

MCFlow: Monte Carlo Flow Models for Data Imputation

Trevor W. Richardson Wencheng Wu Lei Lin Beilei Xu Edgar A. Bernal

Rochester Data Science Consortium, University of Rochester
260 E. Main St. Suite 6108, Rochester, NY 14604

{trevor.richardson, wencheng.wu, lei.lin, beilei.xu, edgar.bernal}@rochester.edu

Abstract

We consider the topic of data imputation, a foundational task in machine learning that addresses issues with missing data. To that end, we propose MCFlow, a deep framework for imputation that leverages normalizing flow generative models and Monte Carlo sampling. We address the causality dilemma that arises when training models with incomplete data by introducing an iterative learning scheme which alternately updates the density estimate and the values of the missing entries in the training data. We provide extensive empirical validation of the effectiveness of the proposed method on standard multivariate and image datasets, and benchmark its performance against state-of-the-art alternatives. We demonstrate that MCFlow is superior to competing methods in terms of the quality of the imputed data, as well as with regards to its ability to preserve the semantic structure of the data.

1. Introduction

Missing data is a widespread problem in real-life machine learning problems. Because most of the existing data analysis frameworks require complete datasets, imputation methods are indispensable to practitioners in the field. As a consequence, data imputation has been the focus of extensive research in recent decades [30, 45]. Moreover, much of the recent research in data imputation has leveraged advanced machine learning techniques. This has resulted in the development of numerous shallow [44, 46, 42] and deep learning frameworks [40, 49, 29, 33] which have continuously pushed the state-of-the-art envelope. A pervasive shortcoming of traditional learning-based imputation methods is that their training relies on fully observed data [15, 40]. A more reasonable assumption, however, is that the training data itself may have missing entries, a limitation addressed by recently proposed methods [49, 29, 33]. In this paper, we propose a data imputation framework that leverages deep generative models and incorporates mechanisms that provide more accurate estimates of the miss-

ing data as compared to existing frameworks. In contrast to competing methods that are based on Generative Adversarial Networks (GANs) [49, 29] and on Deep Latent Variable Models (DLVM) [33], our framework leverages normalizing flow models [10, 11, 23]. In particular, we take advantage of the exact log-likelihood evaluation, latent variable inference, and data sample reconstruction afforded by normalizing flow models in order to explicitly learn complex, high-dimensional data distributions.

We address the causality dilemma that arises when attempting to construct probabilistic models of incomplete data by adopting an alternating learning strategy inspired by model-based multiple imputation approaches [31] such as those based on Expectation Maximization [8] and Monte Carlo Markov Chain [43] techniques. Recent examples of related alternating techniques include applications to training generator networks [19, 14]. The overarching idea behind the proposed framework is to alternately sample the data and update the density estimate until a good approximation of the true empirical distribution is attained. This iterative process is enabled by the exact sampling and density evaluation aspects afforded by the normalizing flow model framework referenced above. We note that although accurate density estimation is central to our framework, and autoregressive models achieve state-of-the-art performance in that field [47, 5, 38], their slow sampling time would be a significance hindrance to our algorithm. In order to address the problem of missing data, we introduce a new imputation algorithm called MCFlow. The proposed iterative learning scheme calls for alternately optimizing two different objective functions: (i) the traditional log-likelihood loss involved in training flow models, required to update the density estimate based on the complete data including imputed values; and, (ii) a maximum-likelihood criterion involved in sampling the latent flow space in order to find the optimal values of the missing data given the current estimate of the data distribution. While the former can be achieved by implementing known backpropagation techniques, we introduce a novel non-iterative approach to address the latter,

which relies on a neural network trained to identify points in the latent space that maximize a density value while minimizing the reconstruction error as computed on the observed entries of the data. This approach can be viewed as an instance of an algorithm that learns to optimize [28, 1] and, although not as effective as iterative [22] or sampling-based methods [43, 20], it is significantly more computationally efficient.

The primary contributions of this work can be summarized as follows:

- a framework for data imputation based on deep generative normalizing flow models;
- an alternating learning algorithm that enables accurate density estimation of incomplete, complex and high-dimensional data by leveraging the efficient sampling and density evaluation attributes of flow models;
- a neural network that learns to optimize in the embedding flow space; and,
- extensive empirical validation of the proposed framework on standard multivariate and image datasets, including benchmarking against state-of-the-art imputation methods.

2. Related Work

An imputation method is classified as a single or multiple imputation method depending upon whether it estimates one or more than one value for each missing entry, respectively [30]. Even though a multitude of single imputation methods have been proposed in the literature [44, 42, 50], multiple imputation methods are often preferred as they provide an assessment of uncertainty [30, 45, 35]. Multiple imputation frameworks often rely on building statistical models for the data, and then drawing samples from them in order to perform the imputation. Early attempts at building multiple imputation frameworks relied on simple parametric models such as Bayesian models [3, 41] as well as mixtures of Gaussians [30, 9]. More recently, with the advent of sophisticated deep generative models, the focus has shifted to investigating more effective ways to leverage the expressivity of the models to address data imputation scenarios.

Initial efforts exploiting deep models for imputation relied on the availability of fully observable training data [40]. More recent work has overcome this shortcoming. Two such publications rely on modifications to the Generative Adversarial Network (GAN) architecture [17]. The Generative Adversarial Imputation Network (GAIN) framework [49] employs an adversarially trained imputer optimized at discriminating between fake and real imputations. While the GAN for Missing Data (MisGAN) approach [29] also implements an adversarially trained imputer, it additionally includes an explicit model for the missing data. Unfortunately, being close GAN relatives, these models can be difficult to train [39, 2]. Furthermore, we believe there is an

advantage to explicitly learning a probability density model of the data in imputation tasks, a feature that GAN-based frameworks generally don't afford.

Models that enable approximate density learning rely on estimating the variational bound [25, 40, 4]. While building such models requires observations of complete data, the Missing Data Importance-Weighted Autoencoder (MI-WAE) framework [33] extends the variational lower bound principle to scenarios where only partial observations are available. Unfortunately, such a family of methods is still intrinsically limited to learning an approximation of the density of the data and can therefore be challenging to optimize [24].

In general, models that enable tractable density learning can be classified into frameworks based on fully visible belief networks (FVBNs) and those leveraging nonlinear independent components analysis (ICA). FVBNs rely on the general product rule of probability which enables calculation of the joint distribution of a set of random variables using the product of conditional probabilities [13, 37]. However, given the sequential nature of the operations involved, sampling the density estimate is not efficient. Nonlinear ICA methods define a set of continuous and invertible nonlinear mappings between two spaces [7, 6, 10, 11, 23]. Their main limitation stems from the fact that the mappings need to be invertible, which may limit their expressivity.

3. Framework

Throughout the paper, we consider the scenario in which data is Missing Completely at Random (MCAR) [31]. Formally speaking, assume that fully observable data points $x \in \mathcal{X} \subseteq \mathbb{R}^n$ are distributed according to $p_X(x)$, the set of binary masks $m \in \{0, 1\}^n$ indicates the location of the missing data entries, and the mask entries are distributed according to $p_M(m|x)$. Data is MCAR when the missingness is independent of the data, that is, when $p_M(m|x) = p_M(m)$. Suppose that we observe data point $\tilde{x}^{(i)}$ with corresponding mask $m^{(i)}$, which means that the k -th entry in $\tilde{x}^{(i)}$, denoted $\tilde{x}_k^{(i)}$, is observed if $m_k^{(i)} = 0$ and missing if $m_k^{(i)} = 1$. Often, it will be convenient to work with the complement of mask $m^{(i)}$, which we denote $\bar{m}^{(i)}$. We use the notation $\tilde{x}_{\bar{m}^{(i)}}^{(i)}$ and $\tilde{x}_{m^{(i)}}^{(i)}$, respectively, to denote the set of observed and unobserved entries in $\tilde{x}^{(i)}$ given mask $m^{(i)}$. Note that it is always possible to determine which entries of $\tilde{x}^{(i)}$ are missing and which are present, which means that $m^{(i)}$ can be uniquely determined from $\tilde{x}^{(i)}$. Assuming $\tilde{x}^{(i)}$ is a partial observation of unknown data point $x^{(i)}$ with entries observed according to mask $m^{(i)}$, we formulate the data imputation task of recovering $x^{(i)}$ as a maximum likelihood problem, namely:

$$x^{(i)*} = \arg \max_{x^{(i)}} \left\{ p_X(x^{(i)}) \right\} \text{ s.t. } x_{\bar{m}^{(i)}}^{(i)} = \tilde{x}_{\bar{m}^{(i)}}^{(i)} \quad (1)$$

Because \log is a monotonic function, Eq. 1 is equivalent to

$$x^{(i)*} = \arg \max_{x^{(i)}} \left\{ \log \left(p_X(x^{(i)}) \right) \right\} \text{ s.t. } x_{\bar{m}^{(i)}}^{(i)} = \tilde{x}_{\bar{m}^{(i)}}^{(i)} \quad (2)$$

Note that, in general, $p_X(\cdot)$ is unknown, but even if that were not the case, solving the optimization task from Eqs. 1 and 2 would be challenging given that the distribution of any data of interest is likely to be highly non-convex and high-dimensional. One of the goals of this work is to make this optimization task feasible. To that end, and for the sake of argument, assume a tractable, explicit density model for $p_X(x)$ exists. Let this model be in the form of a generative network G mapping a sample of interest $x \sim p_X(x)$ lying in space \mathcal{X} to an embedding representation $z \sim p_Z(z)$ in space $\mathcal{Z} \subseteq \mathbb{R}^n$. Further, assume that network G effects a continuous, differentiable and invertible mapping $g : \mathcal{X} \rightarrow \mathcal{Z}$ such that $z = g(x)$ and

$$p_X(x) = p_Z(g(x)) \left| \det \left(\frac{\partial g(x)}{\partial x^T} \right) \right| \quad (3)$$

This model is tractable if $p_Z(z)$ is tractable and if the determinant of the Jacobian of $g(\cdot)$ is tractable [18]. Because $g(\cdot)$ is invertible, exact sample generation from density $p_X(x)$ is possible by drawing a sample $z \sim p_Z(z)$ and computing $x = g^{-1}(z)$. Additionally, computing the density on x involves computing the density of its embedding vectors, and then scaling the result by the Jacobian determinant as in Eq. 3 [11]. Transformations that take the form of Eq. 3 fall under the category of nonlinear independent component analysis. A variety of methods of this type have been proposed in the literature [10, 11, 23]. In this paper, we implement network G in the form of a normalizing flow model, and modify existing architectures to support learning from data with missing entries. It is often computationally convenient to work with the log-likelihood function. Taking the log of both sides of Eq. 3 results in

$$\log(p_X(x)) = \log(p_Z(g(x))) + \log \left(\left| \det \left(\frac{\partial g(x)}{\partial x^T} \right) \right| \right) \quad (4)$$

Given a complete set of training data (that is, data without missing entries), and assuming $g(\cdot)$ is parameterized by a set of parameters θ , learning network G corresponds to finding an optimal set of parameters θ^* such that

$$\theta^* = \arg \max_{\theta} \left\{ \log(p_Z(g_{\theta}(x))) + \log \left(\left| \det \left(\frac{\partial g_{\theta}(x)}{\partial x^T} \right) \right| \right) \right\} \quad (5)$$

If network G were available, then the task of data imputation as formulated in Eq. 2 would become feasible. For instance, standard gradient-based optimization techniques could be applied. Because missing data affects the training

set in the scenarios under consideration, however, direct estimation of the distribution of the data according to Eqs. 3 or 4 is not possible.

This scenario constitutes a causality dilemma. Drawing inspiration from alternating algorithms such as Expectation Maximization (EM) [8], MCFlow iteratively fills in the missing data according to Eq. 2 (i.e., based on the current estimate of the density), and updates the generative network parameters according to Eq. 5 (i.e., based on the generated samples). The intuition behind this approach is that finding the values of the missing entries in the data according to the maximum likelihood condition requires knowledge of the data distribution, and learning a model for the data distribution requires complete knowledge of the data. As will become evident later, filling in for missing values in this iterative process involves drawing samples from the current model for the data distribution. In this sense, the MCFlow framework is more closely related to Monte Carlo versions of the EM algorithm [48, 36] than to its vanilla version. In MCEM implementations, the E step consists in generating samples of the latent variables based on the current estimate of the predictive conditional distribution, and the M step estimates the parameters that maximize the observed posterior. This similarity becomes more apparent if we interpret the imputation tasks from Eqs. 1 and 2 as generating samples from the current approximation of the conditional predictive distribution of the missing data given the observed data and the current model parameters, $p_X(x|\theta, \tilde{x})$ (E step). Furthermore, the optimization task from Eq. 5 can be interpreted as updating the current approximation of the posterior of the model parameters given the observed and imputed values, $p(\theta|x, \tilde{x})$ (M step).

4. Model Architecture

The MCFlow architecture utilizes a hybrid framework comprised of a normalizing flow network G (referred to herein as flow network or model) that is trained in an unsupervised manner, and a feedforward neural network that is trained in a supervised manner. The flow network provides an invertible mapping $g_{\theta}(\cdot)$, where θ are the tunable parameters of the network, between data space \mathcal{X} and embedding space \mathcal{Z} and vice versa. The feedforward network H operates in the embedding space by mapping input embedding vectors to output embedding vectors via function $h_{\phi}(\cdot)$, where ϕ are the tunable parameters of the network. In general, the role of the flow model is to learn the distribution of the data. The role of the feedforward network is to find the embedding vector with the largest possible density estimate (i.e., the likeliest embedding vector) that maps to a data vector whose entries match the observed values (i.e., values at locations indexed by the complement of the mask). A high-level overview of the model is illustrated in Fig. 1.

As with traditional flow implementations, training the

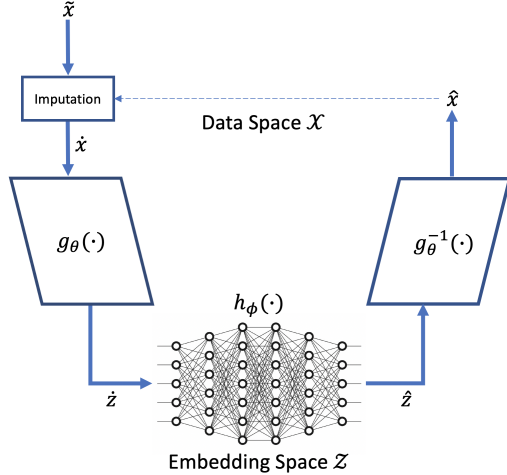


Figure 1. High-level view of the MCFlow architecture.

generative portion of the framework involves finding the set of parameters θ^* that optimizes the mapping $g_\theta(\cdot)$ pursuant to Eq. 5. Because the optimization task from Eq. 5 does not natively support incomplete data, a specific imputation scheme is implemented at initialization. This scheme involves sampling the marginal observed density of the variable that contains the missing data in the case of multivariate, tabular data, and nearest-neighbor sampling in the case of image data. After preprocessing the data as described above, an initial density estimate exists, and the model is now capable of performing data imputation. The estimates for missing values in \tilde{x} are updated during defined subsequent iterations of training with the output of the model, \hat{x} . This is illustrated by the dotted arrow in Fig. 1. The alternating nature of the training process is more clearly illustrated in Fig. 2, which shows how the trained model from the previous epoch is used to update the missing value estimates for use in the current epoch.

The training stage of the generative portion of the framework can be formalized as follows: assume N training samples, $\tilde{x}^{(i)}, i = 0, 1, \dots, N - 1$ with corresponding masks $m^{(i)}$ are available. For every incomplete sample $\tilde{x}^{(i)}$, a full training data sample $\hat{x}^{(i)}$ is computed by combining the observed values with imputed values from imputed sample $\hat{x}^{(i)}$ according to $\hat{x}^{(i)} = \tilde{x}^{(i)} \odot \bar{m}^{(i)} + \hat{x}^{(i)} \odot m^{(i)}$. Here, \odot denotes the Hadamard, or element-wise product between two vectors. With a full training set constructed, learning the optimal set of parameters θ^* in $g_\theta(\cdot)$ is accomplished by minimizing the following cost function:

$$-\frac{1}{N} \sum_{i=0}^{N-1} \log(p_X(\hat{x}^{(i)})) \quad (6)$$

or equivalently, using Eq. 4

$$-\frac{1}{N} \sum_{i=0}^{N-1} \left[\log(p_Z(g_\theta(\hat{x}^{(i)}))) + \log \left(\left| \det \left(\frac{\partial g_\theta(\hat{x}^{(i)})}{\partial \hat{x}^{(i)T}} \right) \right| \right) \right] \quad (7)$$

which is the batch version of Eq. 5. In other words, the optimal parameters of network G are the ones that maximize the log-likelihood of the imputed data, with the imputation mechanism being iteratively updated. The black arrow in Fig. 2 illustrates the direction in which the gradients are backpropagated when training the generative model.

Once an initial density estimate is available, learning the optimal mapping function $h_\phi(\cdot)$ in the embedding space is performed by training feedforward network H . The input to H is the embedding of the imputed training set, namely $\hat{z}^{(i)}, i = 0, 1, \dots, N - 1$, where $\hat{z}^{(i)} = g_\theta(\hat{x}^{(i)})$. Feedforward network H , which maps inputs $\hat{z}^{(i)}$ to outputs $\hat{z}^{(i)}$, is trained by finding the set of parameters ϕ in $h_\phi(\cdot)$ that minimizes the following objective function:

$$\frac{1}{N} \sum_{i=0}^{N-1} \left[\text{MSE}(\hat{x}_{\bar{m}^{(i)}}^{(i)}, \hat{x}_{\bar{m}^{(i)}}^{(i)}) - \lambda \log(p_X(\hat{x}^{(i)})) \right] \quad (8)$$

where $\hat{x}^{(i)} = g_\theta^{-1}(\hat{z}^{(i)})$, $\hat{x}^{(i)} = g_\theta^{-1}(\hat{z}^{(i)})$, $\hat{z}^{(i)} = h_\phi(\hat{z}^{(i)})$, and $\text{MSE}(x, y)$ denotes the mean-squared error operator between vectors x and y . The first cost term in Eq. 8 encourages network H to output an embedding whose reconstruction $g_\theta^{-1}(\hat{z}^{(i)})$ matches the training sample $\hat{x}^{(i)}$ at the observed entries. The second cost term encourages network H to output the vector with the highest density value according to the current density estimate. Both terms combined yield estimates in the form of the likeliest embedding vector that matches the observed values, effectively solving the maximum likelihood objective from Eqs. 1 and 2. The red arrows in Fig. 2 illustrate the computation of the different terms in Eq. 8. The solid segments of the arrows indicate the sections of the framework that contain weights affected by the computations, while the dotted segments indicate where the computations take place and how they are propagated through the framework without affecting any parameters. More specifically, the MSE term is computed in the data space but it only affects weights in neural network H as it is backpropagated. On the other hand, the log-likelihood term is computed in the embedding space of the flow model G and is used to update the weights in H .

We note that optimizing the cost function from Eq. 8 requires repeated log-likelihood evaluation (e.g., in the computation of $\log(p_X(\hat{x}))$), latent variable inference (e.g., in the computation of \hat{z} from \hat{x}), and data sample reconstruction (e.g., in the computation of \hat{x} from \hat{z}), all of which can be effectively computed with flow network G . While other generative models may outperform flow approaches at one task or another, we found flow models to constitute the best compromise for the needs of our framework. Pseudocode

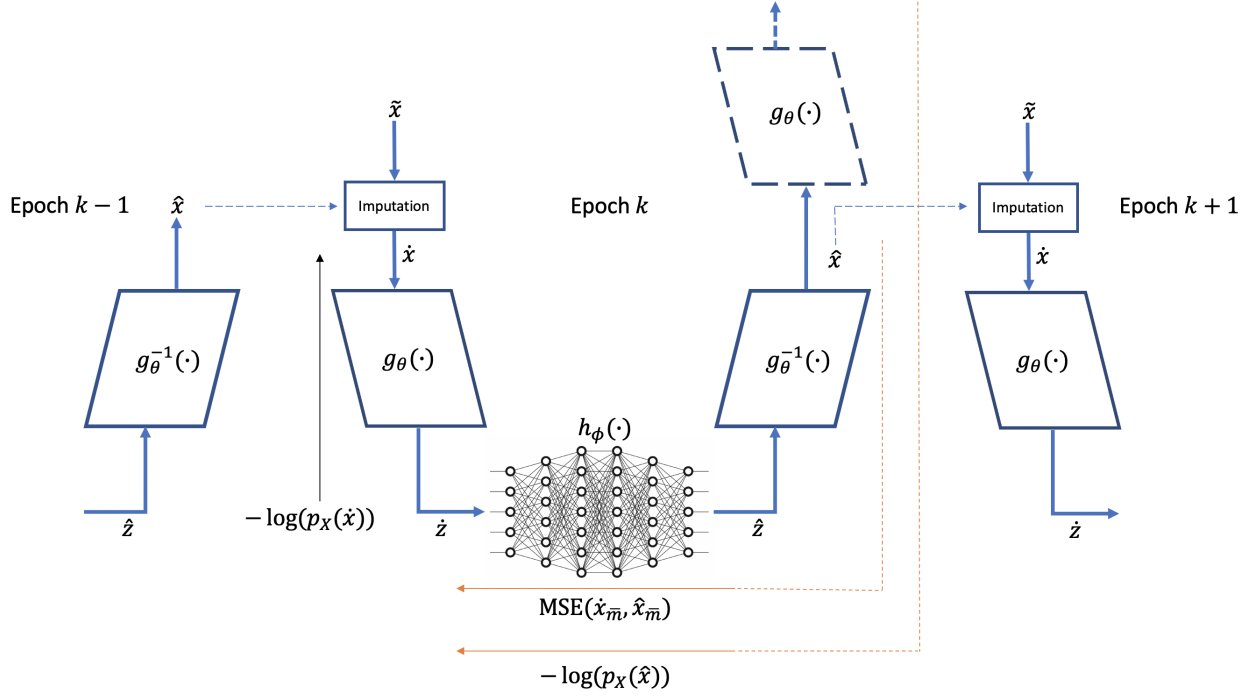


Figure 2. Unrolled view of the architecture and backpropagation process.

for the training procedure is provided in Algorithm 1. Note that we concurrently update the parameters of G and H in order to achieve some computational savings: updating G first and then H would require two separate forward passes through G (one to update G itself and one to compute the mappings involved in updating H).

Algorithm 1 Training Procedure

Data: N data points $\tilde{x}^{(i)}$, $i = 0, 1, \dots, N - 1$ with corresponding masks $m^{(i)}$ indicating location of observed and missing entries

Naively impute missing data in \tilde{x} to produce \hat{x} :

$$\hat{x} \leftarrow \tilde{x} \odot \bar{m} + \hat{x} \odot m$$

for $n = 1$ to $nEpochs$ **do**

 Forward Pass:

$$\hat{z} \leftarrow g_{\theta}(\hat{x})$$

$$\hat{z} \leftarrow h_{\phi}(\hat{z})$$

$$\hat{x} \leftarrow g_{\theta}^{-1}(\hat{z})$$

$$\hat{z} \leftarrow g_{\theta}(\hat{x})$$

 Backpropagation:

 Compute loss according to Eq. 5

 Compute loss according to Eq. 8

 Update θ and ϕ by backpropagating loss

$$\hat{x} \leftarrow \tilde{x} \odot \bar{m} + \hat{x} \odot m$$

Once the model is trained, MCFLOW imputes missing values according to Algorithm 2. The procedure involves encoding the naively imputed data sample \hat{x} into \hat{z} , forward

propagating \hat{z} through the feedforward neural network to obtain \tilde{z} , and reconstructing data sample \hat{x} by decoding \tilde{z} . The final imputed sample \hat{x} is obtained by filling in the missing entries with corresponding entries in \hat{x}

Algorithm 2 Imputation Procedure

Data: Test data point \tilde{x} with corresponding mask m

Naively impute missing data in \tilde{x} to produce \hat{x} :

$$\hat{x} \leftarrow \tilde{x} \odot \bar{m} + \hat{x} \odot m$$

$$\hat{z} \leftarrow g_{\theta}(\hat{x})$$

$$\tilde{z} \leftarrow h(\hat{z})$$

$$\hat{x} \leftarrow g_{\theta}^{-1}(\tilde{z})$$

$$\hat{x} \leftarrow \tilde{x} \odot \bar{m} + \hat{x} \odot m$$

4.1. Implementation Details

The MCFLOW architecture, comprised of various neural network layers, multiple optimizers and competing loss functions, can be implemented in the manner described in this section. A Pytorch implementation of the MCFLOW algorithm is available online.¹ Preprocessing each dataset was done in two steps: 1) initializing $\tilde{x}_{m^{(i)}}^{(i)}$, and 2) scaling the data for training. Each data point $\tilde{x}^{(i)}$ requires the construction of initial estimates for $\tilde{x}_{m^{(i)}}^{(i)}$. This initial step can be performed easily using naive strategies such as zero imputation. We, however, employed two different initialization strategies, both of which estimate initial val-

¹<https://github.com/trevor-richardson/MCFLOW>

ues for $\tilde{x}_{m^{(i)}}^{(i)}$ based on observable values in the dataset. For numerical multivariate datasets, each missing element in $\tilde{x}^{(i)}$ is replaced by sampling the marginal observed density of the variable containing the missing data. For the image datasets, each missing pixel was selected by randomly choosing one entry from the set of nearest observable neighbors of the missing pixel. After completing the selection of values for the missing data points, data was scaled to be in the interval $[0, 1]$ via min-max normalization. For multivariate datasets, only observed values for each variable were used to determine the maximum and minimum values for that variable. For image datasets, the possible maximum of 255 and minimum of 0 were used for each pixel. These steps are required to construct the initial instantiation of \hat{x} .

For all datasets and experiments, the Adam optimizer was used. For multivariate datasets, we used a learning rate of 1×10^{-4} and a batch size of 128. Normalizing flow network G was built using affine coupling layers as introduced by the Real NVP framework [11]. Our implementation of G uses six affine coupling layers and a random masking strategy in contrast to the deterministic masking strategy used in Real NVP. The operations involved in the forward pass of each affine coupling transformation are depicted in Eq. 9:

$$y_D = x_D, \quad (9)$$

$$y_{-D} = x_{-D} \odot \exp(s(x_D)) + t(x_D)$$

In Eq. 9, D ($-D$) represents the set of randomly selected indices which will not (will) be scaled or translated in the current affine transformation. Indices in D were initialized using a binomial distribution with a 50% success rate. These indices remain constant after initialization. In our implementation of G , the s and t functions of every coupling layer are defined by 4-layer fully connected neural networks. Both s and t networks use Leaky ReLu as the activation function in the hidden layers. The final output layer for s and t utilizes the activation functions tanh and linear, respectively.

Network H has five linear layers with the same number of neurons in each layer as the dimensionality of the data. Leaky ReLu was chosen as the activation function between layers. For the image datasets, the following changes were made as a result of resource constraints: First, the batch sizes used for MNIST, CIFAR-10 and CelebA were 128, 128 and 512 respectively; second, for missing data rates above 60%, a learning rate of 1×10^{-3} was employed. A detailed description of the training procedure for MCFlow can be seen in Algorithm 1.

MCFlow periodically updates the missing entries in the training data and resets the θ parameters in the flow model function, $g_\theta(\cdot)$. To that end, we use an exponential scheduling mechanism, where the data update and parameter reset occurs at every epoch that is a power of 2. This implies that in order to perform inference on arbitrary data points,

it is necessary to save the model parameters for both $h_\phi(\cdot)$ and $g_\theta(\cdot)$ after updating \hat{x} according to Eq. 10 but before resetting the model parameters in $g_\theta(\cdot)$:

$$\hat{x} = \tilde{x} \odot \bar{m} + \hat{x} \odot m \quad (10)$$

Based on this exponential scheduling scheme, inference in the MCFlow architecture, as described in Algorithm 2, requires all of the models saved during training. The number of saved models and passes through the architecture is a function of the number of epochs trained, M . More specifically, training MCFlow involves saving the parameters for $h_\theta(\cdot)$ and $g_\theta(\cdot)$ at every epoch that is a power of 2, ultimately requiring $\lceil \log_2(M) \rceil$ saved models in order to properly impute missing data for new samples. The longest the model took to converge was 500 epochs, which required nine saved models to properly inference testing datapoints. Imputing missing data in the test dataset involves performing initial naive imputation followed by full passes through each saved architecture. At the end of this process, the final prediction from MCFlow of x in \hat{x} is returned and performance metrics are recorded.

5. Experimental Results

5.1. Datasets

We evaluate the performance of MCFlow as well as competing methods on three standard, multivariate datasets from the UCI repository [12] and three image datasets. The UCI datasets considered are Default of Credit Card Clients, Letter Recognition and Online News Popularity. The image datasets used in this research are MNIST[27], CIFAR-10 [32] and CelebA [26]. The MNIST dataset contains 28×28 -pixel grayscale images of handwritten digits; we used the standard 60,000/10,000 training/test set partition. CIFAR-10 contains 32×32 -pixel RGB images from ten classes; we used the standard 50,000/10,000 training/test set partition. CelebA contains 178×218 -pixel RGB images of celebrity faces. It was split using the first 162,770 images as the training set and the last 19,962 images for testing. The CelebA images were center cropped and resized to 32×32 pixels using bicubic interpolation.

5.2. Experimental Setup

Imputation performance on multivariate and image data was measured using using root mean squared error (RMSE). We report performance numbers on each test set from models obtained upon convergence of the training process, where convergence is determined based on stabilization of the training losses for both G (Eq. 7) and H (Eq. 8). We emphasize that the training set itself has missing data and that the training loss is computed only on observed data points in order to mimic real-world scenarios as faithfully as possible. Imputation performance on the UCI datasets

Table 1. Imputation Results on UCI Datasets - RMSE (lower is better, 0.2 missing rate)

	Default Credit Card	Online News Popularity	Letter Recognition
MICE	.1763 ± .0007	.2585 ± .0010	.1537 ± .0006
MissForest	.1623 ± .0120	.1976 ± .0015	.1605 ± .0004
Matrix	.2282 ± .0005	.2602 ± .0073	.1442 ± .0006
Auto-Encoder	.1667 ± .0014	.2388 ± .0005	.1351 ± .0009
EM	.1912 ± .0011	.2604 ± .0015	.1563 ± .0012
GAIN	.1441 ± .0007	.1858 ± .0010	.1198 ± .0005
MCFlow	.1233 ± .0012	.1760 ± .0032	.1033 ± .0017

Table 2. Imputation Results on Image Datasets - RMSE (lower is better)

	Missing Rate →	.1	.2	.3	.4	.5	.6	.7	.8	.9
MNIST	GAIN	.11508	.12441	.13988	.14745	.16281	.18233	.20734	.24179	.27258
	MisGAN	.11740	.10997	.11377	.11297	.12174	.13393	.15445	.19455	.27806
	MCFlow	.07464	.07929	.08508	.09187	.10045	.11255	.12996	.15806	.20801
CIFAR-10	GAIN	.10053	.12700	.13248	.11785	.12451	.13130	.13832	.18728	.53728
	MisGAN	.18923	.15223	.14746	.12947	.13027	.14746	.17335	.24060	.31722
	MCFlow	.06012	.06232	.06686	.07215	.08311	.10048	.13132	.15015	.16939
CelebA	GAIN	.06752	.07493	.08367	.08479	.09292	.10608	.11720	.14042	.52050
	MCFlow	.05733	.06243	.06266	.06946	.07261	.07890	.08487	.11073	.12225

Table 3. FID Results on Imputed MNIST Data (lower is better)

Missing Rate →	.1	.2	.3	.4	.5	.6	.7	.8	.9
GAIN	.0696	.4035	1.24	3.277	6.337	12.44	22.91	43.69	92.74
MisGAN	.0529	.1015	.2085	.2691	.3634	.8870	1.324	2.334	6.325
MCFlow	.0521	.0779	.2295	.6097	.8366	.9082	1.951	6.765	15.11

Table 4. Classification Accuracy on Imputed MNIST Data (higher is better)

Missing Rate →	.1	.2	.3	.4	.5	.6	.7	.8	.9
GAIN	.989	.988	.985	.978	.969	.931	.852	.629	.261
MisGAN	.989	.988	.986	.980	.968	.945	.872	.690	.334
MCFlow	.991	.990	.990	.988	.985	.979	.963	.905	.705

was evaluated with a data missingness rate of 20% using five-fold cross validation. We report the mean and standard deviation of the imputation accuracy across all folds for six competing methods: MICE [46], MissForest [42], Matrix Completion (Matrix) [34], Auto-Encoder [16], Expectation-Maximization (EM) [15] and GAIN [49]. For MNIST, CIFAR-10 and CelebA, the testing imputation accuracies are reported across missing rates ranging from 10% to 90% in steps of 10%. Competing methods considered include GAIN [49] and MisGAN [29]. All reported numbers are measured with respect to imputation accuracy over the missing data, namely $\hat{x}_{m^{(i)}}^{(i)}$. Despite considerable efforts to make MisGAN achieve reasonable performance numbers on CelebA, we weren't able to manage it based on the currently published code. No other versions of the code were

available directly from the authors. We also measured the quality of the imputed MNIST images with the Fréchet Inception Distance (FID) [21], which has been shown to correlate well with human perception. Lastly, we measured the ability of the algorithms to preserve semantic content by performing a classification task on imputed imagery with networks pretrained on fully observed data.

5.3. Quantitative Results

Tables 1 and 2 depict MCFlow's performance at MCAR imputation against competing methods. For all datasets and missing rates, the MCFlow framework outperforms all other methods at imputation accuracy from the perspective of the RMSE between the predictions of the model in question and the ground truth values. For the UCI datasets, MCFlow

produces an average reduction in square error of 11% as compared to the existing state of the art, GAIN. In addition, MCFlow produces an average reduction in RMSE of 19%, 38% and 27% as compared to the next best performing method on MNIST, CIFAR-10 and CelebA respectively. Sec. 1 in the Supplementary Material contains imputation results on the training dataset and in terms of PSNR. The results in Table 2 illustrate the pixel-level quality of the imputed imagery; in contrast, Table 3 includes FID performance which is meant to be indicative of quality as perceived by humans. It can be seen that MisGAN outperforms competing methods across most of the missingness range. As will be illustrated in Sec. 5.4, while this means that imputed images more closely resemble attributes of the target image population, they do not necessarily preserve the original semantic content of the partially observed inputs.

In order to illustrate the performance of the methods beyond image quality metrics, we tested their imputation abilities within the context of a larger data processing pipeline where incomplete data is imputed before being processed through a classifier trained on complete data. To that end, we measured the classification performance of a LeNET-based handwritten digit classifier on imputed MNIST data with various degrees of missingness. The LeNET network was pre-trained on MNIST data without missing values. Table 4 contains these results. It can be seen that the imputation results produced by MCFlow have the smallest impact on the semantic content of the imagery, as classification results are consistently higher for images imputed with our method. This phenomenon becomes more evident as the missing data rate increases: the classifier operating on MCFlow-imputed imagery is able to achieve good classification accuracy up to the highest rate of missingness tested, and performs acceptably even in this extreme scenario.

5.4. Qualitative Results

Fig. 3 illustrates the imputation performance of competing methods for the 90% missing data rate on the MNIST dataset. Images along columns (a) and (b) contain the original and observed images, respectively. Pixels that are not observed are assigned a value of 0 for visualization purposes. The imputation models only see the observed images; complete images are included for reference only. Images along columns (c), (d) and (e) include the imputed results after using GAIN, MisGAN and MCFlow, respectively. It can be seen that MCFlow does the best job among the competing methods at preserving and recovering the semantic content of the intended image, which further supports the results from Table 4. Numbers in GAIN-imputed images are, for the most part, illegible. In contrast, MisGAN-imputed images are visually impressive, which is in line with the results from Table 3. One shortcoming of MisGAN, however, is that the imputations produced often

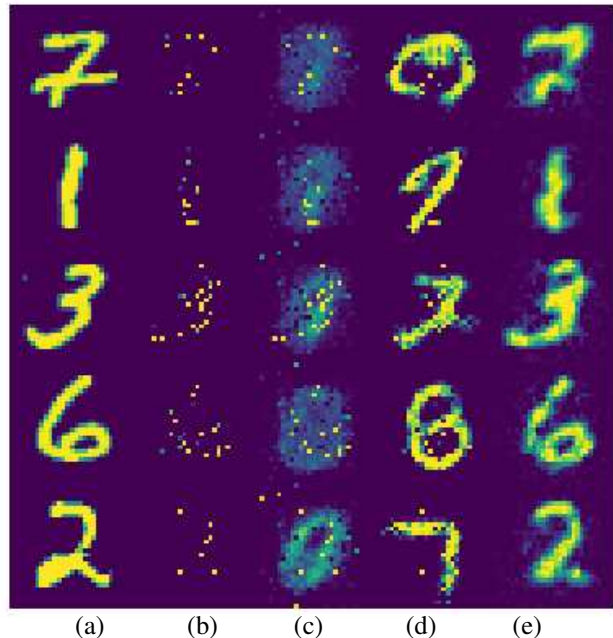


Figure 3. Sample imputation results for MNIST at 0.9 missingness rate: (a) Original, (b) observed, (c) GAIN-imputed, (d) MisGAN-imputed, and (e) MCFlow-imputed images.

fail to represent the digit contained in the original image. Sec. 2 in the Supplementary Material contains additional qualitative results on MNIST and CIFAR-10.

6. Conclusions

We proposed MCFlow, a method for data imputation that leverages a normalizing flow model as the underlying density estimator. We augmented the traditional generative framework with an alternating learning scheme that enables it to accurately learn distributions from incomplete data sets with various degrees of missingness. We empirically demonstrated the superiority of the proposed method relative to state-of-the-art alternatives in terms of RMSE between the imputed and the original data. Experimental results further show that MCFlow outperforms competing methods with regards to preservation of the semantic structure of the data. This is evidenced by the superior classification performance on imputed data achieved by a classifier trained on complete data, which holds across the full range of missing data ratios evaluated. The ability of the method to make sense of and recover the semantic content of the data at every tested missing rate indicates that it is effectively learning the underlying statistical properties of the data, even in extreme paucity scenarios.

Acknowledgements. The authors would like to thank Charles Venuto and Monica Javidnia from the Center for Health and Technology (CHET) at the University of Rochester for their fruitful insight and support.

References

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gomez Colmenarejo, Matthew W. Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. *CoRR*, abs/1606.04474, 2016. [2](#)
- [2] Martín Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. *ArXiv*, abs/1701.04862, 2017. [2](#)
- [3] Vincent Audigier, François Husson, and Julie Josse. Multiple imputation for continuous variables using a bayesian principal component analysis. *Journal of Statistical Computation and Simulation*, 86(11):2140–2156, 2016. [2](#)
- [4] Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *CoRR*, abs/1509.00519, 2015. [2](#)
- [5] XI Chen, Nikhil Mishra, Mostafa Rohaninejad, and Pieter Abbeel. PixelSNAIL: An improved autoregressive generative model. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 864–872, Stockholm, Sweden, 10–15 Jul 2018. PMLR. [1](#)
- [6] Pierre Comon and Christian Jutten. *Handbook of Blind Source Separation: Independent Component Analysis and Applications*. Academic Press, Inc., Orlando, FL, USA, 1st edition, 2010. [2](#)
- [7] Gustavo Deco and Wilfried Brauer. Higher order statistical decorrelation without information loss. In G. Tesauro, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems 7*, pages 247–254. MIT Press, 1995. [2](#)
- [8] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977. [1](#), [3](#)
- [9] Marco Di Zio, Ugo Guarnera, and Orietta Luzzi. Imputation through finite gaussian mixture models. *Comput. Stat. Data Anal.*, 51(11):5305–5316, July 2007. [2](#)
- [10] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *CoRR*, abs/1410.8516, 2014. [1](#), [2](#), [3](#)
- [11] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. 2017. [1](#), [2](#), [3](#), [6](#)
- [12] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. [6](#)
- [13] Brendan J. Frey. *Graphical Models for Machine Learning and Digital Communication*. MIT Press, Cambridge, MA, USA, 1998. [2](#)
- [14] Ruiqi Gao, Yang Lu, Junpei Zhou, Song-Chun Zhu, and Ying Nian Wu. Learning generative convnets via multi-grid modeling and sampling. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 9155–9164. IEEE Computer Society, 2018. [1](#)
- [15] Pedro J. Garcia-Laencina, Jose-Luis Sancho-Gomez, and Anibal R. Figueiras-Vidal. Pattern classification with missing data: A review. *Neural Comput. Appl.*, 19(2):263–282, Mar. 2010. [1](#), [7](#)
- [16] Lovedeep Gondara and Ke Wang. Multiple imputation using deep denoising autoencoders. *ArXiv*, abs/1705.02737, 2017. [7](#)
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014. [2](#)
- [18] Ian J. Goodfellow. NIPS 2016 tutorial: Generative adversarial networks. *CoRR*, abs/1701.00160, 2017. [3](#)
- [19] Tian Han, Yang Lu, Song-Chun Zhu, and Ying Nian Wu. Alternating back-propagation for generator network. In *AAAI*, 2017. [1](#)
- [20] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970. [2](#)
- [21] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017. [7](#)
- [22] C.T. Kelley. *Iterative Methods for Optimization*. Frontiers in Applied Mathematics. Society for Industrial and Applied Mathematics, 1999. [2](#)
- [23] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 10215–10224. Curran Associates, Inc., 2018. [1](#), [2](#), [3](#)
- [24] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4743–4751. Curran Associates, Inc., 2016. [2](#)
- [25] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2013. cite arxiv:1312.6114. [2](#)
- [26] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Cite-seer, 2009. [6](#)
- [27] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. [6](#)
- [28] Ke Li and Jitendra Malik. Learning to optimize. *CoRR*, abs/1606.01885, 2016. [2](#)
- [29] Steven Cheng-Xian Li, Bo Jiang, and Benjamin Marlin. MisGAN: learning from incomplete data with generative adversarial networks. In *International Conference on Learning Representations*, 2019. [1](#), [2](#), [7](#)
- [30] R.J.A. Little and D.B. Rubin. *Statistical analysis with missing data*. Wiley series in probability and mathematical statistics. Probability and mathematical statistics. Wiley, 2002. [1](#), [2](#)

- [31] Roderick J A Little and Donald B Rubin. *Statistical Analysis with Missing Data*. John Wiley & Sons, Inc., New York, NY, USA, 2002. [1](#), [2](#)
- [32] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. [6](#)
- [33] Pierre-Alexandre Mattei and Jes Frellsen. MIWAE: Deep generative modelling and imputation of incomplete data sets. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4413–4423, Long Beach, California, USA, 09–15 Jun 2019. PMLR. [1](#), [2](#)
- [34] Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *J. Mach. Learn. Res.*, 11:2287–2322, Aug. 2010. [7](#)
- [35] Jared Murray. Multiple imputation: A review of practical and theoretical findings. 2018. [2](#)
- [36] Ronald C. Neath. On convergence properties of the Monte Carlo EM algorithm. 2013. [3](#)
- [37] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio, 2016. cite arxiv:1609.03499. [2](#)
- [38] Niki J. Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, and Alexander Ku. Image transformer. 2018. [1](#)
- [39] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks, 2015. [2](#)
- [40] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1278–1286, Beijing, China, 22–24 Jun 2014. PMLR. [1](#), [2](#)
- [41] Yajuan Si and Jerome P. Reiter. Nonparametric bayesian multiple imputation for incomplete categorical variables in large-scale assessment surveys. *Journal of Educational and Behavioral Statistics*, 38(5):499–521, 2013. [2](#)
- [42] Daniel J. Stekhoven and Peter Bühlmann. MissForest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, 28(1):112–118, 10 2011. [1](#), [2](#), [7](#)
- [43] M. Takahashi. Statistical inference in missing data by mcmc and non-mcmc multiple imputation algorithms: Assessing the effects of between-imputation iterations. *Data Science Journal*, 16(37):1–17, 2017. [1](#), [2](#)
- [44] Olga Troyanskaya, Michael Cantor, Gavin Sherlock, Pat Brown, Trevor Hastie, Robert Tibshirani, David Botstein, and Russ B. Altman. Missing value estimation methods for DNA microarrays. *Bioinformatics*, 17(6):520–525, 06 2001. [1](#), [2](#)
- [45] S. van Buuren. *Flexible Imputation of Missing Data. Second Edition*. Chapman and Hall, CRC Press, Boca Raton, FL, 2018. [1](#), [2](#)
- [46] Stef van Buuren and Karin Groothuis-Oudshoorn. Multivariate imputation by chained equations in r. *Journal of Statistical Software*, 45(3):1–67, 2011. [1](#), [7](#)
- [47] Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, pages 1747–1756. JMLR.org, 2016. [1](#)
- [48] Greg C. G. Wei and Martin A. Tanner. A Monte Carlo implementation of the EM algorithm and the poor man’s data augmentation algorithms. *Journal of the American Statistical Association*, 85(411):699–704, 1990. [3](#)
- [49] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. GAIN: Missing data imputation using generative adversarial nets. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5689–5698, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR. [1](#), [2](#), [7](#)
- [50] Zhongheng Zhang. Missing data imputation: focusing on single imputation. *Annals of Translational Medicine*, 4(1), 2016. [2](#)