [40] and schema mapping [45]. We consider MOSAIC to work in tandem with these systems by providing scientists the tools to accurately analyze the data discovered from these open data sources.

While not concerned with data integration is the work of [10, 20] on open world probabilistic databases (PDB). Like MOSAIC, they also make an open world assumption but focus on theoretically understanding how different open world PDB semantics impact the precision and tractability of different classes of queries.

Approximate query processing (AQP) [11, 37] explores how to achieve interactive query response times by allowing for approximate query results. Similar to MOSAIC, many AQP systems aim to accurately answer queries over samples for faster query processing [4, 3]. Further, some systems utilize past or known query answers to improve performance [44, 22, 43]. The important distinction between MOSAIC and standard AQP is that AQP systems either assume they have access to the entire population or assume they have knowledge of the error from querying the sample. MOSAIC makes neither of these assumptions, making standard AQP techniques not applicable. However, if the user had knowledge of the error or assumed the error was normally distributed [22], these AQP systems could be used for SEMI-OPEN query processing in MOSAIC.

While most AQP research assumes a closed world, the work of [47, 41, 29] allow for tuples to be generated that may not exist in the sample. Like MOSAIC, [47] and [29] also use generative models while [41] uses Maximum Entropy to learn a general probability distribution. Note that all systems assume access to the entire population, which does not hold in MOSAIC.

Most similar to MOSAIC’s WGAN approach is [9] where the authors learn the joint distribution of a mixture of Gaussians from lower dimensional projections of data. Sliced WGANs [15] also learn a distribution from 1-dimensional projections of data. Further, [32] use GANs to learn a joint distribution of images in two different domains by weight sharing across two GANs, one for each domain. However, all of these have the population so can take multiple projections. We only have marginals for specific attributes that may not cover all the attributes of the population.

Although not concerned with ad-hoc query answering, the research on sample selection bias correction in machine learning tackles the problem of training and test data being drawn from different distributions; i.e., the training data is biased and not representative of the true population distribution [12, 26, 8]. This work, however, assumes access to a sample from the true distribution (test set) and is mainly

concerned with downstream tasks, such as classification or regression. MOSAIC solves the more general problem of debiasing samples for arbitrary queries using only population metadata and could be used to help machine learning scientists debias their samples.

Social science sample analysis, especially in demography research, is one of the most common use cases of biased sample analysis [38, 13, 50, 19, 18, 33, 46]. The samples are from social media or United States’ surveys, e.g., the American Community Survey, and the aggregates for bias correction are census reports. While there are a few prepackaged libraries that perform specific algorithms, like the Python package IPFN, there is no general purpose technique or data management system for these scientists. Note that demography is not the only field that can benefit from MOSAIC, but it is one of our prime use cases and examples.

The database system most closely related to sample analytics for data science is MauveDB [14]. It allows model-based views over data to support users in building statistical models or abstractions over their data. While the overarching goal of supporting data science matches that of MOSAIC, the underlying mechanisms and data model is different.

7. FUTURE WORK

We want to highlight that MOSAIC is a visionary system, and there remain numerous open challenges for MOSAIC to be fully realized. The following is a list of the more prominent challenges with ideas on solving them.

Multiple Populations. A main area of future work is to allow for multiple global populations. We not only need to keep track of which populations define the sampling mechanisms of which samples but also need to handle when samples exist in multiple populations and have a different sampling mechanism for each population. One possible solution is to treat all populations as disjoint and replicate any overlapping samples. However, we risk losing important debiasing information by this assumption. We leave it as future work to investigate ways of learning an optimal sample reweighting scheme from multiple mechanisms.

Multiple Samples. In Sec. 4, we made the assumption that the query engine receives a single sample for query answering. While simple, this is likely suboptimal as different samples represent different subsets of the population, and relying on one will result in more missed tuples. One solution is to union together all related samples and let IPF or the neural network reweight the tuples accordingly. This might present further problems if certain tuples become overly represented.

Data Integration. In Sec. 4, we also made the assumption that the population attributes are contained in the sample attributes. With real-world data, this is not the case. To relax this assumption, we will need to incorporate data integration query answering and schema mapping techniques to handle the disparate sample attributes. It is an open question whether standard data integration techniques hold or if they will need to be enhanced to handle the open world assumption of MOSAIC.

Metadata. Another aspect of future work is to add a greater variety of metadata. Our system already supports higher dimensional aggregate data, but it is an interesting direction to add metadata such as population constraints or known correlations.

Open World Density Estimation. Our M-WGAN

approach is simply one solution to open query processing. Based on [47] and [32], MOSAIC’s open query answering may be made more accurate by using more advanced sample generation approaches (e.g., [6]) or by learning multiple GANs on meaningful subsets of the data. Our M-WGAN might also be made more accurate with a modified loss term to better ensure sample coverage. Further, it is an open question whether alternative density estimation techniques, like nonparametric kernel density estimation [31], will be more accurate or efficient.

Data Encoding. The last research question is that of data representation. As mentioned in some of the (generative) model-based systems from Sec. 6 and in Sec. 5, it is unclear how best to represent heterogeneous (continuous and categorical) data for query processing. The systems of [34, 41, 29] consider all possible attribute values when building models or answering queries. For small domains, this works well, but when the active domain gets large, these systems break down. [47] takes a similar approach as MOSAIC by encoding categorical variables as binary numbers. While this reduces the dimensionality of one-hot encoding, it introduces various relationships between attribute values that may not exist. For example, suppose the three US states of CA, FL, and NY were encoded in two bits as 11, 01, and 10, respectively. This implies that the sum of FL and NY is equivalent to CA, which is misleading. We think a promising approach is that of [35] where they used a word2vec model to learn *contextualized* encodings of values from the data.

8. CONCLUSION

MOSAIC is a DBMS that treats samples as first-class citizens and allows users to ask open world queries on biased samples of data and get approximate results. We described our proposed extensions to SQL to allow users to ask questions in a declarative manner and then discussed one possible way of doing open world query processing through the use of marginal data. In particular, we presented and gave preliminary results for our novel technique of using a M-WGAN to generate population data.

While we focused on the use case of known population marginals (e.g. the census aggregate reports), the techniques presented are applicable for arbitrary or approximate marginals determined by the user. These marginals can represent any new distribution the user wants the sample to satisfy; e.g., one where classes are rescaled for downstream classification or where sensitive attributes are balanced. MOSAIC takes an important first step in building a system to better support the needs of data scientists by inherently supporting samples and the populations they represent.

Acknowledgements. This work is supported by NSF AITF 1535565 and the Intel Science and Technology Center for Big Data.

9. REFERENCES

- [1] Optimal transport and wasserstein distance (lecture notes). <http://www.stat.cmu.edu/~larry/=sml/Opt.pdf>.
- [2] S. Abiteboul and O. Duschka. Complexity of answering queries using materialized views. 2013.
- [3] S. Acharya, P. B. Gibbons, V. Poosala, and S. Ramaswamy. The aqua approximate query answering system. In *ACM Sigmod Record*, volume 28, pages 574–576. ACM, 1999.
- [4] S. Agarwal, B. Mozafari, A. Panda, H. Milner, S. Madden, and I. Stoica. Blinkdb: queries with bounded errors and bounded response times on very large data. In *Proceedings of the 8th ACM European Conference on Computer Systems*, pages 29–42. ACM, 2013.
- [5] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [6] S. Azadi, C. Olsson, T. Darrell, I. Goodfellow, and A. Odena. Discriminator rejection sampling. *arXiv preprint arXiv:1810.06758*, 2018.
- [7] J.-F. Beaumont. A new approach to weighting and inference in sample surveys. *Biometrika*, 95(3):539–553, 2008.
- [8] S. Bickel, M. Brückner, and T. Scheffer. Discriminative learning for differing training and test distributions. In *Proceedings of the 24th international conference on Machine learning*, pages 81–88. ACM, 2007.
- [9] Q. Cai, C. Zhang, and C. Peng. Learning probability density functions from marginal distributions with applications to gaussian mixtures. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pages 1148–1153. IEEE, 2005.
- [10] I. I. Ceylan, A. Darwiche, and G. Van den Broeck. Open-world probabilistic databases. In *Fifteenth International Conference on the Principles of Knowledge Representation and Reasoning*, 2016.
- [11] S. Chaudhuri, B. Ding, and S. Kandula. Approximate query processing: No silver bullet. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 511–519. ACM, 2017.
- [12] C. Cortes, M. Mohri, M. Riley, and A. Rostamizadeh. Sample selection bias correction theory. In *International conference on algorithmic learning theory*, pages 38–53. Springer, 2008.
- [13] W. E. Deming and F. F. Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, 11(4):427–444, 1940.
- [14] A. Deshpande and S. Madden. Mauvedb: supporting model-based user views in database systems. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 73–84. ACM, 2006.
- [15] I. Deshpande, Z. Zhang, and A. G. Schwing. Generative modeling using the sliced wasserstein distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3483–3491, 2018.
- [16] C. Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.
- [17] P. Eichmann, C. Binnig, T. Kraska, and E. Zraggen. Idebench: A benchmark for interactive data exploration. *arXiv preprint arXiv:1804.02593*, 2018.
- [18] B. Farooq, M. Bierlaire, R. Hurtubia, and G. Flötteröd. Simulation based population synthesis. *Transportation Research Part B: Methodological*, 58:243–263, 2013.
- [19] L. Fiorio, G. Abel, J. Cai, G. Vinué, E. Zagheni, and I. Weber. Migration curves: Using twitter data to illustrate the temporal contours of population flows. Population Association of America Annual Meeting, 2013.
- [20] T. Friedman and G. V. d. Broeck. On constrained open-world probabilistic databases. *arXiv preprint arXiv:1902.10677*, 2019.
- [21] C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. A. Poggio. Learning with a wasserstein loss. In *Advances in Neural Information Processing Systems*, pages 2053–2061, 2015.
- [22] A. Galakatos, A. Crotty, E. Zraggen, C. Binnig, and T. Kraska. Revisiting reuse for approximate query processing. *Proceedings of the VLDB Endowment*, 10(10):1142–1153, 2017.
- [23] I. Goodfellow. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.
- [24] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

- [25] A. Halevy, A. Rajaraman, and J. Ordille. Data integration: the teenage years. In *Proceedings of the 32nd international conference on Very large data bases*, pages 9–16. VLDB Endowment, 2006.
- [26] J. Huang, A. Gretton, K. Borgwardt, B. Schölkopf, and A. J. Smola. Correcting sample selection bias by unlabeled data. In *Advances in neural information processing systems*, pages 601–608, 2007.
- [27] M. Idel. A review of matrix scaling and sinkhorn’s normal form for matrices and positive maps. *arXiv preprint arXiv:1609.06349*, 2016.
- [28] P. G. Kolaitis. Reflections on schema mappings, data exchange, and metadata management. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 107–109. ACM, 2018.
- [29] M. Kulesa, A. Molina, C. Binnig, B. Hilprecht, and K. Kersting. Model-based approximate query processing. *arXiv preprint arXiv:1811.06224*, 2018.
- [30] M. Lenzerini. Data integration: A theoretical perspective. In *Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 233–246. ACM, 2002.
- [31] Q. Li and J. Racine. Nonparametric estimation of distributions with categorical and continuous data. *journal of multivariate analysis*, 86(2):266–292, 2003.
- [32] M.-Y. Liu and O. Tuzel. Coupled generative adversarial networks. In *Advances in neural information processing systems*, pages 469–477, 2016.
- [33] R. Lovelace, M. Birkin, D. Ballas, and E. van Leeuwen. Evaluating the performance of iterative proportional fitting for spatial microsimulation: new tests for an established technique. *Journal of Artificial Societies and Social Simulation*, 18(2), 2015.
- [34] Q. Ma and P. Triantafyllou. Dbest: Revisiting approximate query processing engines with machine learning models. In *Proceedings of the 2019 International Conference on Management of Data*, pages 1553–1570. ACM, 2019.
- [35] R. Marcus, P. Negi, H. Mao, C. Zhang, M. Alizadeh, T. Kraska, O. Papaemmanouil, and N. Tatbul. Neo: A learned query optimizer. *arXiv preprint arXiv:1904.03711*, 2019.
- [36] R. Miller. Open data integration. *Proceedings of the VLDB Endowment*, 11(12):2130–2139, 2018.
- [37] B. Mozafari and N. Niu. A handbook for building an approximate query engine. *IEEE Data Eng. Bull.*, 38(3):3–29, 2015.
- [38] K. Müller. *A generalized approach to population synthesis*. PhD thesis, ETH Zurich, 2017.
- [39] H. Narayanan and S. Mitter. Sample complexity of testing the manifold hypothesis. In *Advances in Neural Information Processing Systems*, pages 1786–1794, 2010.
- [40] F. Nargesian, E. Zhu, K. Q. Pu, and R. J. Miller. Table union search on open data. *Proceedings of the VLDB Endowment*, 11(7):813–825, 2018.
- [41] L. Orr, M. Balazinska, and D. Suciu. Entropydb: a probabilistic approach to approximate query processing. *The VLDB Journal*, 2019.
- [42] L. Orr, M. Balazinska, and D. Suciu. Sample debiasing in the themis open world database system. *Proceedings of the 2020 International Conference on Management of Data*, 2020.
- [43] Y. Park, A. S. Tajik, M. Cafarella, and B. Mozafari. Database learning: Toward a database that becomes smarter every time. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 587–602. ACM, 2017.
- [44] J. Peng, D. Zhang, J. Wang, and J. Pei. Aqp++: connecting approximate query processing with aggregate precomputation for interactive analytics. In *Proceedings of the 2018 International Conference on Management of Data*, pages 1477–1492. ACM, 2018.
- [45] L. Popa, Y. Velegrakis, M. A. Hernández, R. J. Miller, and R. Fagin. Translating web data. In *Proceedings of the 28th international conference on Very Large Data Bases*, pages 598–609. VLDB Endowment, 2002.
- [46] L. Sun and A. Erath. A bayesian network approach for population synthesis. *Transportation Research Part C: Emerging Technologies*, 61:49–62, 2015.
- [47] S. Thirumuruganathan, S. Hasan, N. Koudas, and G. Das. Approximate query processing for data exploration using deep generative models. *arXiv preprint arXiv:1903.10000*, 2019.
- [48] M. Werman, S. Peleg, and A. Rosenfeld. A distance metric for multidimensional histograms. *Computer Vision, Graphics, and Image Processing*, 32(3):328–336, 1985.
- [49] E. Zagheni and I. Weber. You are where you e-mail: using e-mail data to estimate international migration rates. In *Proceedings of the 4th annual ACM web science conference*, pages 348–351. ACM, 2012.
- [50] E. Zagheni and I. Weber. Demographic research with non-representative internet data. *International Journal of Manpower*, 36(1):13–25, 2015.