

SynC: A Unified Framework for Generating Synthetic Population with Gaussian Copula

Colin Wan,¹ Zheng Li,^{2,3*} Alicia Guo,⁴ Yue Zhao⁵

¹Department of Statistical Sciences, University of Toronto, Toronto, ON, Canada

²Northeastern University, Toronto Campus, Toronto, ON, Canada

³Arima Inc., Toronto, ON, Canada

⁴PwC Canada, Toronto, ON, Canada

⁵H. John Heinz III College, Carnegie Mellon University, Pittsburgh, PA, USA

colin.wan@mail.utoronto.ca, winston@arimadata.com, alicia.guo@pwc.com, zhaoy@cmu.edu

Abstract

Synthetic population generation is the process of combining multiple socioeconomic and demographic datasets from different sources and/or granularity levels, and downscaling them to an individual level. Although it is a fundamental step for many data science tasks, an efficient and standard framework is absent due to various reasons like privacy and security. In this study, we propose a multi-stage framework called **SynC** (**S**ynthetic Population via **G**aussian **C**opula) to fill the gap. SynC may first remove potential outliers in the data and then fits the filtered data with a Gaussian copula model to correctly capture dependencies and marginal distributions of the sampled survey data. Finally, SynC leverages predictive models to merge datasets into one and then scales them accordingly to match the marginal constraints. We make three key contributions in this work: 1) propose a novel framework for generating individual level data from aggregated data sources by combining state-of-the-art machine learning and statistical techniques, 2) demonstrate its value as a feature engineering tool, as well as an alternative to data collection in situations where gathering is difficult through two real-world datasets, 3) release an easy-to-use framework implementation for reproducibility, and 4) ensure the methodology is scalable at the production level and can easily incorporate new data.

Introduction

Synthetic population is used to combine socioeconomic and demographic data from multiple sources, such as census and market research, and downscale them to an individual level. Often for privacy reasons, sensitive information such as name, personal contact detail, family history, wealth and health records are considered personally-identifiable, and hence in many jurisdictions, releasing such information is strictly prohibited. As an alternative, however, such data can be released at aggregated regional levels (e.g., only averages or percentages are released for a region with multiple residents). However, practitioners often find individual level data far more appealing, as aggregated data lacks information such as variances and distributions of residents within that region, and therefore an alternative to real population

data is required. For the downscaled synthetic population to be useful, it needs to be fair and consistent. The first condition means that simulated data should mimic realistic distributions and correlations of the true population as closely as possible. The second condition implies that when we aggregate downscaled samples, the results need to be consistent with the original data. A more rigorous analysis is provided in the later section.

Synthetic data generation draws lots of attention from data scientists and is often seen as a privacy-preserving way to augment training data in situations where data collection is difficult. In applications where large-scale data collection involves manual surveys (e.g., demographics), or when the collected data is highly sensitive and cannot be fully released to the public (e.g., financial or health data), synthetically generated datasets become an ideal substitute. For example, due to the Pseudoanonymisation article of the General Data Protection Regulation (2014), organizations across the world are forbidden to release personally identifiable data. As a result, such datasets are often anonymized and aggregated (such as geographical aggregation, where variables are summed or averaged across a certain region). Being able to reverse-engineer the aggregation, therefore, is a key step to reconstruct the lost information.

Common techniques for synthetic population generation are synthetic reconstruction (SR) (1996) and combinatorial optimization (CO) (2001; 2000). Both approaches have specific data requirements and limitations which usually cannot be easily resolved. To address these challenges, we propose a new framework called **SynC** (**S**ynthetic Population with **G**aussian **C**opula) to simulate microdata by sampling features in batches. The concept is motivated by (Jeong et al. 2016) and (Kao et al. 2012), which are purely based on copula and distribution fitting. The rationale behind our framework is that features can be segmented into distinct batches based on their correlations, which reduces the high dimensional problem into several sub-problems in lower dimensions. Feature dependency in high dimensions is hard to evaluate via common methods due to its complexity and computation requirements, and as such, Gaussian copulae, a family of multivariate distributions that is capable of capturing dependencies among random variables, becomes an

*Corresponding Author

ideal candidate for the application.

In this study, we make the following contributions:

1. We propose a novel combination framework which, to the best of our knowledge, is the first published effort to combine state-of-the-art machine learning and statistical instruments (e.g., outlier detection, Gaussian copula, and predictive models) to synthesize population data.
2. We demonstrate SynC’s value as a feature engineering tool, as well as an alternative to data collection in situations where gathering is difficult through two real-world datasets in the automotive and market research industries.
3. To foster reproducibility and transparency, all code, figures and results are openly shared¹. The implementation is readily accessible to be adapted for similar use cases.
4. We ensure the methodology is scalable at the production level and can easily incorporate new data without the need to retrain the entire model.

Related Works

Synthetic Reconstruction

Synthetic reconstruction (SR) is the most commonly used technique to generate synthetic data. This approach reconstructs the desired distribution from survey data while constrained by the marginal. Simulated individuals are sampled from a joint distribution which is estimated by an iterative process to form a synthetic population. Typical iterative procedures used to estimate the joint distribution are iterative proportional fitting (IPF) and matrix ranking. The IPF algorithm fits a n -dimensional contingency table based on sampled data and fixed marginal distributions. The inner cells are then scaled to match the given marginal distribution. The process is repeated until the entries converge.

IPF has many advantages like maximizing entropy, minimizing discrimination information (1968) and resulting in maximum likelihood estimator of the true contingency table (1991). However, IPF is only applicable to categorical variables. The SynC framework incorporates predictive models to approximate each feature, which can be used to produce real-valued outputs as well and probability distribution that can be sampled from to produce discrete features.

Combinatorial Optimization

Given a subset of individual data with features of interest, the motivation behind combinatorial optimization (CO) is to find the best combination of individuals that satisfy the marginal distributions while optimizing a fitness function (2013). CO is typically initialized with a random subset of individuals, and individuals are swapped with a pool of candidates iteratively to increase the fitness of the group. Compared to SR approaches, CO can reach more accurate approximations, but often at the expense of exponential growth of computational power (2013).

Copula-Based Population Generation

Copula is a statistical model used to understand the dependency structures among different distributions (refer to *Proposed Framework* Section for details), and has been widely used in microsimulation tasks (2012). However, downscaling is not possible, and the sampled data stay at the same level of granularity as the input. Jeong et al. discuss an enhanced version of IPF where the fitting process is done via copula functions (2016). Similar to IPF, this algorithm relies on the integrity of the input data, which, as discussed before, can be problematic in real-world settings. Our method, SynC, is less restrictive on the initial condition of the input dataset as it only requires aggregated level data. Therefore SynC is more accessible compared to previous approaches.

Proposed Framework

Problem Description

Throughout this paper, we assume that the accessibility to $X = X_1, \dots, X_M$, where $X_i = [X_i^1, \dots, X_i^D]^T$ is a D^{th} dimensional vector representing input features. Each $m \in M$ represents an aggregation unit containing n_m individuals. When referring to specific individuals, we use $x_{m,k}^d$ to denote the d^{th} feature of the k^{th} individual who belongs to the aggregation unit m . For every feature $d \in D$ and aggregation unit m , we only observe X_m^d at an aggregated level, which is assumed to be an average (i.e. $X_m^d = \sum_{k=1}^{n_m} x_{m,k}^d$) in the case of numerical measurements, or a percentage (i.e. $X_m^d = \sum_{k=1}^{n_m} x_{m,k}^d / n_m$) in the case of binary measurements. We use the term *coarse data* to refer to this type of observations. In applications, aggregation levels can be geopolitical regions, business units or other types of segmentation that make practical sense.

Our proposed framework, SynC, attempts to reverse-engineer the aggregation process to reconstruct the unobserved $x_1^d, \dots, x_{n_m}^d$ for each feature d and aggregation level m . We assume that M is sufficiently large, especially relative to the size of D , so that fitting a reasonably sophisticated statistical or machine learning model is permissible, and we also assume that the aggregation units, n_m , are modest in size so that not too much information is lost from aggregation and reconstruction is still possible. Thus for a $M \times D$ dimensional coarse data X , the reconstruction should have dimensions $N \times D$, where $N = \sum_{m=1}^M n_m$ is the total number of individuals across all aggregation units. This finer-level reconstruction is referred to as the *individual data*.

SynC is designed to ensure the reconstruction satisfies the following three criteria (2003):

- (i) For each feature, the marginal distribution of the generated individual data should agree with the intuitive distribution one expects to observe in reality.
- (ii) Correlations among features of the individual data need to be logical and directionally consistent with the correlations of the coarse data.
- (iii) Aggregating generated individual data must agree with the original coarse data for each $m \in M$ a.

¹<https://github.com/winstonll/SynC>

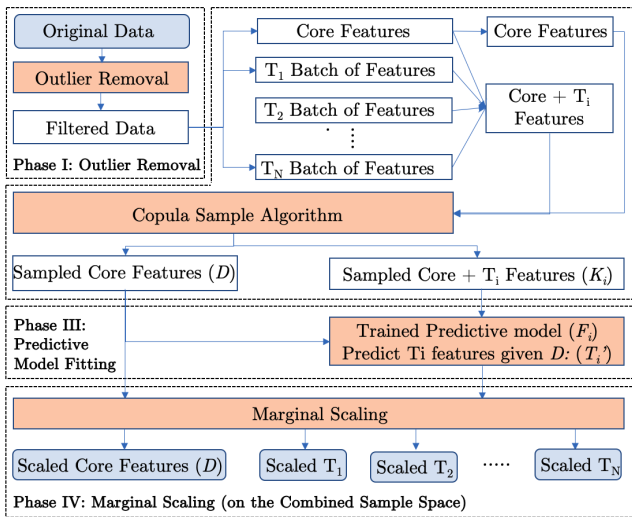


Figure 1: Flowchart of the SynC framework

To illustrate the importance of these criteria, suppose that one wishes to reconstruct individual incomes from a given coarse personal incomes. The first criterion implies that the distribution of sampled individual incomes must align with the observed distribution of income, which is often modelled by the lognormal distribution. The second criterion implies that generated incomes should correlate with other money-related features such as spendings on vacations, luxury goods, and financial investments. The third criterion implies that if we average the generated incomes, we should get the input coarse personal incomes. The main idea of SynC is to simulate individuals by generating features in batches based on core characteristics and scaling the result to maintain hidden dependency and marginal constraints. As illustrated in Fig. 1, SynC contains four core phases and their formal description is provided below.

Phase I: Outlier Removal

Outliers are the deviant samples from the general data distributions that may be a result of recording mistakes, incorrect responses (intentionally or unintentionally) or tabulation errors (2019). The presence of outliers can lead to unpredictable results (2019). Microsimulation tasks are sensitive to outliers (2013), and conducting outlier detection before any analysis is therefore important. Outlier removal methods are often selected by the underlying assumptions and the actual condition of the data. Notably, the necessity of this process depends on the practitioners' definition of abnormality, although it is recommended in most of the cases.

Phase II: Dependency Modeling

We propose to use the copula model to address criteria i) and ii) since copulae, with its wide industrial applications, have been a popular multivariate modeling methodology especially when the underlying dependency is essential. First introduced by Sklar (1959), a copula is a multivariate probability distribution where the marginal probability distribution of each variable is uniform. Let $X = (x_1, x_2, \dots, x_D)$ be

Algorithm 1: Gaussian copula sampling

Data: Coarse Data

Result: Simulated Individual Data

initialization;

M = number of aggregated unit

D = dimension of coarse data

Γ = D -dimensional correlation matrix of features of coarse data

Φ = CDF of a standard normal distribution

F_d^{-1} = inverse CDF of the marginal distribution of the d^{th} feature

for m in $1 \dots M$ **do**

Draw $Z_m = Z_m^1, \dots, Z_m^D \sim N(0, \Gamma)$, where $N(\mu, \Gamma)$ denotes a d -dimensional Normal distribution with mean μ and correlation matrix Γ

for d in $1 \dots D$ **do**

$u_m^d = \Phi(Z_m^d)$

$y_m^d = F_d^{-1}(u_m^d)$

(This implies that Y_m^d follows the desired distribution)

end

Return $Y_m = \{Y_m^1, \dots, Y_m^D\}$

end

Return $Y = \{Y_1, \dots, Y_M\}$

a random vector in R^D , and the marginal cumulative distribution function be $P_i(x) = Pr[x_i < x]$, define U_i as

$$U = (u_1, u_2, \dots, u_D) = (P_1(x_1), P_2(x_2), \dots, P_D(x_D)) \quad (1)$$

A copula of the random vector X is defined as the joint CDF of a random uniform vector U :

$$C(u'_1, u'_2, \dots, u'_D) = Pr(u_1 < u'_1, u_2 < u'_2, \dots, u_D < u'_D) \quad (2)$$

In other words, we can describe the joint distribution of a random vector X using its marginal distributions and some copula functions. Additionally, Sklar's Theorem states that for any set of random variables with continuous CDFs, there exists a unique copula as described above. It allows us to isolate the modeling of marginal distributions from their underlying dependencies. Sampling from copulae is widely used by empiricists to produce desired multivariate samples based on a given correlation matrix and the desired marginal distributions of each of the components. Nelsen (2007) outlines a simpler version of Algorithm 1 for bivariate sampling, which can be generalized to multivariate cases.

In order to properly specify F_d^{-1} , we make a reasonable assumption that the size of each aggregation unit is significant and diverse enough such that $var(X_m^d)$ is approximately constant $\forall m$. This assumption implies that $var(X_m^d)$ can be estimated by $(\sum_{k=1}^m (X_k^d)^2 - \bar{X}^d) / (m - 1)$, the unbiased sample variance estimator. Thus, given μ_m^d is observed from aggregation unit averages, F_d^{-1} can be uniquely specified so long as we are willing to make an assumption on the distribution. For example, if the desired output is a positive continuous variable (such as income), $F_{Y_m^d}(y)$ follows a lognormal distribution with mean μ_m^d and standard deviation σ_m^d . In the case of categorical variables, $F_{Y_m^d}(y)$ follows a beta distribution with parameters α and β such that

Algorithm 2: Batch sampling via Gaussian copula

Data: Coarse Data

Result: Simulated Individual Data

initialization;

X = Input coarse data

B = total number of batches of non-core features

S = predefined set of core features

$T_i = i^{th}$ batch of non-core features

Y = Initial sampled individual data with only core features using Algorithm 1

for i in $1..B$ **do**

$X'_i = X[S \cup T_i]$

$K_i =$ Sampled data by applying Algorithm 1 to X'

$F(T_i | S = s) =$ an approximated distribution function trained on K_i

$Q_i = \arg \max_Q F(Q | S = Y)$

$Y = Y \bowtie Q_i$ where \bowtie is the natural join operator

end

Return Y

$\frac{\alpha}{\alpha+\beta} = \mu_m^d$ and $\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)} = (\sigma_m^d)^2$. Our assumptions imply that μ_m^d and σ_m^d can be derived from coarse data mean and standard deviation for feature d .

$$\mu_m^d = \text{mean of feature } d \text{ in aggregation unit } m \quad (3)$$

$$\sigma_m^d = \sigma^d * \sqrt{M} * \sqrt{n_m} \quad (4)$$

Algorithm 1 incorporates the above assumptions with Gaussian copula to satisfy criteria i) and ii).

Gaussian Copula Being one of the most popular and simple form of copulae, Gaussian Copula is easy to interpret, implement and sample (and will be the copula used in section 4). It is constructed from multivariate normal distribution through probability integral transformation. For a given covariance matrix \mathbb{R} , Gaussian Copula requires the joint distribution of given variables can be expressed as:

$$C_R^{\text{Gauss}}(u) = \Phi_R(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d)) \quad (5)$$

where Φ is the CDF of normal distribution.

$$V = (\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d)) \quad (6)$$

$$c_R^{\text{Gauss}}(u) = \frac{1}{\sqrt{\det R}} \exp\left(-\frac{1}{2} (V \cdot (R^{-1} - I) \cdot V^t)\right) \quad (7)$$

Because of the popularity of the normal distribution and the other advantages mentioned above, Gaussian copula has been the most applied copula. The earliest application of Gaussian Copula was in the finance industry: Frees and Valdez applied it in insurance (1998), and Hull and White used it in credit derivative pricing (2006). Recently, Gaussian Copula's application can be found in many fields such as linkage analysis (2006); functional disability data modeling (2011); power demand stochastic modeling (2012), etc.

Archimedean Copula

After the work of Abel (1826) and Ling (1965), Archimedean Copula is introduced with the definition:

$$C(u_1, \dots, u_p) = \psi\left(\sum_{i=1}^p \psi^{-1}(u_i)\right) \quad (8)$$

where the $\psi : [0, 1] \rightarrow [0, \infty)$ is a real valued function satisfying a few conditions¹.

With different ψ function, Archimedean Copula have many extensions. The popular cases are Clayton's Copula which has been used in bivariate truncated data modelling (2007) and hazard scenarios and failure probabilities modelling (2016), and Frank's Copula, which has been used in storm volume statistics analysis (Salvadori and De Michele 2004) and drought frequency analysis (Wong 2013).

Archimedean Copulae have gained some attention in recent years and their unique generator functions have outstanding performances in particular situations. Yet the fitting process for the generator function can be computational expensive. As SynC preforms data generation in batches and wishes to be computational efficient, we will use Gaussian copula for modeling the dependency among features.

Phase III: Predictive Model Fitting

In theory, one could construct a complex enough correlation matrix with all variables, and then fit a giant copula model using Algorithm 1. However, two reasons make this approach nearly impossible: 1) the curse of dimensionality makes it difficult when D is sufficiently large, and 2) certain columns of X may update more frequently, and it is inefficient to train new models each time. In this section, we introduce an alternative method, Algorithm 2, to resolve this issue with minimal computation requirements.

From X , we select, $S \in X^D$, a subset of X containing the utmost important features in the given context. Next, we apply Algorithm 1 to S and the resulting data set, Y , should satisfy the first two criteria mentioned in **Phase II**.

One crucial question is the choice of *core features*, as it is somewhat arbitrary to decide which variables should be considered core features. The choice should be made specifically for every application, as in some cases the importance and influence of some variables clearly dominate others, while in other cases quality (e.g., sampling method) or availability (e.g., sample size) of data makes some variables better candidates for inclusion in the core feature set.

Secondly, divide the non-core variables into a total of B batches, $T_1 \dots T_B$, based on their correlations. Again, the division is somewhat arbitrary but in most cases, they can be based on content, method of collection or other groupings that make sense to the practitioners. Together with S , all T_i are inputted into Algorithm 2 sequentially. The resulting sampled data set, K_i , contains features from S and T_i .

However, a problem immediately presents: the sampled individuals in K_i do not match Y . Furthermore, individuals sampled in each K_i does not necessarily match the sampled individuals in K_j (for $j \neq i$). To fix it, we can train a predictive model on K_i to approximate the distribution of $P(T_i | S = s)$, and use the model to predict the values of features K_i given their values of the core features from X' . The choice of the predictive model depends on the complexity and nature of the data on a case by case basis.

¹ $(-1)^k d^k \psi(x) / d^k x \geq 0$ for all $x \geq 0$ and $k = 1, \dots, p-2$, and $(-1)^{p-2} \psi^{p-2}(x)$ is non-increasing and convex.

Postal	# Population	Avg Age	% with Mortgage	% Speaks two languages
M5S3G2	467	35.1	0.32	0.69
V3N1P5	269	37.2	0.35	0.67
L5M6V9	41	49.1	0.67	0.43

Table 1: An excerpt of three variables from the census data

Finally, we merge the predicted values with the original data set Y and iterate until all features are covered.

Phase IV: Marginal Scaling

The final step is to address criterion iii), which is to ensure sampled individual data agree with the input coarse data.

If Y^d is categorical with η classes, we first note that the predicted values from **Phase III** for individual k in aggregated unit m , represented by $\mathbf{p}_{m,k}^d = \{p_{m,k,i}^d\}_{i=0}^{\eta}$, is a probability distribution. Hence it is natural to assume $Y^d \sim \text{Multi}(1, \mathbf{p}_{m,k}^d)$. To determine the exact class of individual k , we generate a random sample from the distribution. After initial sampling, the percentage of each category may not match the marginal constraint. To resolve this, SynC first randomly removes individuals from the categories that are over-sampled until their marginal constraints are satisfied. SynC then resamples the same amount of individuals using the normalized probability of the under-sampled categories. It is noted that the above processes are iterated until the desired result is achieved.

If Y^d is continuous, the sampled mean and variance should be in proximity with the original coarse data given the way F_d^{-1} is constructed in **Phase II**. In case of small discrepancy, SynC will horizontally shift each data point by the difference between sample mean and the coarse data.

Results and Discussion

In this section, we demonstrate the quality of SynC using two real-world examples. In both cases, our coarse data comes from the 2016 Canadian National Census, which is collected by Statistics Canada and compiled once every five years. The census is aggregated at the postal code level and made available to the general public. There are 793,815 residential postal codes in Canada (in the format L#L#L#, where L is a letter and # is a digit), with an average of 47 residents per postal code. The dataset contains more than 4,000 variables ranging from demographics, spending habits, financial assets, and social values. Table 1 illustrates a subset of this dataset with 3 postal codes and 4 variables.

Demographics variables from the Census is chosen as the core features because this subset is by far the most comprehensive portion of all parts. All Canadian residents are required to complete this section, while other sections are optional and often complemented by other sources such as private market research data (usually much smaller in sample size and coverage).

In the following two case studies, we validate:

1. the improvements on model accuracy when SynC is used as a feature engineering tool when training data is limited,

Postal	Sex	Age	Ethnicity	Immigration Status	Education	Profession	Marital Status	Family Size	Income	Languages Spoken
V3N1P5	F	19	Latin	Immigrants	No degree	Ed services	Married	5+	<\$10k	1
V3N1P5	F	65+	Chinese	Immigrants	No degree	Food services	Widowed	3	\$10k to \$19k	2
V3N1P5	M	51	Korean	Immigrants	College	Waste management	Separated	1	<\$10k	2
V3N1P5	M	60	Korean	Immigrants	Master	Finance	Married	2	<\$10k	2

Table 2: An excerpt of simulated data for one postal region

PostalCode	Gender	PersonAge	old car VOI	Dealer where old car was purchased	date first email sent	date last email sent	unsold while in LYOL flag Y/N	Finished Full Cadence	Lease terminated while in LYOL	Sold while in LYOL	Lease renewed flag Y/N	new Purchase Date	new car signature	dealer where new car was purchased
V3N1P5	M	53	XXX SUV 3	Brian Jessel Dealership	4/9/2018	10/18/2018	N	Y	Y		Y	7/10/2018	Lease	XXX Laval
H7T1T4	M	68	XXX SUV 4	XXX Dealership in Laval	2/20/2018	6/20/2018	N							
L9X0S4	M	45	XXX Sedan 3	Georgian XXX	5/6/2019	7/8/2019	N	Y	Y					
H9R1A5	F	45	XXX Sedan 3	XXX Canbec	1/17/2019	7/24/2019	N	Y						
L4S1W5	M	52	XXX SUV 3	XXX Autobus	3/29/2018	10/11/2018	N	Y			Y	10/31/2018	XXX SUV 3	Lease
H7P0B9	F	43	XXX Sedan 3	XXX Laval	4/1/2019	6/25/2019	N							

Table 3: An excerpt of auto leasing data

- the quality of SynC by comparing a sample of the generated individual data against real residents in these locations, and therefore leveraging it as an alternative to traditional market research.

SynC as a Feature Engineering Tool

To assess the performance of SynC as a feature engineering tool, we collaborate with a global automotive company (hereafter referred to as the "Client") that specializes in producing high-end cars to build a predictive model to better assist their sales team in identifying which of their current customers who have a leased vehicle are interested in buying the car in the next 6 months. This type of analysis is extremely important in marketing, as contacting customers can be both expensive (e.g. hiring sales agents to make calls) and dangerous (e.g. potentially irritating and leading to unsubscriptions of emails or services). Therefore building accurate propensity models can be extremely valuable.

Our experiment contains three steps. 1) We work with the Client to identify internal sales data and relevant customer profiles such as residential postal code, age, gender, 2) we apply *probabilistic matching* to join sampled data together with the Client's internal data to produce an augmented version of the training data, and 3) we run five machine learning models on both the original and the augmented data, and evaluate the effectiveness of SynC.

Data Description There are 7,834 customers who currently lease a specific model of car from the Client in Canada. Our Client is interested in predicting who are more likely to make a purchase in the next 6 months. In Table 5, we attach an excerpt of the Client's sale records. For security reasons, names and emails are removed.

To augment this dataset, the Client selects 30 variables from the Census, which included information they do not know but could be useful features. The variables includes, demographics (personal and family), financial situations and how they commute to work. We include an excerpt of the sampled individual data obtained by apply SynC to the Census in Table 5; both datasets are available upon request.

Probabilistic Matching One challenge of using SynC as a feature engineering tool is the fact that synthetic population is anonymous. In most applications, enterprise level data sources are individually identifiable through an identifier which may be unique to the company (e.g. customer ID for multiple products/divisions at one company) or multiple companies (e.g. cookie IDs which can be shared between multiple apps/websites). This makes merging different data sources very easy as the identifier can be used as the primary key. By construct, however, synthetic population individuals are model generated from anonymous data, and therefore cannot be joined by traditional means. Here we present an alternative method called *Probabilistic Matching*.

Because SynC produces anonymous information (i.e. data that cannot be attributed to a specific person, such as name, email address or ID), we use age, gender, ethnicity, and profession as good identifiers to the real population. In table 3 we show the first few customers supplied by our industry partner. We also provide the list of synthetically generated

	LR	DT	RF	SVM	NN
Original Data	61.5%	63.9%	70.4%	69.3%	68.8%
Augmented Data	66.2%	71.1%	73.0%	80.6%	73.9%

Table 4: Comparisons of accuracy measures of 5 different classifiers trained on original data and synthetic population augmented data

persons for postal code V3N1P5 in Table 2, and use this as an example to demonstrate how probabilistic matching can be done on the first customer.

Client data show that a 53 year old male, who lives in the area of V3N1P5 made a lease purchase. In this case, we have three indicative measurement for this customer - this buyer is *53 years old, male* and lives in *V3N1P5* (an area in Burnaby, British Columbia, Canada). In our synthetic data, the closet match would be the tenth person, as two of the three indicators (postal code and genders) match precisely, and the last indicators matches closely (age of 54 vs. 53, which is a difference of 1 year). We can conclude that this customer, who leased an SUV of model type 3, is likely to be ethnically Chinese, an immigrant with a bachelor's degree, an income of between \$90k to \$99k and speaks 2 languages.

Notably, the third person is also a very close match. Both postal code and gender matches perfectly, and age matches closely as well (51 vs. 53, which is a difference of 2 years rather than 1 year). It would also make sense to set a threshold (such as within 5 years difference in age), and aggregate all matched records via averaging (for numerical variables) or voting/take proportions (for categorical variables).

Method Evaluation We train 5 different classifiers on both the partner's data, as well as the augmented dataset to predict whether a customer buys the leased vehicle. We train Logistic Regression (LR), Decision Tree (DT), Random Forest with 500 trees (RF), SVM with Radial Basis Kernel and 2-Layer Neural Network (NN), and their performances can be found in Table 4.

In all five cases, augmented data produces a higher classification accuracy than the raw data from our industry partner. Accuracy increases range from slightly over 2.5% (RF) to as much as 11% (SVM), with an average increase of 6.2%. This increase is both technically significant, as well as practically meaningful, as the Client would easily apply this model to their business and achieve grow their sales.

This case study has shown that Synthetic Population is an effective way to engineer additional features to situations where the original training data is limited. As explained in early sections, SynC takes coarse datasets and generate estimates of individuals that are likely to reside within the given postal code area. Although SynC does not produce real people, the generated "synthetic" residents both closely resembles the behaviours of true population and is also consistent with the available sources. We demonstrate that it is a viable data augmentation technique.

SynC as an Alternative to Market Research

The validity is hard to assess due to lack of true individual data to benchmark. After all, if individual data was attain-

able then SynC would not be needed in the first place. Therefore, We propose an alternative evaluation metric which requires 1) a set of surveyed response representing the population and 2) simulated populations from the postal region of the surveyed individuals. For a surveyed dataset containing D features from T individuals across M postal codes (each containing n_m surveyed individuals), M simulated populations each with T_m (the population of m^{th} postal region) individuals are generated using the proposed framework.

The evaluation function, \mathcal{L} is defined as the following:

$$\mathcal{L} = \frac{1}{M} \sum_{m=1}^M \frac{1}{n_m} \sum_{k=1}^{n_m} \frac{1}{D} \sum_{d=1}^D \mathbb{I}\{x_{m,k}^d = y_{m,k}^d\} \quad (9)$$

and

$$\{y_{m,k}\}_{k=1}^{n_m} = \arg \max_{\{y_{m,k}\}_1^{n_m}} \sum_{k=1}^{n_m} \sum_{d=1}^D L(x_{m,k}^d, y_{m,k}^d) \quad (10)$$

where

$$L(x, y) = \begin{cases} \mathbb{I}\{x_{m,k}^d = y_{m,k}^d\}, & \text{if } x_{m,k}^d \text{ is categorical} \\ \frac{|x_{m,k}^d - y_{m,k}^d|}{x_{m,k}^d}, & \text{if } x_{m,k}^d \text{ is continuous} \end{cases} \quad (11)$$

Intuitively, \mathcal{L} measures the average similarity across all surveyed postal codes between the person from true population and the set of simulated individuals in each region that resembles them the most.

We collect responses from 30 individuals in Toronto, Canada. For simplicity, the surveyed questions only contain five core features (Age, Sex, Ethnicity, Education, Income), seven additional personal features (Immigration Status, Marital Status, Family Size, Profession etc.) and four spending behaviour feature (Favourite Store for Cloth/Food/Grocery/Furniture). The five demographics features are selected as core variables and the others are grouped into two batches, personal and spending behaviour. Using the SynC framework along with a 2-layered neural network in **Phase III**, a sample is produced and an excerpt is shown in Table 2.

Table 5 presents three comparison pairs between a few surveyed individuals and their simulated counterparts (pairs are highlighted by color). With this sample, $\mathcal{L} = 82\%$. On a closer look it is easy to see that the generated individuals are more realistic than what the score can reflect. The metric function uses an indicator function for each feature to evaluate the accuracy, which does not capture the potential distance among categories (e.g. $Income_{50k-59k}$ should be ‘‘closer’’ to $Income_{60k-69k}$ compare to $Income_{150k+}$).

Conclusion and Future Directions

In this work, we propose a novel framework, SynC, for generating individual level data from aggregated data sources, using state-of-the-art machine learning and statistical methods. To show the proposed framework’s effectiveness and boost reproducibility, we provide the code and data on the Canada National Census example described in *Proposed Method* Section. We also present a real-world business use case to demonstrate its data augmentation capabilities.

Postal	Sex	Age	Ethnicity	Immigration Status	Education	Profession	Marital Status	Family Size	Income	Favorite Store
M2M4L9	M	60	Korean	Immigrants	Master	Finance	Married	2	<\$10k	H&M
M2M4L9	F	58	Other	Non-immigrants	College	Health care	Married	3	\$50k to \$59k	Hot Renfrew
M2M4L9	M	26	Chinese	Immigrants	HighSchool	Retail	Never married	3	<\$10k	Old Navy
M2M4L9	M	<14	Mid Eastern	Immigrants	N/A	N/A	N/A	N/A	N/A	N/A
M2M4L9	F	21	Chinese	Immigrants	Bachelor	Retail	Never married	3	;\$10k	Gap
M2M4L9	M	51	Chinese	Immigrants	No degree	Other services	Never married	2	\$10k to \$19k	Gap

Table 5: An expert of simulated data for one postal region

Type	Postal	Age	Sex	Education	Ethnicity	Family Size	Immigration Status	Marital Status	Income	Profession	Favorite Store
Simulated	M2M4L9	47	F	Bachelor	Chinese	2	Immigrants	Married	\$40k to \$49k	Scientific/Technical	Gap
Surveyed	M2M4L9	49	F	Bachelor	Chinese	3	Immigrants	Married	\$5k to \$59k	Scientific/Technical	H&M
Simulated	M2M4L9	46	M	Bachelor	Chinese	2	Immigrants	Married	<10k	Scientific/Technical	Gap
Surveyed	M2M4L9	51	M	Bachelor	Chinese	3	Immigrants	Married	\$5k to \$59k	Scientific/Technical	Gap
Simulated	M4Y1G3	22	M	Bachelor	Filipino	3	Immigrants	Married	\$70k to \$79k	Scientific/Technical	Banana Rep
Surveyed	M4Y1G3	23	M	Bachelor	Chinese	2	Immigrants	Never married	\$5k to \$59k	Finance	Banana Rep

Table 6: An comparison between simulated samples and surveyed samples

As a first attempt to formalize the problem, we see multiple areas where future works can improve upon. First of all, our method relies on Gaussian copulae and this can be further extended by leveraging other families of copulae to better model the underlying dependency structures. Secondly, we use beta and log-normal distributions to approximate marginal distributions for categorical and continuous variables, respectively, and other families of distributions could be considered (e.g., the κ -generalized model (2016) can be used for money related distributions). Thirdly, a better metric can be designed to measure how accurate SynC generated data is compared to actual population. For example, earth mover's distance can be a better alternative. Fourthly, including more geographical features can potentially increase the accuracy of predicted individual level traits. Lastly, it is valuable to conduct more large-scale analysis with real-world datasets.

References

- Barthelemy, J., and Toint, P. L. 2013. Synthetic population generation without a sample. *Transportation Science* 47(2):266–279.
- Beckman, R. J.; Baggerly, K. A.; and McKay, M. D. 1996. Creating synthetic baseline populations. *Transportation Research Part A: Policy and Practice* 30(6):415–429.
- Clementi, F.; Gallegati, M.; Kaniadakis, G.; and Landini, S. 2016. κ -generalized models of income and wealth distributions. *Review Network Economics* 225:1959–1984.
- Council of European Union. 2014. Council regulation (EU) no 269/2014. <http://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1416170084502&uri=CELEX:32014R0269>.
- Dobra, A.; Lenkoski, A.; et al. 2011. Copula gaussian graphical models and their application to modeling functional disability data. *The Annals of Applied Statistics* 5(2A):969–993.
- Frees, E. W., and Valdez, E. A. 1998. Understanding relationships using copulas. *North American actuarial journal* 2(1):1–25.
- Huang, Z., and Williamson, P. 2001. A comparison of synthetic reconstruction and combinatorial optimisation approaches to the creation of small-area microdata. *University of Liverpool*.
- Hull, J. C., and White, A. D. 2006. Valuing credit derivatives using an implied copula approach. *The Journal of Derivatives* 14(2).
- Ireland, C. T., and Kullback, S. 1968. Contingency tables with given marginals. *Biometrika* 55(1):179–188.
- Jeong, B.; Lee, W.; Kim, D.-S.; and Shin, H. 2016. Copula-based approach to synthetic population generation. *PloS one* 11(8):e0159496.
- Kao, S.-C.; Kim, H. K.; Liu, C.; Cui, X.; and Bhaduri, B. L. 2012. Dependence-preserving approach to synthesizing household characteristics. *Transportation Research Record* 2302(1):192–200.
- Li, M.; Boehnke, M.; Abecasis, G. R.; and Song, P. X.-K. 2006. Quantitative trait linkage analysis using gaussian copulas. *Genetics* 173(4):2317–2327.
- Little, R. J., and Wu, M.-M. 1991. Models for contingency tables with known margins when target and sampled populations differ. *Journal of the American Statistical Association* 86(413):87–95.
- Lojowska, A.; Kurowicka, D.; Papaefthymiou, G.; and van der Sluis, L. 2012. Stochastic modeling of power demand due to evs using copula. *IEEE Transactions on Power Systems* 27(4):1960–1968.
- Münnich, R., and Schürle, J. 2003. On the simulation of complex universes in the case of applying the german microcensus. *DACSEIS research paper series No. 4*.
- Nelsen, R. B. 2007. *An introduction to copulas*. Springer Science & Business Media.
- Passow, B. N.; Elizondo, D.; Chiclana, F.; Witheridge, S.; and Goodyer, E. 2013. Adapting traffic simulation for traffic management: A neural network approach. In *International Conference on Intelligent Transportation Systems*, 1402–1407. IEEE.
- Salvadori, G., and De Michele, C. 2004. Analytical calculation of storm volume statistics involving pareto-like intensity-duration marginals. *Geophysical Research Letters* 31(4).
- Salvadori, G.; Durante, F.; De Michele, C.; Bernardi, M.; and Petrella, L. 2016. A multivariate copula-based framework for dealing with hazard scenarios and failure probabilities. *Water Resources Research* 52(5):3701–3721.
- Sklar, M. 1959. Fonctions de repartition an dimensions et leurs marges. *Publ. inst. statist. univ. Paris* 8:229–231.
- Voas, D., and Williamson, P. 2000. An evaluation of the combinatorial optimisation approach to the creation of synthetic microdata. *International Journal of Population Geography* 6(5):349–366.
- Wang, A. 2007. The analysis of bivariate truncated data using the clayton copula model. *The international journal of biostatistics* 3(1).
- Wong, W.; Wang, X.; and Guo, Z. 2013. Optimizing marker planning in apparel production using evolutionary strategies and neural networks. *Optimizing decision making in the apparel supply chain using artificial intelligence (AI): form production to retail. Woodhead Publishing Series in Textiles* 106–131.
- Wong, G. 2013. A comparison between the gumbel-hougaard and distorted frank copulas for drought frequency analysis. *International Journal of Hydrology Science and Technology* 3(1):77–91.
- Zhao, Y.; Nasrullah, Z.; Hryniewicki, M. K.; and Li, Z. 2019. Lscp: Locally selective combination in parallel outlier ensembles. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, 585–593. SIAM.
- Zhao, Y.; Nasrullah, Z.; and Li, Z. 2019. Pyod: A python toolbox for scalable outlier detection. *Journal of Machine Learning Research* 20(96):1–7.