



## INFORMS Journal on Data Science

Publication details, including instructions for authors and subscription information:  
<http://pubsonline.informs.org>

### Stochastic Neighbourhood Components Analysis

Graham Laidler, Lucy E. Morgan, Nicos G. Pavlidis, Barry L. Nelson

To cite this article:

Graham Laidler, Lucy E. Morgan, Nicos G. Pavlidis, Barry L. Nelson (2025) Stochastic Neighbourhood Components Analysis. INFORMS Journal on Data Science 4(3):248-264. <https://doi.org/10.1287/ijds.2023.0018>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact [permissions@informs.org](mailto:permissions@informs.org).

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2025, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

# Stochastic Neighbourhood Components Analysis

Graham Laidler,<sup>a,\*</sup> Lucy E. Morgan,<sup>b</sup> Nicos G. Pavlidis,<sup>b</sup> Barry L. Nelson<sup>c</sup>

<sup>a</sup>STOR-i Centre for Doctoral Training, Lancaster University, Lancaster LA1 4YW, United Kingdom; <sup>b</sup>Department of Management Science, Lancaster University, Lancaster LA1 4YW, United Kingdom; <sup>c</sup>Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208

\*Corresponding author

Contact: [gllaidler@gmail.com](mailto:gllaidler@gmail.com), <https://orcid.org/0009-0007-6269-6815> (GL); [l.e.morgan@lancaster.ac.uk](mailto:l.e.morgan@lancaster.ac.uk), <https://orcid.org/0000-0003-0987-2350> (LEM); [n.pavlidis@lancaster.ac.uk](mailto:n.pavlidis@lancaster.ac.uk), <https://orcid.org/0000-0002-0301-5350> (NGP); [nelsonb@northwestern.edu](mailto:nelsonb@northwestern.edu), <https://orcid.org/0000-0002-1325-2624> (BLN)

Received: September 13, 2023

Revised: July 18, 2024; October 24, 2024

Accepted: November 11, 2024

Published Online in Articles in Advance:  
May 5, 2025

<https://doi.org/10.1287/ijds.2023.0018>

Copyright: © 2025 INFORMS

**Abstract.** Distance metric learning is a fundamental task in data mining and is known to enhance the performance of various distance-based algorithms. In this paper, we consider stochastic training data in which repeated feature vectors can belong to different classes, a scenario in which existing methods of metric learning are known to struggle. This type of data is common in stochastic simulations, where multidimensional, recurrent system states are subject to inherent randomness. Classification models on such high-resolution simulation-generated data play a critical role in real-time decision making across diverse applications. This paper presents and implements a stochastic version of the popular neighbourhood components analysis. We demonstrate its behaviour on stochastic data using simulation models and reveal its advantages when used for nearest neighbour classification. Meanwhile, the assumptions of stochastic labelling and repeated feature vectors extend to data from various domains, suggesting that the method can attain broad impact. For example, beyond its applications to system control and decision making with digital twin simulation, it may enhance the analysis of data from sensor networks, recommender systems, and crowdsourced platforms, where stochasticity and recurring feature patterns are typical.

**Funding:** This work was supported by the Engineering and Physical Sciences Research Council-funded STOR-i Centre for Doctoral Training at Lancaster University [Grant EP/L015692/1]. In addition, Barry L. Nelson's work was supported by the U.S. National Science Foundation [Grant DMS-1854562].

**Data Ethics & Reproducibility Note:** The code capsule is available on Code Ocean at <https://doi.org/10.24433/CO.0189724.v5> and in the e-Companion to this article (available at <https://doi.org/10.1287/ijds.2023.0018>).

**Keywords:** distance metric learning • stochastic data • discrete-event simulation • simulation analytics • nearest neighbours

## 1. Introduction

Stochastic, discrete-event simulation (DES) is a valuable tool for the design of systems that evolve through time and must perform well in the face of uncertainty. By “design,” we refer to one-time decisions, such as the layout of work centres in a factory, the number of beds allocated to postsurgery patients in a hospital, the securities to include in a portfolio, or the capacities of warehouses in a supply chain. Simulation is less well-suited for real-time control decisions because of long execution times, but is being pushed in that direction by the need for high-resolution, but timely, decisions in applications such as digital twins (dos Santos et al. 2022). A digital twin is a simulation that shadows a real-world system to support control or recovery decision as needed. This tension between simulation execution time and real-world decision time is the genesis of *offline simulation for online application*: extensive simulations are executed in advance and the results processed to be exploited for real-time decisions (Hong and Jiang 2019). This involves the careful mining of simulation state information; we refer to the time-dynamic trajectory of simulation state variables as the simulation *sample path*. Recent applications of sample path mining include financial risk assessment (Jiang et al. 2020), personalised medicine (Shen et al. 2021), and assortment optimisation (Keslin et al. 2024). This work aims to more fully exploit the properties of simulation sample paths to enable higher-resolution predictions than previous methods.

We consider the task of classifying a future system state. Section 4.1 presents an analysis of a queueing system looking at the probability of the system becoming blocked, a commonly encountered issue across many industries. Section 4.2 describes a manufacturing context in which product completion times are classified as early or late at the start of a product cycle, based on simulation-generated completion times from the nearest neighbours

of the current system state. At its highest level of detail, the system state is completely descriptive of the system at any moment in time, and so identifying nearest neighbours on the system state provides a logical basis for predictions (Laidler et al. 2020). For this reason, we propose in this paper a method of distance metric learning (DML) tailored to the characteristics of sample path data.

Extensive research efforts in the field of DML have resulted in a wealth of notable formulations (Kulis 2013). However, existing algorithms are typically built with an implicit assumption of having unique feature vectors for training, an assumption that conflicts with the nature of sample path data. The typically discrete representation and recurrent nature of the system state in a DES model provide repeated feature vector observations, and calculating objective functions as a sum over individual points becomes inefficient. More significantly, inherent simulation randomness results in stochastic classifications. Repeated feature vectors belonging to different classes will often introduce conflicting constraints when existing DML methods are applied and can lead to  $k$  nearest neighbours ( $k$ NN) committing “substantial errors” (Suárez et al. 2021, p. 313). In the current work, we address this shortcoming by proposing a method that we refer to as stochastic neighbourhood components analysis (SNCA). This is based on the neighbourhood components analysis (NCA) of Goldberger et al. (2005), which builds a probabilistic model of class assignment. We modify NCA to the context of repeated data points and stochastic classifications.

Although providing a large family of problems, the simulation domain is not the beginning and the end of our application. Simulation sample paths identify a characteristic type of data that arises in numerous contexts. Crowdsourced data sets, for example, often result in multiple annotators providing various responses to the same labelling task (Vaughan 2018). Tasks such as sentiment analysis or media ratings are subjective in nature and certainly result in stochastic labelling. Beyond this, however, many real-life situations experience naturally stochastic behaviour. Medical prognoses (Suo et al. 2018), financial forecasts (Cao and Tay 2001), and sporting events (Horvat and Job 2020), for example, are often modelled with discrete features and encounter natural variability in their observed outcomes. In summary, repeated feature vectors and stochastic labelling characterise a large body of classification problems, and so we see a DML formulation tailored to this context as a useful contribution.

The remainder of this paper is organised as follows. Section 2 establishes a relevant background in simulation analytics and DML. In Section 3, we give a characterisation of the data and the proposed methodology for SNCA, including attention to the theoretical convergence of the optimal solution. We present experimental results to assess the performance of SNCA on sample path data in Section 4, before the paper concludes with a brief summary in Section 5.

## 2. Related Work

Numerous texts provide comprehensive studies on the subject of stochastic simulation (Law and Kelton 2007, Nelson and Pei 2021). In this section, we focus attention on the emerging topic of simulation analytics (Nelson 2016), which provides a motivation for the current work. We also introduce the task of DML and present a focused review of the existing literature.

### 2.1. Simulation Analytics

Discrete-event simulation represents the operation of a real system as a sequence of events occurring at discrete points in time. Each event, such as a customer arrival or a machine failure, triggers a change in a set of variables describing the system state, including, for instance, a queue size or a machine status. The trajectory of state variables, or sample path, from a DES model hides a wealth of insight into the system behaviour. Nelson (2016) recognised modern capacity for the storage of sample path data as paving the way for machine learning solutions and coined the term *simulation analytics*. Essentially, analysis of simulation has traditionally centred around static summaries of long-run performance, with average performance *over* time being prioritised above dynamic performance *through* time. However, simulation analytics aims to prioritise the latter, recognising that a far more complete picture of system behaviour can be unlocked by directing machine learning efforts to the simulation sample path. The current work adds to this endeavour.

We summarise some existing publications on the theme of simulation analytics. A number of papers have emerged in which sample path data are used to build metamodels for dynamic predictions. In the context of a queueing network, Ouyang and Nelson (2017) propose a two-stage logistic regression modelling approach in which the state and time aspects of the sample path are treated separately, whereas Jiang et al. (2020) use a logistic regression model to dynamically predict the risk of financial portfolios. Morgan and Barton (2022), meanwhile, show that Fourier analysis can successfully detect changes in the trajectories of individual state variables to discriminate between congested and uncongested systems. Moving to distance-based methods, Lin et al. (2019) suggest a  $k$ NN approach to provide performance predictions based on simulation time. A first

investigation of the potential for DML in simulation analytics was provided in Laidler et al. (2020), with the independent variable being extended to a multidimensional description of the system state. However, a tailored DML approach was not offered, and this is instead the main contribution of the current work.

As an accessible source of measurement-error-free data, sample paths provide a rich environment for machine learning. The attempts described above have applied off-the-shelf algorithms. However, to achieve the full potential of learning in this context, it is necessary to account for the distinctive properties of simulation sample path data. This presents an opportunity for machine learning research, and this work takes a step in that direction.

It is encouraging to note that beyond the sphere of research, applications and opportunities for machine learning in simulation are also filtering through to practitioners. Commercial providers of simulation software, including, for example, Simul8 (<https://www.simul8.com>), Simio (<https://www.simio.com>), and AnyLogic (<https://www.anylogic.com>), have recognised its benefits and provide easy integration with machine learning models, including with runtime execution. Crucially, the support is already in place for research advances in simulation analytics to quickly deliver impact to simulation users.

## 2.2. Distance Metric Learning

Distance calculations among data points comprise a fundamental aspect of many machine learning tasks. DML has been shown to improve the performance of nearest neighbour predictions (Weinberger and Saul 2009), clustering (Xing et al. 2002), and information retrieval (McFee and Lanckriet 2010), and it is therefore treated as an important objective in its own right (Kulis 2013). In a fully supervised setting, real-valued, multivariate feature vectors,  $\mathbf{x}_i \in \mathbb{R}^d$ , are supplied with a class label,  $y_i$ , and the goal is to adapt a pairwise distance function over the feature vectors such that nearby points are more likely to belong to the same class. A common approach to DML is to learn a generalised Mahalanobis distance. Our work in this paper conforms to this framework, and this becomes the focus of our review.

Defining the squared distance between  $\mathbf{x}_i$  and  $\mathbf{x}_j$  as  $d_M^2(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i - \mathbf{x}_j)^\top M(\mathbf{x}_i - \mathbf{x}_j)$ , the Mahalanobis distance (Mahalanobis 1936) is defined by setting  $M = \Sigma^{-1}$ , where  $\Sigma$  is the covariance matrix of the data. In the metric learning literature, a broader family of Mahalanobis metrics is parameterised by restricting  $M$  to the set of  $d \times d$  positive semidefinite matrices, which we will denote by  $\mathbb{S}_+^d$ . The task of learning a Mahalanobis metric is then generally expressed as an optimisation problem of the following form:

$$\min_{M \in \mathbb{S}_+^d} \ell(M, \mathcal{D}_n) + \lambda r(M).$$

Here,  $\ell$  is a loss function relating the suitability of  $M$  to the supervision brought by the data,  $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ , and  $r(M)$  describes some regularisation on the values of  $M$ , with the regularisation parameter  $\lambda \geq 0$ . Many formulations have been proposed, with differences arising in their choice of loss functions and regularisation. For a comprehensive review of established methods, refer to surveys provided by Kulis (2013), Bellet et al. (2013), and, more recently, Li and Tian (2018). Here, we limit our discussion to the main directions of the research and establish a relevant background for the current work.

A common approach converts supervision into the form of similarity judgments among tuples of training points. For instance, sets  $\mathcal{S}$  and  $\mathcal{D}$  are often used to contain pairs of training points deemed to be similar and dissimilar, respectively. Under this framework, the loss function,  $\ell(M, \mathcal{D}_n)$ , is used to encode violations of the desired relationships under the metric induced by  $M$ . Construction of a linear loss function leads to convex formulations and allows the problem to be treated as a semidefinite program. The earliest example of this is Xing et al. (2002), with the following intuitive formulation:

$$\min_{M \in \mathbb{S}_+^d} \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} d_M^2(\mathbf{x}_i, \mathbf{x}_j) \quad \text{s.t.} \quad \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{D}} d_M(\mathbf{x}_i, \mathbf{x}_j) \geq \gamma.$$

A succession of notable semidefinite programming formulations have since emerged, including the widely used large margin nearest neighbors (LMNN; Weinberger and Saul 2009). Motivated by  $k$ NN, LMNN seeks local neighbourhoods of each instance to be populated with other instances of the same class, whereas those of different classes are repelled by a “large” margin. LMNN has been the subject of a number of extensions. Torresani and Lee (2007) combine LMNN with dimensionality reduction of the feature space and explore nonlinear kernel methods, whereas Kedem et al. (2012) suggest alternative extensions to a nonlinear metric. Further notable works also making use of supervision sets  $\mathcal{S}$  and  $\mathcal{D}$  offering an information-theoretic perspective, include Bar-Hillel et al. (2003) and Davis et al. (2007).

The approach of populating supervision sets with pairwise relationships is less natural in our context of repeated feature vectors and stochastic classifications. In particular, we might expect the same pair to appear a number of times in *both*  $\mathcal{S}$  and  $\mathcal{D}$ , leading to either inefficient or infeasible formulations. Instead, the stochastic nature of the class labels leads us to prefer a probabilistic formulation, and we hence turn our attention to existing probabilistic methods.

An influential formulation known as neighbourhood components analysis (Goldberger et al. 2005) proceeds to model the nearest neighbour probabilities with a softmax normalisation over distances. A point,  $\mathbf{x}_i$ , identifies its nearest neighbour as  $\mathbf{x}_j$  with probability  $p_{ij}$ :

$$p_{ij} = \frac{\exp\{-\|A\mathbf{x}_i - A\mathbf{x}_j\|_2^2\}}{\sum_{k \neq i} \exp\{-\|A\mathbf{x}_i - A\mathbf{x}_k\|_2^2\}}, \quad p_{ii} = 0. \quad (1)$$

This imposed distribution provides a continuous relaxation to the deterministic point mass distribution on the nearest neighbour. With  $\mathbf{x}_i$  inheriting its classification from a neighbour selected from this distribution, the quantity  $p_i = \sum_{j: y_j=y_i} p_{ij}$  represents the probability of  $\mathbf{x}_i$  having the class assignment of  $y_i$ . The suggested objectives are to maximise the expected leave-one-out accuracy of the nearest neighbour classifier under this distribution, or the log-likelihood function:

$$\max_A \sum_{i=1}^n p_i, \quad \text{or} \quad \max_A \sum_{i=1}^n \log(p_i).$$

The optimisation of NCA is performed over an unconstrained matrix,  $A$ , which reveals a Mahalanobis matrix,  $M$ , via the relation  $A^T A = M$ . Any  $M \in \mathbb{S}_+^d$  permits this decomposition such that a Mahalanobis metric can equivalently be viewed as a projection of Euclidean space under a linear transformation,  $\mathbf{x} \rightarrow A\mathbf{x}$ . With this perspective, optimising the transformation matrix,  $A$ , provides an alternative perspective to the task of Mahalanobis metric learning. Dimensionality reduction can be directly sought by restricting the number of rows in  $A$ , whilst costly projections onto  $\mathbb{S}_+^d$  are avoided. However, an optimisation of  $A$  often comes at the expense of convexity: a loss function that is convex in  $M$  is typically nonconvex in  $A$ . This is no concern in the case of NCA, because the objective function is nonconvex in either parameterisation.

Underneath NCA lies a predictive model of class distributions:  $p_i$  can be seen as a kernel density estimator for the conditional probability of  $y_i$  given  $\mathbf{x}_i$  (Devroye and Wagner 1980). The form of the softmax transformation (1) disguises a Gaussian kernel function, with the common bandwidth parameter being absorbed into the scale of  $A$ . Weinberger and Tesauro (2007) more explicitly present a metric learning method in this context, albeit for a continuous target variable. In particular, kernel regression estimates for the target variable are constructed in the same form as above, and a metric is sought to minimise the mean squared error.

A similar probabilistic model was proposed by Peltonen and Kaski (2005), in which an expression of the data log-likelihood is maximised by evaluating conditional class probability estimators of a similar form as those of NCA. Indeed, the estimates introduced by NCA establish a popular framework that has been adopted by a number of other authors, and also provides inspiration for our own approach. Globerson and Roweis (2005) obtain a convex objective function, whereas Tarlow et al. (2013) explicitly extend the NCA objective to reflect the expected  $k$ NN accuracy for any choice of  $k$ . Further related extensions of NCA include that by Wang and Tan (2014), who consider the context of noisy labels arising from measurement error. Our own context involves true noise arising from a probabilistic labelling mechanism, and in this direction, Yang (2020) adapts NCA to the context of probabilistic labels, in which a full probabilistic class distribution is assumed to be known for each instance. Our own approach assumes no such knowledge.

The NCA framework provides a useful foundation for probabilistic metric learning, although it is not designed to apply when input observations are characterised by *repeating* features and *stochastic* outcomes. In this context, the objective function is heavily influenced by points' identical *self-neighbours*, which impede learning by remaining unchanged under any metric, and the pointwise formulation amounts to inefficient optimisation. We proceed in the next section to introduce a multinomial framework for our data setting and propose an NCA-inspired metric learning formulation to serve it.

### 3. Methodology

We begin by describing a data framework in which repeated feature vectors and stochastic labelling characterise a set of classification problems. To establish a mechanism generating observation pairs, we assume a distribution,  $q_{\mathbf{X}, Y}$ , over a pair of discrete random variables,  $\mathbf{X}$  and  $Y$ . The input variable,  $\mathbf{X}$ , takes values from a set,  $\mathcal{X} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_m\}$ , containing a finite number of vectors in  $\mathbb{R}^d$ , which we refer to as states, whereas  $Y$  represents

a class label drawn from a set,  $\mathcal{Y}$ , containing at least two unordered values. Our task is to find a pairwise distance function acting on the states in  $\mathcal{X}$  which in some way reflects the similarity of the conditional class distributions of  $Y|X$ . We receive supervision from a finite training sample,  $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ , drawn independently from  $q_{X,Y}$ .

In many applications, the training sample will not represent the entirety of  $\mathcal{X}$ . Although we expect repeats of *some* states, many others may be observed only once or not at all. Therefore, in classifying a new input,  $\mathbf{x}^*$ , we are not expecting to rely solely on observed training repetitions of  $\mathbf{x}^*$  and an empirical class distribution. Instead, we require a distance metric to point us to neighbouring states that, through a similarity of class distributions, will make good predictors. As the basis for this, we seek a model that can learn from the more frequently attended states to extend a conditional structure for  $Y|X$  across the whole set of  $\mathcal{X}$ .

The discrete nature of  $\mathcal{X} \times \mathcal{Y}$  allows us to view  $\mathcal{D}_n$  as a multinomial sample. We introduce a random variable,  $C_l^y(n) = |\{i \in \{1, 2, \dots, n\} : \mathbf{X}_i = \mathbf{b}_l, Y_i = y\}|$ , to represent the observed frequency of the pair  $(\mathbf{b}_l, y)$  in a data sample of size  $n$ . The random vector,  $\{C_l^y(n)\}_{l=1, \dots, m}^{y \in \mathcal{Y}}$ , follows a multinomial distribution with parameters  $n$  and  $\{q(\mathbf{b}_l, y)\}_{l=1, \dots, m}^{y \in \mathcal{Y}}$ , and we can represent  $\mathcal{D}_n = \{c_l^y\}_{l=1, \dots, m}^{y \in \mathcal{Y}}$  as a realisation of this random vector. Conditional on  $\mathcal{D}_n$ , the maximum likelihood estimates of the multinomial parameters are provided by  $\hat{q}(\mathbf{b}_l, y) = c_l^y/n$ . Using  $c_l$  to represent  $\sum_{y \in \mathcal{Y}} c_l^y$ , we also introduce marginal and conditional data distributions with  $\hat{q}(\mathbf{b}_l) = c_l/n$  and  $\hat{q}(y|\mathbf{b}_l) = c_l^y/c_l$  provided  $c_l > 0$ , respectively. Although these distributions depend on  $\mathcal{D}_n$ , we omit this to simplify the notation.

The remainder of this section is organised as follows. Section 3.1 introduces a probabilistic framework and objective function for SNCA. This establishes a model for the conditional distribution of  $Y|X$ . The relationship of SNCA to NCA is described in Section 3.2, before Section 3.3 presents a convergence result relating to the asymptotic behaviour of the optimal SNCA solution. Section 3.4 contains details of the optimisation algorithm applied in this paper.

### 3.1. A Probabilistic Framework

In the style of NCA, we impose a probability distribution for neighbour assignment based on a softmax function over distances. In contrast to NCA, however, this is applied over multinomial cells as opposed to over individual  $\mathbf{x}_i$ . Under a transformation matrix,  $A$ , a cell with  $\mathbf{X} = \mathbf{b}_l$  identifies the cell  $(\mathbf{b}_h, y)$  as its nearest neighbour with probability  $p^A((\mathbf{b}_h, y)|\mathbf{b}_l)$ :

$$p^A((\mathbf{b}_h, y)|\mathbf{b}_l) = \frac{c_h^y \exp\{-\|A\mathbf{b}_l - A\mathbf{b}_h\|_2^2\}}{\sum_{k \neq l} c_k \exp\{-\|A\mathbf{b}_l - A\mathbf{b}_k\|_2^2\}}, \quad p^A((\mathbf{b}_l, y)|\mathbf{b}_l) = 0.$$

Again, we omit the conditioning on  $\mathcal{D}_n$  from the notation. Assuming classifications to be inherited from a cell selected by this distribution, the probability of classifying  $\mathbf{X} = \mathbf{b}_l$  as class  $y$  is given by

$$p^A(y|\mathbf{b}_l) = \sum_{h=1}^m p^A((\mathbf{b}_h, y)|\mathbf{b}_l).$$

Under any matrix,  $A$ , this specifies a complete conditional distribution,  $p_{Y|X}^A$ , which is constructed as a weighted average over observed class distributions,  $\hat{q}_{Y|X}$ . That is, we can write

$$p^A(y|\mathbf{b}_l) = \sum_{h=1}^m p_{lh} \hat{q}(y|\mathbf{b}_h),$$

where the weight,  $p_{lh} = \sum_{y \in \mathcal{Y}} p^A((\mathbf{b}_h, y)|\mathbf{b}_l)$ , given to the observed class distribution of the neighbour  $\mathbf{b}_h$  decreases exponentially with the distance under  $A$  of  $\mathbf{b}_h$  from  $\mathbf{b}_l$ . With this perspective, we treat our task as an optimisation problem to find a transformation,  $A$ , to align the model distribution,  $p_{Y|X}^A$ , with the data distribution,  $\hat{q}_{Y|X}$ .

Obtaining the modelled joint distribution as  $p^A(\mathbf{b}_l, y) = \hat{q}(\mathbf{b}_l) p^A(y|\mathbf{b}_l)$ , the multinomial likelihood,  $L(A; \mathcal{D}_n)$ , and log-likelihood,  $\log L(A; \mathcal{D}_n)$ , of  $\mathcal{D}_n$  can be expressed as the following functions of  $A$ :

$$L(A; \mathcal{D}_n) = \frac{n!}{\prod_l \prod_y c_l^y!} \prod_{l=1}^m \prod_{y \in \mathcal{Y}} p^A(\mathbf{b}_l, y)^{c_l^y},$$

$$\log L(A; \mathcal{D}_n) = \log(n!) - \sum_{l=1}^m \sum_{y \in \mathcal{Y}} \log(c_l^y!) + \sum_{l=1}^m \sum_{y \in \mathcal{Y}} c_l^y \log(p^A(\mathbf{b}_l, y)).$$

Recalling that  $\hat{q}(\mathbf{b}_l, y) = c_l^y/n$ , we see that a maximisation of the log likelihood corresponds to the maximisation of a function,  $g_n(A)$ , and observe a connection to a Kullback–Leibler (KL) divergence:

$$\begin{aligned} \max_A \log L(A; \mathcal{D}_n) &= \max_A g_n(A), \\ \text{where } g_n(A) &= \sum_{l=1}^m \sum_{y \in \mathcal{Y}} \hat{q}(\mathbf{b}_l, y) \log(p^A(\mathbf{b}_l, y)) \\ &= \sum_{l=1}^m \sum_{y \in \mathcal{Y}} \hat{q}(\mathbf{b}_l, y) \log\left(\frac{p^A(\mathbf{b}_l, y)}{\hat{q}(\mathbf{b}_l, y)}\right) + \sum_{l=1}^m \sum_{y \in \mathcal{Y}} \hat{q}(\mathbf{b}_l, y) \log(\hat{q}(\mathbf{b}_l, y)) \\ &= -D_{\text{KL}}(\hat{q}_{\mathbf{X}, \mathbf{Y}} \| p_{\mathbf{X}, \mathbf{Y}}^A) - H(\hat{q}_{\mathbf{X}, \mathbf{Y}}). \end{aligned}$$

Here,  $H$  represents the Shannon (1948) entropy. Because  $H(\hat{q}_{\mathbf{X}, \mathbf{Y}})$  is independent of  $A$ , the maximisation of  $g_n(A)$  corresponds to a minimisation of the KL divergence of  $p_{\mathbf{X}, \mathbf{Y}}^A$  from  $\hat{q}_{\mathbf{X}, \mathbf{Y}}$ . We receive a similar interpretation in relation to the conditional distributions:

$$\begin{aligned} g_n(A) &= \sum_{l=1}^m \sum_{y \in \mathcal{Y}} \hat{q}(\mathbf{b}_l) \hat{q}(y | \mathbf{b}_l) \log(\hat{q}(\mathbf{b}_l) p^A(y | \mathbf{b}_l)) \\ &= \sum_{l=1}^m \hat{q}(\mathbf{b}_l) \sum_{y \in \mathcal{Y}} \hat{q}(y | \mathbf{b}_l) \log p^A(y | \mathbf{b}_l) + \sum_{l=1}^m \hat{q}(\mathbf{b}_l) \log \hat{q}(\mathbf{b}_l) \\ &= -\mathbb{E}_{\hat{q}_{\mathbf{X}}} [D_{\text{KL}}(\hat{q}_{\mathbf{Y} | \mathbf{X}} \| p_{\mathbf{Y} | \mathbf{X}}^A)] - \mathbb{E}_{\hat{q}_{\mathbf{X}}} [H(\hat{q}_{\mathbf{Y} | \mathbf{X}})] - H(\hat{q}_{\mathbf{X}}). \end{aligned}$$

Therefore, maximisation of  $g_n(A)$  minimises the expected KL divergence of the class conditional distributions of the model from the data, under  $\hat{q}_{\mathbf{X}}$ . This perspective perhaps most convincingly aligns with our intuition for a desirable metric, whilst making use of a common probability-based divergence measure.

### 3.2. Relationship to NCA

The formulations of NCA and SNCA exactly coincide when applied to training data without repeated observations, that is, training data for which  $\sum_{y \in \mathcal{Y}} c_l^y \leq 1$  holds for all  $l = 1, 2, \dots, m$ . When this condition does not hold, a difference arises. Expressing both formulations as

$$\max_A \sum_l \sum_y \hat{q}(\mathbf{b}_l, y) \log p^A(y | \mathbf{b}_l),$$

the difference occurs in the model estimates of the conditional probabilities,  $p^A(y | \mathbf{b}_l)$ . Whenever  $\hat{q}(\mathbf{b}_l, y) > 0$ , these estimates are as follows:

$$\text{NCA} : p_{(N)}^A(y | \mathbf{b}_l) = \frac{c_l^y - 1 + \sum_{h \neq l} c_h^y \exp\{-\|A\mathbf{b}_l - A\mathbf{b}_h\|_2^2\}}{c_l - 1 + \sum_{h \neq l} c_h \exp\{-\|A\mathbf{b}_l - A\mathbf{b}_h\|_2^2\}}, \quad (2)$$

$$\text{SNCA} : p_{(S)}^A(y | \mathbf{b}_l) = \frac{\sum_{h \neq l} c_h^y \exp\{-\|A\mathbf{b}_l - A\mathbf{b}_h\|_2^2\}}{\sum_{h \neq l} c_h \exp\{-\|A\mathbf{b}_l - A\mathbf{b}_h\|_2^2\}}. \quad (3)$$

The difference between (2) and (3) stems from the inclusion and exclusion, respectively, of a point's repeats in the softmax construction of its neighbour distribution. Because identical points will always have a distance of zero, the terms  $c_l^y - 1$  and  $c_l - 1$  appear regardless of  $A$  and effectively act as bias to  $p_{(N)}^A(y | \mathbf{b}_l)$ . As a result, the influence of  $A$  on the objective function of NCA is reduced. The adjustment made by SNCA is to exclude identical neighbours from consideration in the softmax functions, which allows the SNCA objective to solely reflect the quality of distance relationships among nonidentical points.

We also recognise inconsistencies arising in the NCA estimates. Specifically, we may not have  $\sum_{y \in \mathcal{Y}} p_{(N)}^A(y | \mathbf{b}_l) = 1$ . As  $n$  increases, the discrepancy reduces, and we can in fact identify an asymptotically optimal solution. In particular, a solution,  $A$ , with infinite column norms, such that nonidentical points are infinitely far apart, will yield  $p_{(N)}^A(y | \mathbf{b}_l) = (c_l^y - 1)/(c_l - 1) \rightarrow \hat{q}(y | \mathbf{b}_l)$  as  $n \rightarrow \infty$ . This identifies a noninformative solution to which NCA becomes susceptible, under its log-likelihood objective function, as the data size increases. This essentially represents the situation of fully overfitting to the data, a hazard that SNCA avoids by introducing implicit leave-one-out validation during learning.

### 3.3. Convergence of Optimal Solutions

In Section 3.1, we understood  $g_n(A)$  in relation to the observed distribution,  $\hat{q}_{\mathbf{X}, \mathbf{Y}}$ , of a finite sample. In this section, we seek to establish convergence guarantees with respect to the underlying true distribution,  $q_{\mathbf{X}, \mathbf{Y}}$ .

The strong law of large numbers (Kolmogorov and Širjaev 1992) ensures consistency of  $\hat{q}_{\mathcal{X},\mathcal{Y}} \xrightarrow{a.s.} q_{\mathcal{X},\mathcal{Y}}$  as  $n \rightarrow \infty$ , and by the continuous mapping theorem, we obtain the pointwise limits of  $g_n(A)$  over fixed  $A \in \mathbb{R}^{d \times d}$  as  $n \rightarrow \infty$ :

$$g_n(A) \xrightarrow{a.s.} \sum_{l=1}^m \sum_{y \in \mathcal{Y}} q(\mathbf{b}_l, y) \log \left( \frac{q(\mathbf{b}_l) \sum_{h \neq l} q(\mathbf{b}_h, y) \exp\{-\|A\mathbf{b}_l - A\mathbf{b}_h\|_2^2\}}{\sum_{k \neq l} q(\mathbf{b}_k) \exp\{-\|A\mathbf{b}_l - A\mathbf{b}_k\|_2^2\}} \right).$$

We use  $g(A)$  to denote the limit of  $g_n(A)$  given above, and consider the (set of) maximisers of  $g(A)$ , denoted by  $A_g^*$ , to be the true system-optimal solutions. To address the asymptotic behaviour of a finite sample solution,  $\hat{A}_g$ , we refer to a result relating to the consistency of sample average approximation estimators (Shapiro et al. 2021, theorem 5.3). Adopting the notation introduced here, the result is as follows.

**Theorem 1.** *Suppose there exists a compact subset  $C \subset \mathbb{R}^{d \times d}$  such that*

- i.  $A_g^*$  is nonempty and contained in  $C$ ,
- ii. for sufficiently large  $n$ , with probability one,  $\hat{A}_g$  is nonempty and contained in  $C$ ,
- iii.  $g(A)$  is finite valued and continuous on  $C$ , and
- iv.  $g_n(A) \xrightarrow{a.s.} g(A)$  as  $n \rightarrow \infty$ , uniformly on  $C$ .

*Then  $g_n(\hat{A}_g) \xrightarrow{a.s.} g(A_g^*)$  and  $\mathbb{D}(\hat{A}_g, A_g^*) \xrightarrow{a.s.} 0$  as  $n \rightarrow \infty$ .*

Here,  $\mathbb{D}(B, C)$  describes the deviation of the set  $B$  from the set  $C$ , defined as  $\mathbb{D}(B, C) = \sup_{x \in B} \{\inf_{x' \in C} \|x - x'\|\}$ . Central to this theorem is the assumption of a compact subset  $C$  in which  $A_g^*$  is contained. This excludes the possibility of infinite-valued solutions, which, although theoretically possible, we do not encounter in practice. For simplicity, we solve an optimisation problem in which the search space is not bounded, and assume that i and ii can be satisfied.

To address condition iii, we consider the argument of the logarithms to ensure that  $g(A)$  is finite valued. These arguments represent the limits of  $p^A(\mathbf{b}_l, y)$  as  $n \rightarrow \infty$  for all  $\mathbf{b}_l$  and  $y$ . Although these naturally reside in  $[0, 1]$ , we require them in the slightly stricter domain of  $[a, 1]$ , where  $a > 0$ . That is, we seek to exclude the possibility of  $p^A(\mathbf{b}_l, y)$  converging to zero when  $q(\mathbf{b}_l, y) > 0$ . This can be ensured with the following mild condition on  $q_{\mathcal{X},\mathcal{Y}}$ :

$$\sum_{h \neq l} q(\mathbf{b}_h, y) > 0 \quad \forall \mathbf{b}_l \in \mathcal{X}, y \in \mathcal{Y} \text{ with } q(\mathbf{b}_l, y) > 0. \quad (4)$$

In other words, each class,  $y \in \mathcal{Y}$ , must be observable in conjunction with at least two states in  $\mathcal{X}$ . Alongside the assumption that  $A \subset C$ , condition (4) ensures that  $g(A)$  is finite valued on  $C$ , and because it is also continuous, condition iii is satisfied. A justification of the uniform convergence required by condition iv can be found in Online Appendix A.1.

Whereas we understand  $\hat{A}_g$  as minimising an expected KL divergence of the class distributions of the model from the data, we consider  $A_g^*$  as relating to this expectation under the true data-generating distribution. Theorem 1 provides a helpful assurance that as the training sample increases and the empirical distribution,  $\hat{q}_{\mathcal{X},\mathcal{Y}}$ , approaches its truth, our metric solutions should also approach their truth.

### 3.4. Optimisation

To implement SNCA, the differentiability of  $g_n(A)$  allows us to apply a gradient-based optimiser. Recalling that  $p_{lh} = \sum_{y \in \mathcal{Y}} p^A((\mathbf{b}_h, y) | \mathbf{b}_l)$  and denoting  $\mathbf{b}_{lh} = \mathbf{b}_l - \mathbf{b}_h$ , we have the following gradient:

$$\frac{\partial g_n(A)}{\partial A} = 2A \sum_{l=1}^m \sum_{y \in \mathcal{Y}} \hat{q}(\mathbf{b}_l, y) \sum_{h \neq l} \mathbf{b}_{lh} \mathbf{b}_{lh}^\top p_{lh} \left( 1 - \frac{\hat{q}(y | \mathbf{b}_h)}{p^A(y | \mathbf{b}_l)} \right).$$

Details for deriving this expression can be found in Online Appendix A.2.

In the metric learning literature, imposing some regularisation on the parameters of the metric is a common and accepted practice (Kulis 2013). In this work, we choose to impose orthogonality among the rows of  $A$ . Encouraging orthogonality has recently been recognised as an effective regularisation for metric learning (Dutta et al. 2020), and besides leading to a smaller set of projection vectors which bear less redundancy, benefits in mitigating the effects of class imbalance and avoiding overfitting have also been recognised (Xie et al. 2018). We choose not to require unit length vectors, because a flexible scaling allows the predictive relevance of the various

orthogonal directions to be reflected. Our optimisation problem is thus expressed as follows:

$$\begin{aligned} & \max_{A \in \mathbb{R}^{r \times d}} g_n(A) \\ & \text{s.t. } AA^\top \text{ diagonal} \\ & \quad r \leq d \end{aligned}$$

Optimisation over the full space of feasible matrices described in this formulation can be computationally challenging, and scale poorly with the problem dimension,  $d$ . Therefore, we propose a tailored optimisation scheme in two stages. In the first stage, an orthogonal row-wise optimisation approach is used to identify a good solution with rank  $r \leq d$ , which is used in the second stage as the initial solution to a constrained optimisation problem over the full matrix. This establishes a computationally efficient procedure leading to a good locally optimal solution. This approach is summarised in Algorithm 1.

**Algorithm 1** (SNCA)

**Input:**  $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$

*Stage 1 – Row-wise construction of initial solution*

- 1:  $j \leftarrow 1$
- 2:  $A_0 \leftarrow []$  ▷ initialise empty solution matrix
- 3:  $P \leftarrow I_d$  ▷ initialise null space projection matrix
- 4:  $\text{old} \leftarrow -\infty$  ▷ initialise old objective value
- 5:  $\text{new} \leftarrow g_n(P)$  ▷ initialise new objective value
- 6:  $\text{obj\_values} \leftarrow ()$  ▷ initialise empty objective value vector
- 7: **while**  $\text{new} > \text{old}$  and  $j \leq d$  **do**
- 8:  $\mathbf{a}_j^* \leftarrow \arg \max_{\mathbf{a} \in \mathbb{R}^{1 \times d}} g_n \left( \begin{bmatrix} A_0 \\ \mathbf{a}P \end{bmatrix} \right)$  ▷ optimise row  $j$  in null space of previous rows
- 9:  $\text{old} \leftarrow \text{new}$  ▷ update old objective value
- 10:  $\text{new} \leftarrow g_n \left( \begin{bmatrix} A_0 \\ \mathbf{a}_j^* \end{bmatrix} \right)$  ▷ update new objective value
- 11:  $\text{obj\_values} \leftarrow (\text{obj\_values}, \text{new})$  ▷ track objective value improvements
- 12: **if**  $\text{new} > \text{old}$  **then** ▷ check if row  $j$  improves objective value
- 13:  $A_0 \leftarrow \begin{bmatrix} A_0 \\ \mathbf{a}_j^* \end{bmatrix} \in \mathbb{R}^{j \times d}$  ▷ add row  $j$  to solution matrix
- 14:  $\hat{\mathbf{a}}_j^* \leftarrow \mathbf{a}_j^* / \|\mathbf{a}_j^*\|_2$
- 15:  $P \leftarrow (I_d - \hat{\mathbf{a}}_j^* \hat{\mathbf{a}}_j^{*\top})P$  ▷ update projection matrix with null space of row  $j$
- 16:  $j \leftarrow j + 1$  ▷ move to next row
- 17: **end if**
- 18: **end while** ▷ stop when row  $j$  does not improve objective value

*Stage 2 – Orthogonal optimisation*

- 19:  $r \leftarrow j - 1$  ▷ default rank is the number of rows in solution  $A_0$ . Alternatively, select  $r < j - 1$  based on

the objective value improvements in  $\text{obj\_values}$ . Then  $A_0 \leftarrow \begin{bmatrix} \mathbf{a}_1^* \\ \vdots \\ \mathbf{a}_r^* \end{bmatrix} \in \mathbb{R}^{r \times d}$ .

- 20:  $A^* \leftarrow \arg \max_{A \in \mathbb{R}^{r \times d}} g_n(A)$  s.t.  $AA^\top$  diagonal ▷ optimisation with initial solution  $A_0$

**Return:**  $A^*$

In the first stage of the optimisation procedure (Algorithm 1, Lines 1–18), the row-wise construction of an initial solution is made by optimising each row in the null space of the previously determined rows. Performing this null space projection is a computationally inexpensive step and leads to orthogonal rows without the need for optimisation constraints. This row-by-row approach ensures that the most informative directions are learned first, and we thus observe diminishing row norms as incrementally less improvement is found with each added dimension. Correspondingly, the diminishing improvements to the objective value allow a natural stopping rule when further improvement cannot be made, giving an appropriate  $r \leq d$  and bringing explicit dimensionality reduction to the solution. In high-dimensional applications, the diminishing marginal gains to the objective value may not justify the computational cost of further row-wise optimisation, and a stricter stopping rule, requiring objective improvements to exceed a percentage threshold, may be applied at the user’s discretion.

The optimised orthogonal directions from the first stage generate a good solution, albeit one that may not be locally optimal. Therefore, this solution is used to initialise a constrained optimisation task in the second stage (Algorithm 1, Lines 19–20), leading to a final solution with a local optimality guarantee. Despite its good initial solution, however, the complexity of this second-stage optimisation can be significant, depending on parameters such as the problem dimension,  $d$ , and the matrix rank,  $r$ . For applications with high dimension, obtaining a lower-rank solution may be more advantageous than the diminishing improvements from the first stage, given the added complexity that an increased rank passes to the second stage, as well as to subsequent tasks such as nearest neighbour searches. Thus, choosing a lower rank,  $r$ , can be beneficial, and the objective values from the first-stage optimisation may be used to produce a diminishing returns curve to guide this decision. In the experiments presented in Section 4, the problem dimension is low enough that the time complexity of the second-stage optimisation is not prohibitive, and manual rank reduction is not required. In general, the orthogonal row-wise construction in the first stage naturally yields a low-rank solution, making further reduction of the rank unnecessary unless time constraints are a concern.

Although the second-stage optimisation offers a theoretical guarantee, we have observed in practice that it tends to provide only minor refinement to the first-stage solution. Although this may not always be the case, its theoretical benefits might therefore be weighed against time constraints in practical application.

Because we are optimising a nonconvex objective function, falling into “bad” local optima remains a concern in the first stage of row-wise optimisation. Therefore, in the experiments presented in Section 4, we make multiple restarts for each row’s optimisation. We choose a Latin hypercube space-filling design using software provided by Urquhart et al. (2020) to generate a set of initialisations and use the limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm for the optimisation with a convergence limit of  $\epsilon = 10^{-6}$  on the gradient. The orthogonality constraint in the second-stage optimisation is imposed using Riemannian optimisation. All implementation code is made available via the accompanying Code Ocean capsule (<https://doi.org/10.24433/CO.0189724.v5>).

## 4. Experiments on Simulated Data

In this section, we explore the performance of SNCA using data generated by stochastic simulation. We introduce a simple queueing model with a small and comprehensible state space, before presenting a more realistic application in the simulation of a semiconductor fabrication plant.

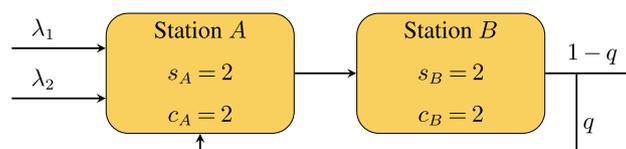
### 4.1. A Tandem Queueing System

Consider the simple queueing system represented in Figure 1. Two stations are arranged in tandem, with station  $I \in \{A, B\}$  containing  $s_I$  identical servers and  $c_I$  total capacity. Customers of type  $k \in \{1, 2\}$  arrive to station  $A$  according to a Poisson process with rate  $\lambda_k$  and experience service times at station  $I$  that are exponentially distributed with rate  $\mu_{Ik}$ . The parameter  $q$  is used to denote the probability of rework.

This two-station, two-type setup suggests a four-dimensional system state,  $\mathbf{X} = (X_{A1}, X_{A2}, X_{B1}, X_{B2})$ , where  $X_{Ik}$  represents the number of type  $k$  customers at station  $I$ . The memoryless property of the exponential distribution makes elapsed service times irrelevant, and setting  $s_I = c_I$  for  $I \in \{A, B\}$  removes queueing space from both stations and ensures that  $\mathbf{X}$  fully characterises the observable system state. We set  $s_I = c_I = 2$  for  $I \in \{A, B\}$ , which allows  $\mathbf{X}$  to take 36 distinct states.

We refer to station  $I \in \{A, B\}$  as being blocked whenever  $X_{I1} + X_{I2} = c_I$ , because no customers can enter the station in this state. In the logic of the system, customers completing their service in station  $A$  will occupy space there while station  $B$  is blocked. More significantly, customers arriving to a blocked station  $A$  (including those returning from station  $B$ ) are immediately lost from the system. As such, blocking in station  $A$  represents a critical state for the system, and our interest in this example is in predicting the future probability of this event based on the currently observed state. In particular, we observe the system state,  $\mathbf{X}(t)$ , at time  $t$ , and predict whether

**Figure 1.** (Color online) A Tandem Queueing System with Feedback



station  $A$  will be blocked at time  $t + T$ . Let

$$Y(t) = \begin{cases} 1 & \text{if } X_{A1}(t+T) + X_{A2}(t+T) = c_A, \\ 0 & \text{otherwise.} \end{cases}$$

Simulating the system provides the data  $\mathcal{D}_n = \{(\mathbf{X}(t_i), Y(t_i))\}_{i=1}^n$ , and we choose the recording times such that  $t_{i+1} = t_i + T$  in order to obtain nonoverlapping observations. To provide a roughly even distribution for  $Y$ , we simulate the system with the parameters  $\lambda_1 = \lambda_2 = 2, \mu_{A1} = 2, \mu_{A2} = 3, \mu_{B1} = 1, \mu_{B2} = 2$ , and  $q = 0.2$ , and we set the prediction horizon to  $T = 0.25$ . Choosing a smaller value of  $T$  would move toward a more deterministic problem, whereas a longer horizon would reduce the influence of  $\mathbf{X}(t)$  and converge each state’s conditional probability of  $Y = 1$  toward the system’s long-run blocking probability for station  $A$ . With the specified system parameters, setting  $T = 0.25$  provides a suitable middle ground in which the set of system states exhibits a wide range of class probabilities.

**4.1.1. Metric Embeddings.** On a data set of size  $n = 100,000$ , we obtained the following SNCA solution:

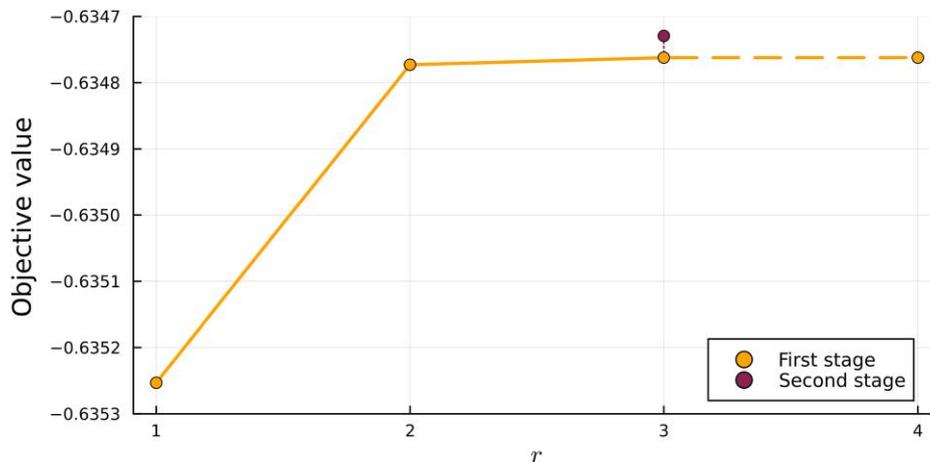
$$A = \begin{bmatrix} 7.53 & 6.92 & 3.17 & 1.85 \\ 0.31 & -0.71 & 0.25 & 1.01 \\ -0.15 & 0.03 & 0.31 & -0.01 \end{bmatrix}.$$

The transformation matrix consists of orthogonal rows representing data directions that are informative with respect to the class distributions. Following the optimisation scheme detailed in Section 3.4, rows are appended to the (first-stage) solution sequentially, and as such, their relevance decreases and their magnitudes shrink, pointing us to an appropriate lower-dimensional solution. The trajectory of the objective function value throughout the optimisation process is shown in Figure 2. In this example, a dominant first dimension provides most of the projection variation, with a second and a third dimension making smaller contributions. A fourth dimension was not found to provide any further improvements to the objective value.

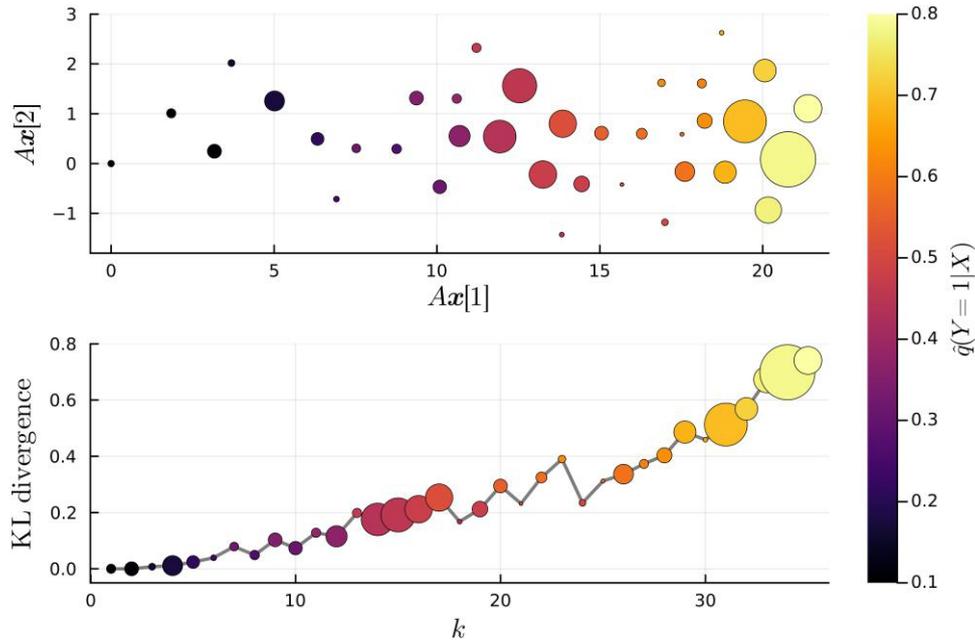
We consider the first row of the solution as the most informative, and we can derive interpretation from the magnitudes in this row acting on the original data dimensions. The dimension  $X_{A1}$  is deemed slightly more relevant than  $X_{A2}$ , which we can understand, because the system was simulated with  $\mu_{A1} < \mu_{A2}$ . This means that customers represented by  $X_{A1}$  tend to have longer service times and so are more likely to still occupy station  $A$  when  $Y$  is observed. Customers in station  $B$  can indirectly increase the risk of station  $A$  becoming blocked, and in the same way,  $\mu_{B1} < \mu_{B2}$  leads to a first dimension to which  $X_{B1}$  makes a greater contribution than  $X_{B2}$ .

To visualise the effect of SNCA, we plot the 36 unique states projected into the solution space, showing the first two dimensions that capture the majority of the variation (Figure 3, top). Recall that our distance metric is represented by Euclidean distance in this transformed space. The projection plot shows the SNCA metric arranging the states with regard to the similarity of their stochastic class distributions, which are depicted by the shading scale. The state  $\mathbf{X} = (0, 0, 0, 0)$  represents the empty system and is leftmost in the projection plot. The lower plot explicitly lays out the states in the order of their distance from  $(0, 0, 0, 0)$  and shows that the KL divergence of a neighbouring state’s class distribution from that of  $(0, 0, 0, 0)$  tends to increase as the neighbour moves further away.

**Figure 2.** (Color online) The Objective Value Improvement Throughout the SNCA Algorithm



**Figure 3.** (Color online) (Top) The 36 Unique States Projected in the First Two Dimensions of the SNCA Solution Matrix and (Bottom) the KL Divergence from the Class Distribution of the State  $(0,0,0,0)$  to That of Its  $k$ th Nearest Neighbour



Note. The shading of the points relates to their empirical class distribution, and the size is proportional to their observed frequency.

To allow an appropriate comparison for plotting, we obtain the following two-dimensional NCA solution on the same data set:

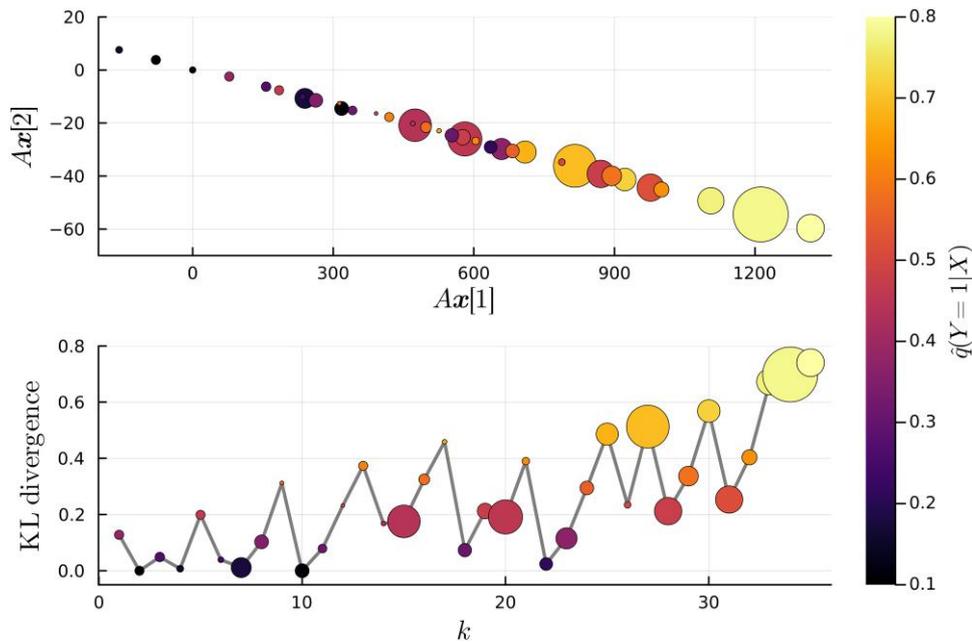
$$A = \begin{bmatrix} 341.23 & 234.95 & 317.98 & -78.4 \\ -15.27 & -10.11 & -14.54 & 3.8 \end{bmatrix}.$$

The projection and KL divergence plots corresponding to Figure 3 for this NCA solution are shown in Figure 4. The comparison between Figures 3 and 4 highlights the deficiency of NCA when applied to discrete stochastic data, and the suitable corrections made by SNCA. Notably, the main direction of the NCA solution is less effective at discriminating the class distributions than the main direction of SNCA (see the top panels of Figures 3 and 4,  $x$  direction). This highlights the effectiveness of the SNCA adjustment to exclude a point's repeats from its probabilistic neighbour distribution. Additionally, the usefulness of the orthogonality constraints introduced in the SNCA implementation is apparent, as the two rows of the NCA solution are seen to represent very similar directions. In the lower plots of Figures 3 and 4, the 35 remaining states are displayed along the  $x$ -axis in the order of their distance from the state  $(0,0,0,0)$ , with the KL divergence of their class distribution from that of  $(0,0,0,0)$  shown on the  $y$ -axis. Thus, we are looking for a positive correlation between the two axes; SNCA noticeably achieves this pattern more successfully than NCA. These figures demonstrate that the adjustment made by SNCA to exclude self-neighbours is effective; it forces the algorithm to find good-quality, nonidentical neighbours. NCA does not have this incentive, because identical  $X$  observations dominate its neighbour distributions and the influence of nonidentical neighbours on its objective function is diminished.

The reference state of  $(0,0,0,0)$  used for the lower plots of Figures 3 and 4 was not specially selected. Figure 5 shows this relationship between the distance ordering and the KL divergence averaged with respect to  $\hat{q}_X$  over all 36 states. Across the whole input space, SNCA has consistently arranged states' neighbours with broadly increasing divergence of their class distributions, whereas NCA is less reliable in this regard. For the purpose of nearest neighbour predictions, the closer neighbourhoods are the most relevant, and we see SNCA bringing a desirable structure in terms of the similarity of class distributions within the closer neighbourhoods.

**4.1.2.  $k$ NN Classification.** As a further verification of metric performance, we display results of  $k$ NN classification under the learned distance metrics. To have a clearer interpretation of  $k$  in our data context, we consider  $k$  as the number of distinct neighbours, and include all of their repeats. In other words, we classify an individual point,  $x_i$ , by the majority class of its  $k$  nearest neighbour states, and additionally include any states that are tied in

**Figure 4.** (Color online) The Equivalent Plots from Figure 3 for a Two-Dimensional NCA Solution

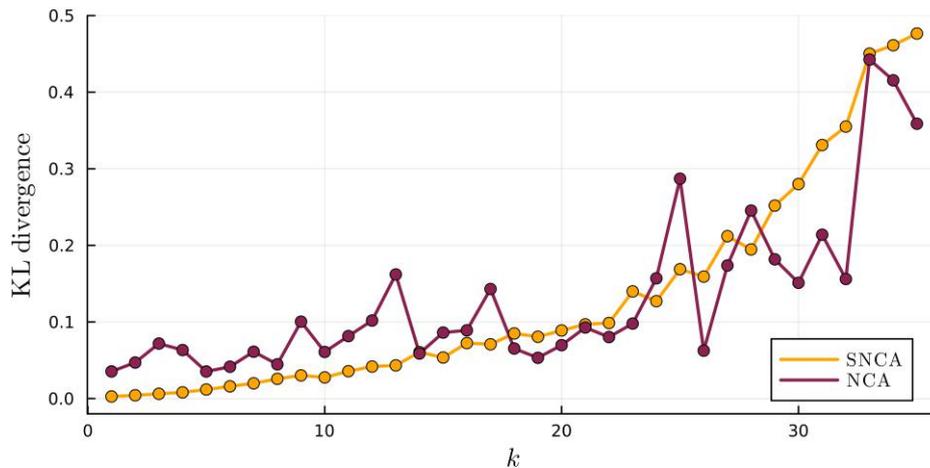


*Notes.* The top panel shows the 36 unique states projected by the solution matrix. The bottom panel shows the KL divergence from the class distribution of the state  $(0, 0, 0, 0)$  to that of its  $k$ th nearest neighbour.

distance with the  $k$ th. Further, we exclude the other observations of  $x_i$ 's own state from the neighbour search. This would be the nearest neighbour state under any metric, so its inclusion does not help a comparison across metrics. Further details of the  $k$ NN procedures used in this section, and the generation of the figures, are found in Online Appendix A.3.

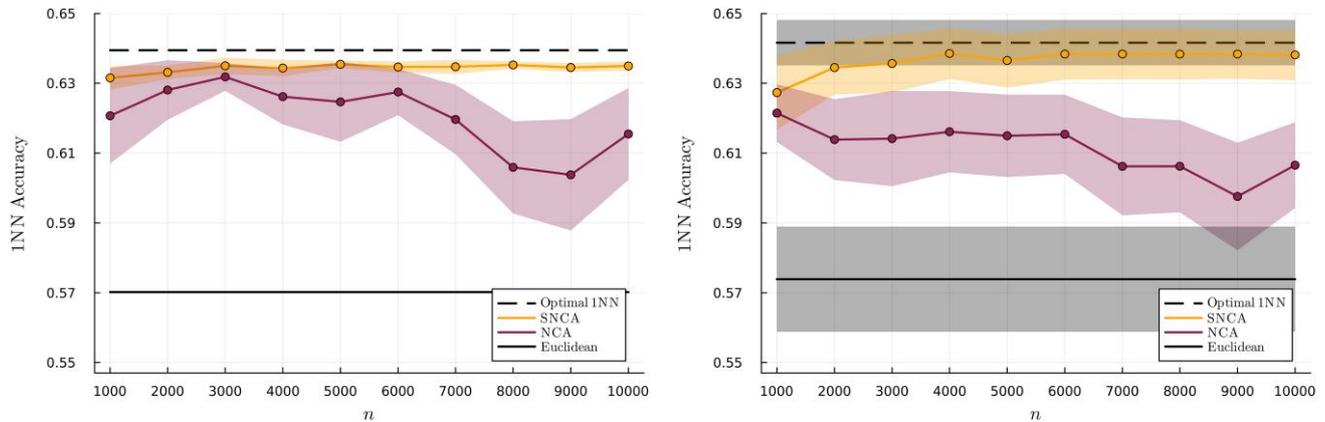
Because this example carries only a small number of distinct  $X$  values, we only show results for  $k = 1$ . Figure 6 concerns metrics trained on data sets of size  $n$  and shows on the left the leave-one-out correct classification rates on a large ( $n = 100,000$ ) test data set. This is to approximate the expected classification accuracy on data drawn randomly from  $q_{X,Y}$ . As benchmarks for comparison, we include the 1NN classification rate achieved using a Euclidean distance metric, and a 1NN classifier that omnisciently selects the optimal neighbour, defined as the state whose class distribution has minimum KL divergence from that of the query point. We are encouraged to

**Figure 5.** (Color online) The Average KL Divergence from a Point's Class Distribution to the Class Distribution of its  $k$ th Nearest Neighbour, Under the SNCA and NCA Embeddings



*Note.* This is equivalent to the lower plots of Figures 3 and 4 averaged with respect to  $\hat{q}_X$  over all states.

**Figure 6.** (Color online) (Left) The Leave-One-Out 1NN Correct Classification Rate Under Metrics Learnt from a Training Set of Size  $n$  and (Right) 1NN Classification Performance on Test States Withheld from the Metric Training Following Sixfold Cross-Validation



Note. The shaded regions indicate  $\pm 1.96$  standard errors around the average classification accuracy.

see SNCA still performing well with small training data sets, and achieving classification rates close to optimal given the level of stochasticity in the system.

In practical problems with large input spaces, we do not expect to exhaust  $\mathcal{X}$  in our training sample. For this reason, we are particularly interested in classification performance on states that are unrepresented in the metric training. Figure 6 (right) uses sixfold cross-validation on the 36 unique states, through which the metrics were trained on data covering 30 states, and the classification accuracy recorded on the remaining 6. We are encouraged to see strong 1NN classification performance by SNCA on unseen states. For each  $n$  marked in the plots in Figure 6, the metric training was repeated with 10 independent data sets of size  $n$ , with the ribbons showing pointwise  $\pm 1.96$  standard errors around the average classification rate. We see much greater variability in the NCA solutions as compared with SNCA.

The  $k$ NN performance of SNCA does not noticeably benefit by increasing the training set beyond a size of  $n = 3,000$ . We can be confident that this is enough training data and the metric solution is not overfitted. In practice, we advise increasing the training set size for SNCA until the classification accuracy on test data and the interpretations of the learned metric appear consistent.

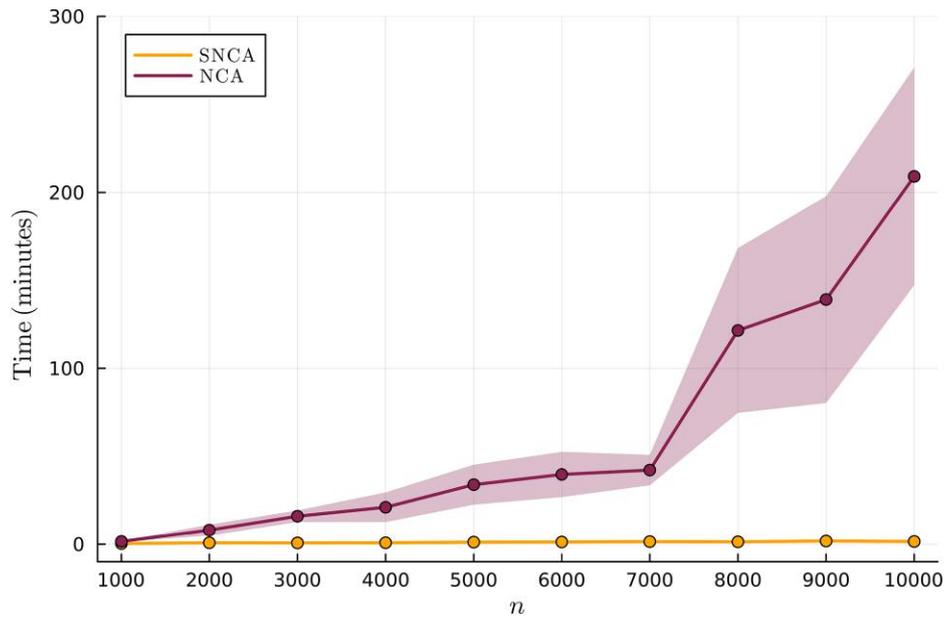
**4.1.3. Execution Time.** To demonstrate the computational benefit of SNCA with discrete stochastic data, Figure 7 shows the mean execution time using a 1.6 GHz Intel Core i5 processor for the metric training of SNCA and NCA across the 10 training sets of size  $n$  used in Figure 6 (left),  $\pm 1.96$  standard errors. The reported times include 10 initialisations for each optimisation. With the small state space in this example, all states are sampled in the smallest data set size of  $n = 1,000$ . After this, increasing the value of  $n$  only refines the empirical class distributions and does not increase the burden on SNCA. In contrast, a pointwise metric learning method such as NCA, which involves summation over  $n$  individual points, understandably suffers as  $n$  increases. It should be noted that when training NCA on the data set of size  $n = 100,000$  to illustrate the metric embedding in Section 4.1.1, the pointwise formulation was not feasible. Instead, we employed a modified algorithm in which identical points were grouped together in the calculations of the objective function and the gradient.

## 4.2. A Wafer Fab Model

As a more realistic example for discrete stochastic data, we explore the simulation model of a wafer fabrication facility (fab) described by Kayton et al. (1996). Wafers are released into this system at fixed intervals and follow a sequence involving several processing steps at a number of stations. Aspects of reentrant flow and machines offering different processing capacities and unpredictable breakdown patterns complicate the management of product flow through wafer fabs, and a simulation model becomes an indispensable tool. For this example, we used Simul8 software to imitate the model described by Kayton et al. (1996).

To create a classification problem, we consider the practical task of predicting an early or late completion based on the system state on a wafer's release. To restrict the feasible state space and encourage repeated observations, we take an incomplete view of the system state, with  $\mathbf{X}$  containing simply the queue sizes at each of the facility's

**Figure 7.** (Color online) Metric Learning Execution Time on a Training Set of Size  $n$

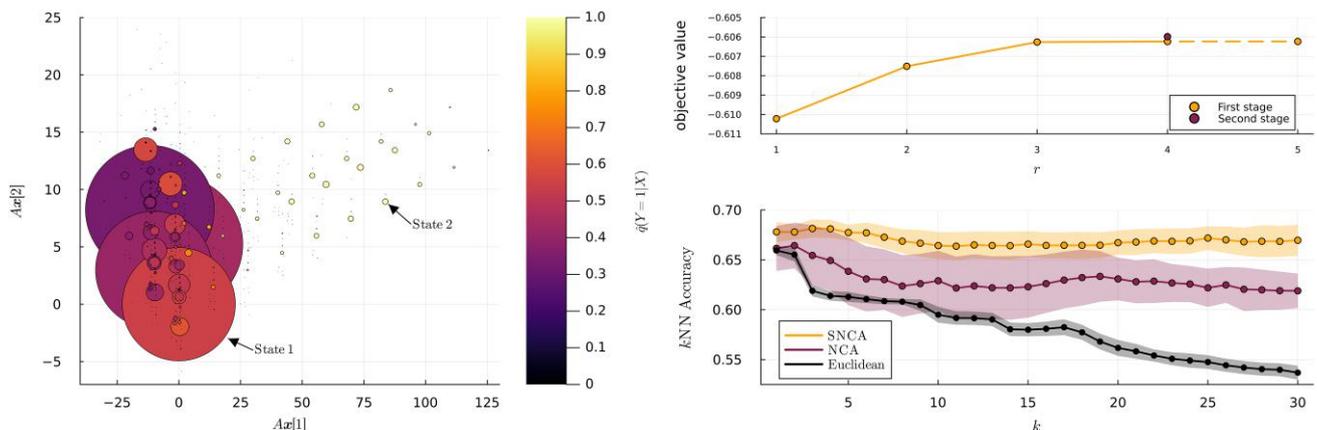


Note. The shaded regions indicate  $\pm 1.96$  standard errors following 10 independent trials for each  $n$ .

11 stations. The size of  $\mathcal{X}$  is further limited by a total capacity constraint on the system, whereby new wafers cannot be released if the total work in process exceeds a threshold. The classification,  $Y$ , is observed as the wafer leaves the facility either before or after its due date, where the due date is preassigned by a constant inflation factor on its pure processing time. In this problem, we use a classification of  $Y = 0$  to indicate an early or on-time wafer completion, whereas  $Y = 1$  implies a late completion. Running the simulation provides an observation,  $(X, Y)$ , for every wafer that passes through the system. A classification model from these simulated data can be applied on the actual wafer fab to alert planners to possible late deliveries.

The system design yields around 21,000 feasible combinations of the 11 queue sizes, although the distribution over this state space is highly imbalanced. Figure 8 (left) shows a data set of size  $n = 18,402$  projected in the first two dimensions of an SNCA solution, where size is proportional to the observed frequency of the states. There is a small number of commonly visited states, although many more are observed very few times, and overall, only

**Figure 8.** (Color online) (Left) States from the Wafer Fab Model Projected in the First Two Rows of an SNCA Solution; (Top Right) Objective Value Improvement Throughout the SNCA Algorithm; and (Bottom Right)  $k$ NN Correct Classification Rate on Rarely Seen States



Notes. In the left panel, the shading of the points relates to their observed class proportions, and the size is proportional to their observed frequency. In the bottom right panel, the shaded regions indicate  $\pm 1.96$  standard errors.

Downloaded from informs.org by [98.253.32.243] on 05 February 2026, at 04:43 . For personal use only, all rights reserved.

a small percentage of the feasible space has been represented in the sample. To aid understanding of this plot, consider the point lying on the origin, labelled as state 1. In the original context, this represents the system state in which all queues are empty. This is a commonly observed state for the system. It is also a difficult state to classify; its shade indicates the  $Y = 1$  class has an observed proportion of around 0.55. In other words, a wafer released while the system has empty queues is marginally more likely to be completed late than on time. Although counterintuitive, this is because the first station is a batch station, which requires at least two wafers to be processed simultaneously. Therefore, a wafer released when the first queue is empty must invariably wait at least until the following wafer is released. Other states shown by the shading scale to have lower proportions of late completions represent states in which the first queue is nonempty.

The labelled state 2 in Figure 8 (left) represents the state in which six wafers occupy the queue at station 3 and all other queues are empty. In the system logic, the machine at station 3 is prone to breakdowns, which leads to the occasional buildup of this station queue. This represents a rare yet critical system condition and accounts for the branch of lesser-observed states extending away from the main cluster. The class distributions of these states are heavily skewed in favour of  $Y = 1$ , as wafers entering a blocked system are highly likely to finish late. SNCA rightly identifies this third queue size as an effective indicator of the system performance and lets this variable dominate the first projection dimension.

Of primary interest is to understand metric performance with regard to states that are underrepresented in the training sample, because these are the states for which we are more reliant on neighbours to make a prediction. Therefore, as a more targeted approach, we show the  $k$ NN accuracy on the rarely seen states. In particular, we consider states observed in the data set with fewer than 25 repetitions. Figure 8 (bottom right) shows the average  $k$ NN accuracy on these states  $\pm 1.96$  standard errors, after repeating the experiment with 10 simulated data sets. Additional details on the generation of this plot are provided in Online Appendix A.3.

To provide context to the  $k$ NN accuracy results, the system was simulated to give roughly equal proportions of early and late wafers. Employing  $k$ NN on the simplified state description used in this example provides additional predictive power, and the benefit of a trained distance metric is evident from the comparison with Euclidean distance. Reaching around 68% accuracy following metric learning represents a marked and useful achievement considering the stochasticity of the system. This example is broadly representative of the type of data we aim to cater to with SNCA, and the further performance improvement brought over NCA is encouraging.

### 4.3. Discussion

The proposed SNCA constitutes a metric learning method designed under the assumption that data contain repeating input features with stochastic classifications, a scenario that is characteristic of discrete-event simulations, our intended application. The probabilistic formulation uses a multinomial framework to accommodate observations from a discrete and finite data space. We recognise that SNCA represents a *linear* and *global* metric learning method, in that it learns a single linear projection to apply to all data points. It is not designed to accommodate nonlinear or value-dependent relationships among the predictors and the response; the suitability of this limitation should be considered with each application.

Nonetheless, the assumptions of SNCA are broadly characteristic of classification problems within stochastic simulation, and this section presented examples from this domain to explore its performance on such data. Initially, a simple simulation model was chosen to provide an example with a small state space and intuitive system behaviour against which to interpret the SNCA metric. Figures showed the desirable property of the SNCA solution in arranging states in the projected space according to the similarity of their various class distributions.

We note that applications in the simulation domain often have very low sampling cost, and additional trials on this initial example indicate the computational advantage of the SNCA formulation with increasing data set size. However, the advantages seen in typical applications with larger state spaces will be less pronounced, because the small state space in this example was comprehensively sampled such that incremental sampling provided no additional contributions to the objective function summation. Further, the scaling of SNCA with dimensionality has not been explicitly explored, although we recognise that the suggested optimisation approach (Section 3.4) scales well owing to the row-wise construction of a good initial solution in the first stage, and the scalability of the second-stage matrix optimisation relies upon the gradient descent procedure. Although we thus find practical advantages through the suggested optimisation scheme, our main interest in this paper is toward the behaviour of the SNCA projection when applied to simulation-generated data.

A second simulation example was presented to provide a more realistic data set with greater dimension and state space. This example also showed good projection behaviour from SNCA and improved performance compared with NCA when using the distance metric for  $k$ NN classification. Compared with NCA across all

examples we have looked at offering discrete stochastic data, SNCA was found to give equivalent or improved  $k$ NN performance.

## 5. Conclusion

In this work, we proposed a method of distance metric learning that we refer to as stochastic neighbourhood components analysis. The aim of SNCA is to extend the reach of metric learning to scenarios involving repeated feature vectors with stochastic labelling. This is a type of data characteristic of classification problems arising from the sample paths of stochastic simulation. Interest in probing the dynamic behaviour of simulation is growing, and to this purpose, we demonstrated SNCA for nearest neighbour classification on sample path data. This is a useful application for SNCA; the ability to anticipate stochastic behaviour provides valuable assistance to real-time planning and control in a live system and is well supported by the emergence of digital twin simulation. In addition, an SNCA solution matrix can reveal the influential drivers of a system's stochastic behaviour. Experimental results show SNCA to be effective and to bring further improvement to  $k$ NN performance over NCA when applied to this type of data.

We recognise numerous contexts to which the model of discrete stochastic data extends itself. For example, crowdsourcing represents an increasingly attractive solution to data generation and will often result in an instance being assigned to various classes. Although we have focused our motivation on simulation analytics, the scope of SNCA stretches further, and we see possibility for a wider impact.

## Acknowledgments

The authors thank the editors and anonymous reviewers for their insightful comments and constructive feedback, which have improved the quality and clarity of this manuscript.

## References

- Bar-Hillel A, Hertz T, Shental N, Weinshall D (2003) Learning distance functions using equivalence relations. Fawcett T, Mishra N, eds. *Proc. 20th Internat. Conf. Machine Learn.* (AAAI Press, Washington, DC), 11–18.
- Bellet A, Habrard A, Sebban M (2013) A survey on metric learning for feature vectors and structured data. Preprint, submitted June 28, <https://arxiv.org/abs/1306.6709>.
- Cao L, Tay FE (2001) Financial forecasting using support vector machines. *Neural Comput. Appl.* 10(2):184–192.
- Davis JV, Kulis B, Jain P, Sra S, Dhillon IS (2007) Information-theoretic metric learning. Ghahramani Z, ed. *Proc. 24th Internat. Conf. Machine Learn.* (Association for Computing Machinery, New York), 209–216.
- Devroye LP, Wagner T (1980) Distribution-free consistency results in nonparametric discrimination and regression function estimation. *Ann. Statist.* 8(2):231–239.
- dos Santos CH, Montevechi JAB, de Queiroz JA, de Carvalho Miranda R, Leal F (2022) Decision support in productive processes through des and abs in the digital twin era: A systematic literature review. *Internat. J. Production Res.* 60(8):2662–2681.
- Dutta UK, Harandi M, Sekhar CC (2020) Unsupervised deep metric learning via orthogonality based probabilistic loss. *IEEE Trans. Artificial Intelligence* 1(1):74–84.
- Globerson A, Roweis S (2005) Metric learning by collapsing classes. Weiss Y, Schölkopf B, Platt JC, eds. *Advances in Neural Information Processing Systems*, vol. 18 (MIT Press, Cambridge, MA), 451–458.
- Goldberger J, Roweis ST, Hinton GE, Salakhutdinov RR (2005) Neighbourhood components analysis. Saul L, Weiss Y, Bottou L, eds. *Advances in Neural Information Processing Systems*, vol. 17 (MIT Press, Cambridge, MA), 513–520.
- Hong LJ, Jiang G (2019) Offline simulation online application: A new framework of simulation-based decision making. *Asia-Pacific J. Oper. Res.* 36(6):1–22.
- Horvat T, Job J (2020) The use of machine learning in sport outcome prediction: A review. *WIREs Data Mining Knowledge Discovery* 10(5):e1380.
- Jiang G, Hong LJ, Nelson BL (2020) Online risk monitoring using offline simulation. *INFORMS J. Comput.* 32(2):356–375.
- Kayton D, Teyner T, Schwartz C, Uzsoy R (1996) Effects of dispatching and down time on the performance of wafer fabs operating under the-ory of constraints. *Proc. 19th IEEE/CPMT Internat. Electronics Manufacturing Tech. Sympos.* (Institute of Electrical and Electronics Engineers, Piscataway, NJ), 49–56.
- Kedem D, Tyree S, Sha F, Lanckriet GR, Weinberger KQ (2012) Non-linear metric learning. Pereira F, Burges CJ, Bottou L, Weinberger KQ, eds. *Advances in Neural Information Processing Systems*, vol. 25 (Curran Associates, Red Hook, NY), 2573–2581.
- Keslin G, Nelson BL, Pagnoncelli BK, Plumlee M, Rahimian H (2024) Ranking and contextual selection. *Oper. Res.*, ePub ahead of print October 3, <https://doi.org/10.1287/opre.2023.0378>.
- Kolmogorov AN, Shiryaev AN (1992) *Selected Works of A. N. Kolmogorov: Volume 2, Probability Theory and Mathematical Statistics* (Kluwer Academic Publishers, Berlin, Heidelberg).
- Kulis B (2013) Metric learning: A survey. *Foundations Trends Machine Learn.* 5(4):287–364.
- Laidler G, Morgan LE, Nelson BL, Pavlidis NG (2020) Metric learning for simulation analytics. Bae KH, Feng B, Kim S, Lazarova-Molnar S, Zheng Z, Roeder T, Thiesing R, eds. *Proc. 2020 Winter Simulation Conf.* (Institute of Electrical and Electronics Engineers, Piscataway, NJ), 349–360.
- Law AM, Kelton WD (2007) *Simulation Modeling and Analysis*. 3rd ed. (McGraw-Hill, New York).
- Li D, Tian Y (2018) Survey and experimental study on metric learning methods. *Neural Networks* 105(C):447–462.

- Lin Y, Nelson BL, Pei L (2019) Virtual statistics in simulation via  $k$  nearest neighbors. *INFORMS J. Comput.* 31(3):576–592.
- Mahalanobis P (1936) On the generalised distance in statistics. *Proc. Natl. Inst. Sci. India* 2:49–55.
- McFee B, Lanckriet GR (2010) Metric learning to rank. Fürnkranz J, Joachims T, eds. *Proc. 27th Internat. Conf. Machine Learn.* (Omnipress, Madison, WI), 775–782.
- Morgan LE, Barton RR (2022) Fourier trajectory analysis for system discrimination. *Eur. J. Oper. Res.* 296(1):203–217.
- Nelson BL (2016) Some tactical problems in digital simulation for the next 10 years. *J. Simulation* 10(1):2–11.
- Nelson BL, Pei L (2021) *Foundations and Methods of Stochastic Simulation: A First Course*, 2nd ed. (Springer Nature, Cham, Switzerland).
- Ouyang H, Nelson BL (2017) Simulation-based predictive analytics for dynamic queueing systems. Chan WKV, D'Ambrogio A, Zacharewicz G, Mustafee N, Wainer G, Page E, eds. *Proc. 2017 Winter Simulation Conf.* (Institute of Electrical and Electronics Engineers, Piscataway, NJ), 1716–1727.
- Peltonen J, Kaski S (2005) Discriminative components of data. *IEEE Trans. Neural Networks* 16(1):68–83.
- Shannon CE (1948) A mathematical theory of communication. *Bell System Tech. J.* 27(3):379–423.
- Shapiro A, Dentcheva D, Ruszczyński A (2021) *Lectures on Stochastic Programming: Modeling and Theory* (Society for Industrial and Applied Mathematics, Philadelphia).
- Shen H, Hong LJ, Zhang X (2021) Ranking and selection with covariates for personalized decision making. *INFORMS J. Comput.* 33(4):1500–1519.
- Suárez JL, García S, Herrera F (2021) A tutorial on distance metric learning: Mathematical foundations, algorithms, experimental analysis, prospects and challenges. *Neurocomputing* 425:300–322.
- Suo Q, Ma F, Yuan Y, Huai M, Zhong W, Gao J, Zhang A (2018) Deep patient similarity learning for personalized healthcare. *IEEE Trans. Nanobioscience* 17(3):219–227.
- Tarlow D, Swersky K, Charlin L, Sutskever I, Zemel R (2013) Stochastic  $k$ -neighborhood selection for supervised and unsupervised learning. Dasgupta S, McAllester D, eds. *Proc. 30th Internat. Conf. Machine Learn.*, Proceedings of Machine Learning Research, vol. 28 (JMLR.org), 199–207.
- Torresani L, Lee K (2007) Large margin component analysis. Schölkopf B, Platt JC, Hoffman T, eds. *Advances in Neural Information Processing Systems*, vol. 19 (MIT Press, Cambridge, MA), 1385–1392.
- Urquhart M, Ljungskog E, Sebben S (2020) Surrogate-based optimisation using adaptively scaled radial basis functions. *Appl. Soft Comput.* 88(C):106050.
- Vaughan JW (2018) Making better use of the crowd: How crowdsourcing can advance machine learning research. *J. Machine Learn. Res.* 18(193):1–46.
- Wang D, Tan X (2014) Robust distance metric learning in the presence of label noise. *Proc. 28th AAAI Conf. Artificial Intelligence* (AAAI Press, Palo Alto, CA), 1321–1327.
- Weinberger KQ, Saul LK (2009) Distance metric learning for large margin nearest neighbor classification. *J. Machine Learn. Res.* 10(9):207–244.
- Weinberger KQ, Tesauro G (2007) Metric learning for kernel regression. Meila M, Shen X, eds. *Proc. 11th Internat. Conf. Artificial Intelligence Statist.*, Proceedings of Machine Learning Research, vol. 2 (JMLR.org), 612–619.
- Xie P, Wu W, Zhu Y, Xing E (2018) Orthogonality-promoting distance metric learning: Convex relaxation and theoretical analysis. Dy J, Krause A, eds. *Proc. 35th Internat. Conf. Machine Learn.*, Proceedings of Machine Learning Research, vol. 80 (JMLR.org), 5403–5412.
- Xing EP, Ng AY, Jordan MI, Russell S (2002) Distance metric learning, with application to clustering with side-information. Becker S, Thrun S, Obermayer K, eds. *Advances in Neural Information Processing Systems*, vol. 15 (MIT Press, Cambridge, MA), 521–528.
- Yang X (2020) Essays on distance metric learning. Unpublished PhD thesis, University College London, London.