

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 PSEUDO-NONLINEAR DATA AUGMENTATION: A CON- STRAINED ENERGY MINIMIZATION VIEWPOINT

Anonymous authors

Paper under double-blind review

## ABSTRACT

We propose a simple yet novel data augmentation method for general data modalities based on energy-based modeling and principles from information geometry. Unlike most existing learning-based data augmentation methods, which rely on learning latent representations with generative models, our proposed framework enables an intuitive construction of a geometrically aware latent space that represents the structure of the data itself, supporting efficient and explicit encoding and decoding procedures. We then present and discuss how to design latent spaces that will subsequently control the augmentation with the proposed algorithm. Empirical results demonstrate that our data augmentation method achieves competitive performance in downstream tasks compared to other baselines, while offering fine-grained controllability that is lacking in the existing literature.

## 1 INTRODUCTION

*Data augmentation* has advanced considerably in recent years, driven largely by the increasing use of generative models (Kingma & Welling, 2014; Chadebec et al., 2022; Antoniou, 2017; Trabucco et al., 2024) to meet the demand for larger and more diverse datasets (Feng et al., 2021; Wong et al., 2016). Beyond traditional domains such as images, these methods have been extended to a wide range of modalities. Despite their promise, however, generative-model-based augmentation faces several fundamental challenges. First, data augmentation is most valuable when training data is scarce, yet in such cases, we typically lack a pre-trained foundational model for the target domain. This creates a paradox: before we can augment the data, we must first train a generative model—reintroducing the very problem of limited data. Second, even when suitable foundational models are available, the computationally intensive nature of deep generative methods poses practical obstacles. Since effective augmentation often requires generating data of the same order as the original dataset, the cost of large-scale generation can quickly become prohibitive. Third, augmenting data with generative models raises concerns about their interpretability and controllability (Guidotti et al., 2018). Consequently, even when these models perform well, the lack of understanding of the underlying transformations of the augmented data makes it difficult to control the generated outputs, which poses a significant risk in the case of high-stakes scenarios (Rudin, 2019).

In this work, we propose a new data augmentation framework that addresses the challenges outlined above by providing a **learning-free**, **efficient**, and **controllable** algorithm applicable across diverse data modalities. Our approach builds on the well-established theory of *energy-based models* (Xie et al., 2016), together with recent advances in *log-linear models on partially ordered sets (posets)* (Sugiyama et al., 2016; 2017) and *information geometry* (Amari, 2016; Amari & Nagaoka, 2000; Ay et al., 2017). Conceptually, our framework resembles an autoencoder (Kingma & Welling, 2014). We begin by parametrizing data as discrete probability distributions on a *curved* statistical manifold  $\mathcal{S}$ . The data is then *encoded* into a chosen “latent space”  $\mathcal{B} \subseteq \mathcal{S}$  via *forward projection*. Within this latent space, simple augmentation procedures informed by the encoded data are applied. Finally, the resulting “augmented representation” is *backward projected* to the local data space  $\mathcal{D} \subseteq \mathcal{S}$ , yielding new augmented data. As our proposed algorithm exploits the duality of the projection in the statistical manifold  $\mathcal{S}$ , where it is linear in the manifold’s intrinsic coordinates yet non-linear in the ambient space, we hence term it *pseudo-non-linear* data augmentation.

This design offers three key advantages. First, it is learning-free: the sub-manifold structure is constructed explicitly, allowing direct control over the properties of the augmented data without the

need to train a generative model. Second, it is computationally efficient: both forward and backward projections can be formulated as convex programs and solved with efficient first-order methods such as gradient descent. Third, it provides controllability: leveraging prior knowledge about relationships among features, one can adjust the choice of  $S$  and the sub-manifold of projection to tailor the statistical properties of the augmented data. Our contributions are summarized as follows:

- We propose a novel framework for modeling structured data (e.g., tensors) within a statistical manifold using energy-based models. This framework captures the intrinsic geometry of data and enables the design of geometry-aware algorithms.
- We develop the *pseudo-nonlinear* data augmentation algorithm under this framework. The method is **learning-free**, **efficient**, and **controllable**, and applies broadly across different data modalities.
- We empirically validate the effectiveness of our approach, showing that it achieves competitive or superior performance compared to both generative-model-based baselines (e.g., autoencoders) and classical augmentation methods across multiple datasets and modalities.

## 2 RELATED WORK

### 2.1 LEARNING-BASED DATA AUGMENTATION

Data augmentation has proven to be effective in enhancing deep learning training by increasing dataset size, improving model robustness (Rebuffi et al., 2021), and introducing implicit regularization (Hernández-García & König, 2018). These techniques have been applied across various modalities, including text (Shorten et al., 2021; Feng et al., 2021; Li et al., 2022a) and images (Shorten & Khoshgoftaar, 2019; Mumuni & Mumuni, 2022; Wang et al., 2017). Due to the generality and the popularity, much of the recent progress in data augmentation for general modalities has been driven by advancements in generative models, such as autoencoders (Kingma & Welling, 2014; Chadebec et al., 2022), generative adversarial networks (Antoniou, 2017), and diffusion models (Trabucco et al., 2024). Despite the progress, to date, there is no fully satisfactory solution for the two challenges mentioned for generative-model-based data augmentation, i.e., efficiency and controllability. For example, the design of controllable GANs (Li et al., 2022b; She et al., 2021) and efficient flow-based models (Geng et al., 2025) remains an active area of research, and exploration in this direction is largely limited to specific domains such as images.

### 2.2 LEARNING-FREE DATA AUGMENTATION

Learning-free data augmentation methods (Maharana et al., 2022) that do not rely on generative models are particularly appealing because they construct an explicit, low-dimensional “latent space” in which data can be represented and augmented with fine-grained and intuitive control and interpretability. These latent spaces are typically derived from classical *dimension reduction* techniques such as Principal Component Analysis (PCA)(Wold et al., 1987) and Singular Value Decomposition (SVD)(Stewart, 1993). In general, these methods identify an optimal *linear* subspace and then perform augmentations that respect this (Euclidean) geometry of the ambient subspace. Beyond their simplicity and transparency, linear methods also provide useful insights; for instance, PCA reveals principal directions that capture the dominant modes of variation in the data.

A major limitation of linear dimension reduction for augmentation, however, is the difficulty of the *inverse* problem: reconstructing high-dimensional data from low-dimensional representations without a learned decoder is often non-trivial. Some works attempt to circumvent this issue by leveraging linear dimension reduction only indirectly for augmentation (Abayomi-Alli et al., 2020; Sirakov et al., 2024). Other approaches, such as mixup (Zhang et al., 2018), bypass dimension reduction entirely and operate directly in the original data ambient space. While popular in practice due to simplicity, these methods are typically heuristic, application-specific, and harder to generalize.

*Non-linear* generalizations of dimension reduction is often referred to as *manifold learning* (Meilă & Zhang, 2024), which provides an alternative route. Methods such as t-SNE (Hinton & Roweis, 2002; Van der Maaten & Hinton, 2008), Isomap (Tenenbaum et al., 2000), and UMAP (McInnes et al., 2018) aim to exploit the manifold hypothesis, which posits that high-dimensional data lie near a lower-dimensional manifold embedded in the ambient space. Their goal is to uncover this manifold and produce a smooth embedding that captures the intrinsic geometry of the data.

108 In principle, manifold learning could avoid the inverse problem by recovering a low-dimensional  
 109 manifold with minimal information loss, making it an attractive candidate for data augmentation. In  
 110 practice, however, this goal is rarely achieved without incorporating learning mechanisms (Duque  
 111 et al., 2020; Coifman & Lafon, 2006; Williams & Seeger, 2000; Vladymyrov & Carreira-Perpinán,  
 112 2013; Han et al., 2022). A workaround is to exploit the latent space learned by a model already  
 113 trained on a downstream task, as in manifold mixup (Verma et al., 2019). Yet these approaches again  
 114 sacrifice interpretability and controllability, compared to fully learning-free augmentation methods.  
 115

### 116 3 PRELIMINARY

#### 118 3.1 DUALLY-FLATNESS IN INFORMATION GEOMETRY

120 Information geometry studies the structure of *statistical manifolds*  $\mathcal{S}$  within the space of probability  
 121 distributions. In this paper, we are primarily concerned with the space of an exponential family  
 122  $\{p_\theta(x) \mid \theta \in \mathbb{R}^D\}$ , where each  $p_\theta$  denotes a probability density function parameterized by  $\theta$ . We  
 123 focus on the key concept in this field, *dually-flatness*, in this preliminary, while directing readers to  
 124 Appendix A and Amari (2016) for more comprehensive details.

125 The starting point is the observation that the *log-partition function*  $\psi(\theta)$  (also known as the *cumulant*  
 126 *generating function* in statistics and *free energy* in physics) of an exponential family with density  $p_\theta$   
 127 is convex in the *natural parameter*  $\theta \in \mathbb{R}^D$ . This convexity induces a natural coordinate system,  
 128  $\theta$ , on  $\mathcal{S}$ , defining both the Riemannian metric  $g = \nabla^2 \psi(\theta)$  and the Bregman divergence (Bregman,  
 129 1967)  $D_\psi(p_\theta, p_{\theta'})$ . With these structures, the manifold  $(\mathcal{S}, g)$  is flat, meaning that any curve  $\theta(t) =$   
 130  $at + b$  (where  $a, b \in \mathbb{R}^D$  are constants) is a geodesic and lies entirely within  $\mathcal{S}$ . This flatness is  
 131 known as *e-flatness*, and the geodesics are referred to as *e-geodesics* or *primal-geodesics*.

132 The dual structure arises from the *Legendre transform* (Legendre, 1787), which generates the dual  
 133 function  $\psi^*(\eta)$ , where  $\eta \in \mathbb{R}^D$  is the *expectation parameter*. This dual function is also convex,  
 134 giving rise to the expectation coordinate system  $\eta$ , the dual Riemannian metric  $g^*$ , and also the dual  
 135 Bregman divergence  $D_{\psi^*}$  which is the well-known Kullback-Leibler divergence  $D_{\text{KL}}$  (Eq.(3)). The  
 136 corresponding flatness is termed *m-flatness*, with *m-geodesics* or *dual-geodesics* as its geodesics.

137 **Remark 3.1.** *An e-flat (m-flat) sub-manifold can be defined by forcing linear constraints on the  $\theta$*   
 138 *coordinates ( $\eta$  coordinates)* (Amari, 2016, Chapter 2).

140 *Dually-flatness* emerges from the interplay between these two structures. Specifically, for any point  
 141 (distribution)  $p$  in  $\mathcal{S}$ , there is a unique point  $p^*$  on an *e*-flat sub-manifold  $\mathcal{B} \subseteq \mathcal{S}$  that minimizes  
 142 the dual Bregman divergence  $D_{\psi^*}(p, q) = D_{\text{KL}}(p, q)$  (Amari, 2016, Theorem 1.5). This process,  
 143 known as the *m-projection*, can be efficiently solved via convex optimization (Appendix A.2). The  
 144 dual holds when switching *e* and *m*. Projection is a central tool in information geometry with  
 145 profound implications for understanding the geometry of  $\mathcal{S}$ , which we will use later.

#### 147 3.2 STATISTICAL MANIFOLD ON POSETS

148 A set  $\Omega$  is a *partially ordered set (poset)* if it is equipped with a *partial order* “ $\leq$ ”, a relation  
 149 satisfying the following for all  $x, y, z \in \Omega$ : 1.)  $x \leq x$  (reflexivity); 2.)  $x \leq y$  and  $y \leq x$  implies  
 150  $x = y$  (antisymmetry); 3.)  $x \leq y$  and  $y \leq z$  implies  $x \leq z$  (transitivity). We focus on finite posets  
 151  $\Omega$  with a bottom element  $\perp$  such that  $\perp \leq x$  for all  $x \in \Omega$  to prevent some technical difficulties.

152 Given such a poset  $\Omega$ , consider a discrete random variable  $X$  with finite support  $\Omega$  with its prob-  
 153 ability mass function  $p: \Omega \rightarrow \mathbb{R}_{\geq 0}$ ,  $p(x) = \Pr(X = x)$  for  $x \in \Omega$ . For a discrete probability  
 154 distribution  $p$  over a poset  $\Omega$ , the *log-linear model on posets* recursively defines  $\theta: \Omega \rightarrow \mathbb{R}$  as  
 155  $\log p(x) = \sum_{y \leq x} \theta(y)$  for all  $x \in \Omega$  (Sugiyama et al., 2017). This model belongs to the expo-  
 156 nential family, with  $\theta$  corresponding to the natural parameters, except for  $\theta(\perp)$ , which coincides  
 157 with the partition function. Thus, all discrete probability distributions over  $\Omega$  form a  $(|\Omega| - 1)$ -  
 158 dimensional dually-flat statistical manifold  $\mathcal{S} := \{p: \Omega \rightarrow \mathbb{R}_{\geq 0} \mid \sum_{x \in \Omega} p(x) = 1\}$ , with the dual  
 159 coordinate systems  $(\theta, \eta)$  depend on the poset structure.

161 **Remark 3.2.** *One can think of  $\theta(x)$  for each element  $x \in \Omega$  as specifying the energy (i.e.,  $p(x)$ ) for  
 x, and the relation between  $\theta(x)$ 's on different elements x's is specified by the poset structure.*

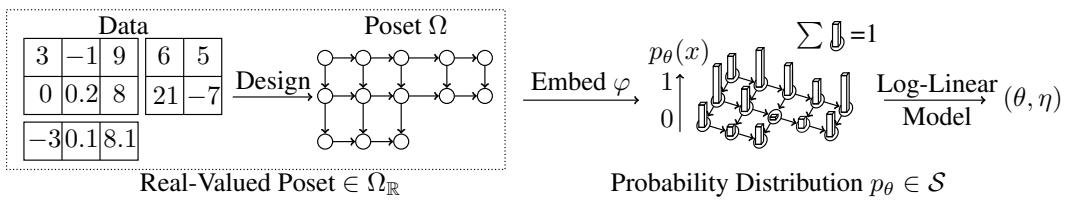
162 **4 PSUEDO-NONLINEAR DATA AUGMENTATION**  
 163

164 We first present our proposed framework in Section 4.1 and the projection algorithms in Section 4.2,  
 165 then we combine and apply them to data augmentation in Section 4.3. Finally, we discuss two  
 166 important features of the proposed method regarding **controllability** and **efficiency** in Section 4.4.  
 167 Throughout this section, we will use *positive tensors* as our running example (Example 4.1).

168  
 169 **4.1 LOG-LINEAR MODEL ON POSETS FRAMEWORK**  
 170

171 Motivated by Remark 3.2, by associating each element  $x \in \Omega$  with a feature dimension, we can  
 172 specify the geometry among features—i.e., design the poset structure—using prior knowledge or  
 173 natural structure present in the data. The resulting models define an energy-aware, curved statistical  
 174 space that faithfully reflects the prescribed relationships among features.

175 More specifically, given a dataset  $\{z_i\}_{i=1}^n$ , we first embed the data into a statistical manifold  $\mathcal{S}$   
 176 by leveraging the log-linear model on posets, which provides a geometric structure induced by the  
 177 energy-based modeling. The process works in three steps: 1.) models each  $z_i$  as a *real-valued*  
 178 *poset*; 2.) embeds the real-valued poset into the statistical manifold  $\mathcal{S}$  by viewing it as a probability  
 179 distribution; 3.) computes the corresponding two coordinate representations using the log-linear  
 180 model on posets. See Figure 1 for an illustration. We now explain each step in detail below.



181  
 182  
 183  
 184  
 185  
 186  
 187  
 188  
 189 Figure 1: Given structured data, we first design a poset  $\Omega$  that reflects the structure or the relationship  
 190 between features. The resulting real-valued poset is then embedded into the statistical manifold  $\mathcal{S}$  as  
 191 a discrete probability distribution  $p_\theta(x)$  via an embedding  $\varphi$ . Finally, the log-linear model on posets  
 192 provides the dually-flat coordinates  $(\theta, \eta)$  for  $p_\theta$ , which can be computed efficiently (Section 3.2).

193  
 194 **Real-Valued Poset.** In the typical machine learning pipeline, inputs are often constrained to be  
 195 vectors or matrices, which fail to accommodate more complex data. In contrast, posets are flexible  
 196 enough to capture data with structures, including vectors, matrices, or tensors. In general, any data  
 197 structure that admits a natural partial order can be modeled by a poset: for instance, a  $D$ -dimensional  
 198 vector  $z \in \mathbb{R}^D$  (i.e., 1<sup>st</sup>-order tensor) admits a structure that can be modeled by the poset  $\Omega := [D]$   
 199 with the partial order being the natural order between positive integers. Similarly, other common  
 200 data structures, such as matrices or tensors, can be treated in the same way. On the other hand, as  
 201 noted at the beginning of the section, a custom poset structure  $\Omega$  can also be specified, either with  
 202 or without the natural structure, to reflect prior knowledge of the relation between features.

203 We can then define the *real-valued poset*, which is a mapping from the poset  $\Omega$  to the set of real  
 204 numbers  $\mathbb{R}$  such that each entry (element) of the data structure (poset)  $x \in \Omega$  is associated with a  
 205 feature in  $\mathbb{R}$ . We denote the set of real-valued posets as  $\Omega_{\mathbb{R}}$ . In the  $D$ -dimensional vector example,  
 206  $\Omega = [D]$ , each element  $x \in \Omega$  corresponds to one of the  $D$  dimensions. Associating a real number  
 207 to each dimension then corresponds to an element in  $\Omega_{\mathbb{R}}$ .

208  
 209 **Embedding.** To embed the data  $\{z_i \in \Omega_{\mathbb{R}}\}_{i=1}^n$ , which are now modeled as real-valued posets,  
 210 to the statistical manifold  $\mathcal{S}$  which concerns with discrete probability distributions, we want an  
 211 embedding  $\varphi: \Omega_{\mathbb{R}} \rightarrow \mathcal{S}$  and its inverse  $\varphi^{-1}: \mathcal{S} \rightarrow \Omega_{\mathbb{R}}$ , such that  $\sum_{x \in \Omega} (\varphi(z_i))_x = 1$  for all  $z_i$  with  
 212  $\dim(\mathcal{S}) = D - 1$ .<sup>1</sup> In other words,  $\varphi(z_i)$  gives a probability mass function of a discrete random  
 213 variable over the poset, where  $(\varphi(z_i))_x$  is the probability of sampling  $x \in \Omega$  when sampled from  
 214  $\varphi(z_i)$ , representing the *energy* of the feature  $x$ . From the perspective of energy-based modeling, as  
 215 noted in Remark 3.2,  $\varphi$  can be naturally induced, e.g., for tabular frequency data, or customized to

<sup>1</sup>One can also consider the manifold of positive measures of dimension  $D$  and avoid the potential scaling issues. For simplicity, we omit this trivial extension in the presentation.

reflect the desired energy relationship between features. We note that in the latter case, a joint design of  $\varphi$  with the poset structure  $\Omega$  is often more expressive to describe a desirable energy structure.

**Dually-Flat Coordinates.** Finally, from the log-linear model on posets introduced in Section 3.2, for each  $z'_i := \varphi(z_i) \in \mathcal{S}$ , we associate the dually-flat coordinates  $\theta(z'_i) \in \mathbb{R}^{D-1}$  and  $\eta(z'_i) \in \mathbb{R}^{D-1}$ . Such coordinate systems are defined with respect to the underlying poset structure  $\Omega$ , providing a representation that captures the prescribed geometric structure among features of the data. We next illustrate the framework with one canonical example, positive tensors, in Example 4.1:

**Example 4.1** (Positive tensor). A  $d^{\text{th}}$ -order tensor  $T \in \mathbb{R}^{I_1 \times \dots \times I_d} =: \mathbb{R}^D$  is a multidimensional array with real entries for every index vector  $v = (i_1, \dots, i_d) \in [I_1] \times \dots \times [I_d] =: \Omega$  where for each  $k$ ,  $[I_k] := \{1, 2, \dots, I_k\}$  for a positive integer  $I_k$ . Tensors with entries all being positive are called positive tensors, denoted as  $P \in \mathbb{R}_{\geq 0}^{I_1 \times \dots \times I_d}$ .

A natural partial order “ $\leq$ ” between two index vectors  $v = (i_1, \dots, i_d)$ ,  $w = (j_1, \dots, j_d)$  for tensors is that  $v \leq w$  if and only if  $i_k \leq j_k$  for all  $k \in [d]$ . Finally, for positive tensors, a simple embedding  $\varphi: \mathbb{R}_{\geq 0}^{I_1 \times \dots \times I_d} \rightarrow \mathcal{S}$  where  $P' := \varphi(P): \Omega \rightarrow \mathbb{R}_{\geq 0}$  such that  $P'_v := P_v / \sum_{w \in \Omega} P_w$  for all  $v \in \Omega$  can be defined with a natural empirical inverse (see Remark 4.3).

Example 4.1 applies to many common data modalities. For example, a color image can be represented as a 3<sup>rd</sup>-order tensor, where the first two dimensions correspond to height and width, and the third dimension encodes the color channels. Time-series data likewise admit a tensorial poset structure, with the temporal dimension inducing a natural ordering among features at discrete time steps. Despite this flexibility, our framework comes with one notable limitation:

**Remark 4.2** (Invariance). Because the framework relies on specifying a partial order over the index set, it is not naturally equipped to model invariances under index permutations. This can introduce unwanted bias, for instance, in settings such as graphical data. Nonetheless, a key advantage of being learning-free and fully transparent is that the source of such bias is explicit, allowing one to identify and, when desired, mitigate it by appropriately modifying the modeling choices.

To conclude, we emphasize that the notion of *energy* in our framework acts as a modality-specific potential function that quantifies the stability or plausibility of a data configuration under the chosen poset embedding. Once the poset is specified, both the energy function and the induced manifold structure follow directly, yielding a clean and computationally convenient representation that supports subsequent computations and algorithm design, which we will discuss next.

## 4.2 FORWARD AND BACKWARD PROJECTION

We now demonstrate how to incorporate the log-linear on poset framework with projection theory to conduct data augmentation. Our algorithm mimics the architecture of autoencoders, focusing on two building blocks: the *encoder*  $\text{Enc}(\cdot)$  and the *decoder*  $\text{Dec}(\cdot)$ . First, for the encoding step, we formally explain how projection theory can be applied to perform dimension reduction and obtain compact representations within our framework. Next, for the decoding step, we introduce our proposed algorithm, termed *backward projection*, which serves as the *inverse* of dimension reduction. Figure 3 provides an intuitive geometric illustration for our proposed algorithms.

**Forward Projection.** The embedding from  $\Omega_{\mathbb{R}}$  to  $\mathcal{S}$  introduced in Section 4.1 maintains the dimensionality. To achieve dimension reduction, we leverage the projection theory: by projecting  $z'_i = \varphi(z_i)$  onto a low-dimensional flat sub-manifold called *base sub-manifold*  $\mathcal{B} \subseteq \mathcal{S}$  with  $\dim(\mathcal{B}) \ll \dim(\mathcal{S})$ , we obtain the desired *encoding*

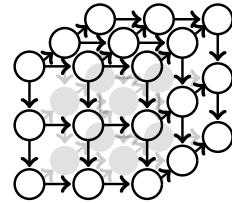


Figure 2: Natural poset structure of 3<sup>rd</sup>-order tensors in  $\mathbb{R}^{3 \times 3 \times 3}$ .

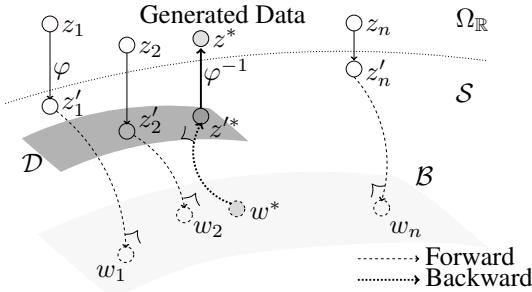


Figure 3: Illustration of forward and backward projection. Here,  $w_i$ : latent representation of the original data  $z_i$ , obtained from forward projection to  $\mathcal{B}$ ;  $w^*$ : augmented latent representation;  $w^* \mapsto z^*$ : backward projection to  $\mathcal{D}$ , obtained from the original data of the nearest neighbor(s) of  $w^*$  in the latent space.

270  $\text{Enc} := \text{Proj}_{\mathcal{B}} \circ \varphi: \Omega_{\mathbb{R}} \rightarrow \mathcal{B}$  that maps the data to a low-dimensional latent representation. Note  
 271 that the encoding  $\text{Enc}(\cdot)$  is smooth and well-defined as the projection is unique when  $\mathcal{B}$  is flat and  
 272 minimizing either the primal or the dual Bregman divergence, depending on either  $\mathcal{B}$  is  $e$ - or  $m$ -flat.  
 273

274 **Backward Projection.** One of the technical burdens is that the encoding  $\text{Enc}(\cdot)$  is not invertible,  
 275 hence a perfect decoding  $\text{Dec}(\cdot)$  is mathematically impossible, even when  $\text{Enc}(\cdot)$  only involves a  
 276 simple linear projection in Euclidean space. Here, we propose a simple, geometrically intuitive, and  
 277 data-centric solution that aims to find the *inverse* of the projection with theoretical guarantees.  
 278

279 The high-level intuition is simple: we assume that similar data will result in similar projections.  
 280 Hence, given a point in the low-dimensional latent representation space, we try to “project it back”  
 281 to approximate the original dataset by exploiting the fact that we have access to the inverse of the  
 282 dataset’s projection, which is the dataset itself. Specifically, we can artificially create a sub-manifold  
 283  $\mathcal{D}$  around a subset of the dataset that captures the local geometric structure of the dataset around the  
 284 latent representation, and *backward* project onto it.  
 285

286 Formally, assuming that we have access to the embedded dataset  $\{z'_i = \varphi(z_i)\}_{i=1}^n$  and their pro-  
 287 jected result  $\{w_i = \text{Proj}_{\mathcal{B}}(z'_i)\}_{i=1}^n$  for some flat base sub-manifold  $\mathcal{B}$ . To find the inverse of some  
 288 given point  $w^* \in \mathcal{B}$  assuming it comes from the projection on  $\mathcal{B}$ , we first find  $w^*$ ’s  $k$ -nearest  
 289 neighbors among  $w_i$ ’s, obtaining a size  $k$  index set  $N \subseteq [n]$  with  $|N| = k$ . Then we create  
 290 a flat sub-manifold  $\mathcal{D}$  called *local data sub-manifold* based on the pre-images  $z'_i$ ’s of these  $w_i$ ’s,  
 291 and project  $w^*$  on  $\mathcal{D}$  to obtain the *inverse*  $z^* = \text{Proj}_{\mathcal{B}}^{-1}(w^*) := \text{Proj}_{\mathcal{D}}(w^*)$ . The construction  
 292 of  $\mathcal{D}$  can be arbitrary, in particular, one can easily control the degree of freedom of the result-  
 293 ing  $z^*$ : for instance, from Remark 3.1, given the nearest neighbor  $z'_{i^*}$ , one can define an  $e$ -flat  
 294  $\mathcal{D} := \{\theta \in \mathbb{R}^{\dim(\mathcal{S})} \mid (\theta)_x = (\theta(z'_{i^*}))_x \text{ for some } x \in \Omega\}$  by fixing some indexes of  $\theta$  to be the  
 295 corresponding  $\theta$ -coordinate values of  $z'_{i^*}$ .  
 296

297 Algorithm 1 in Appendix B summarizes this procedure, which we termed *backward projection*.  
 298 With access to  $\text{Proj}_{\mathcal{B}}^{-1}(\cdot)$ , decoding is simply  $\text{Dec} := \varphi^{-1} \circ \text{Proj}_{\mathcal{B}}^{-1}: \mathcal{B} \rightarrow \Omega_{\mathbb{R}}$ . Algorithm 1  
 299 is a geometrically intuitive, data-centric algorithm with desirable theoretical guarantees such as  
 300 divergence minimizing when projecting on the constructed local data sub-manifold  $\mathcal{D}$ .  
 301

### 4.3 PSUEDO-NONLINEAR DATA AUGMENTATION

302 With all the building blocks in place, we can now formally describe the proposed data augmentation  
 303 algorithm, which consists of: 1.) encoding, 2.) augmenting, and 3.) decoding.  
 304

305 **Encoding.** As described in Section 4.2, the encoding  $\text{Enc} := \text{Proj}_{\mathcal{B}} \circ \varphi$  is simply a combination  
 306 of the embedding followed by a projection. Notation-wise, we write  $w_i := \text{Enc}(z_i)$ .  
 307

308 **Augmenting.** To generate an augmented data  $z^*$ , we first generate a new representation  $w^*$  in the  
 309 latent space, which in our case, is a pre-specified flat base sub-manifold  $\mathcal{B}$ . This  $w^*$  can be generated  
 310 in various ways, such as controlled perturbations of the original representations or a linear mixture  
 311 of two arbitrary original representations.  
 312

313 **Decoding.** As described in Section 4.2, the decoding  $\text{Dec} := \varphi^{-1} \circ \text{Proj}_{\mathcal{B}}^{-1}$  is simply a combination  
 314 of backward projection (Algorithm 1) with the inverse of the embedding. Notation-wise, we write  
 315  $z^* := \text{Dec}(w^*) = \varphi^{-1}(z^*)$  where  $z^* := \text{Proj}_{\mathcal{B}}^{-1}(w^*) := \text{Proj}_{\mathcal{D}}(w^*)$ .  
 316

317 The proposed method integrates the *nonlinear* forward and backward projections as encoding and  
 318 decoding, which we summarize the above in Algorithm 2 with an illustration given by Figure 3.  
 319

320 **Remark 4.3 (Positive tensor).** *With Algorithm 2, the empirical inverse  $\varphi^{-1}$  for the positive tensors*  
 321 *in Example 4.1 can be defined as the inverse of the average of the scaling among nearest neighbors.*  
 322

### 4.4 SUB-MANIFOLDS FOR POSITIVE TENSORS

323 It is evidence that the proposed method is **learning-free**. In this section, we describe how to con-  
 324 struct flat sub-manifolds (for  $\mathcal{B}$  and  $\mathcal{D}$ ) to **control** the augmentation process, and how these sub-  
 325 manifolds naturally admit **efficient projection algorithms**. For clarity and concreteness, we focus  
 326 our discussion on the case of positive tensors, while noting that the principles and arguments extend  
 327 to broader settings for other data modalities, poset structures, and the design of the energy potential.  
 328

**Designing Sub-Manifolds.** We start by discussing an intrinsic trade-off of choosing the dimension of  $\mathcal{B} \subseteq \mathcal{S}$  we aim to forward project on. Clearly, more information about the data is preserved after forward projection onto  $\mathcal{B}$  as  $\dim(\mathcal{B})$  increases. Hence, the quality of the backward projection  $\text{Proj}_{\mathcal{B}}^{-1}(\cdot)$  (Algorithm 1) increases along with  $\dim(\mathcal{B})$ . However, in the extreme case when  $\dim(\mathcal{B}) \approx \dim(\mathcal{S})$ , Algorithm 2 becomes less effective as the augmenting step now suffers from the curse of dimensionality. Previous studies on such a trade-off of choosing  $\dim(\mathcal{B})$  (Sugiyama et al., 2018; Ghalamkari et al., 2024) reveal how one should construct flat base sub-manifolds. In particular, the *many-body tensor approximation* (Ghalamkari et al., 2024; Derun & Sugiyama, 2025) aims to capture a *hierarchy* of mode interactions with different  $\dim(\mathcal{B})$  for positive tensors within the log-linear model on posets. Specifically, the  $\ell$ -body approximation considers projection on the following sub-manifold:

$$\mathcal{M}_\ell := \{\theta \in \mathbb{R}^{\dim(\mathcal{S})} \mid \theta_x = 0 \text{ for all } \mathbf{non} \text{ } \ell\text{-body parameters } x \in \Omega\}, \quad (1)$$

where the  $\ell$ -body parameter corresponds to  $\ell$  non-one indices, acting as a generalization of one-body and two-body parameters (Ghalamkari et al., 2024). Intuitively speaking, an  $\ell$ -body parameter captures the interaction among  $\ell$  different modes. Hence, when  $\mathcal{B} = \mathcal{M}_\ell$ , all interactions between modes of orders higher than  $\ell$  are neglected. This approach provides a principled way of designing the latent space with a clear understanding of what each dimension signifies.

On the other hand, a dual-like trade-off exists for the local data sub-manifold  $\mathcal{D}$ . Recall that the goal of backward projection is to project back to the “local data” space  $\mathcal{D}$  given by a set  $N$  of  $k$  nearest neighbors of a generated latent representation. When  $\dim(\mathcal{D})$  increases, backward projecting onto  $\mathcal{D}$  has a higher degree of freedom, which is desirable for data augmentation. However, in the extreme case when  $\dim(\mathcal{D}) \approx \dim(\mathcal{S})$ , the backward projection becomes unconstrained and potentially generates gibberish results. Hence, a natural construction of  $\mathcal{D}$  is to consider the “dual” notion of  $\mathcal{M}_\ell$ , where we now allow all  $\mathbf{non} \ell$ -body parameters to vary while fixing every  $\ell$ -body parameter to be the average of the  $\theta$  values among  $N$ :

$$\mathcal{M}_\ell^*(N) := \left\{ \theta \in \mathbb{R}^{\dim(\mathcal{S})} \mid \theta_x = \frac{1}{k} \sum_{i^* \in N} (\theta(z'_{i^*}))_x \text{ for all } \ell\text{-body parameters } x \in \Omega \right\}. \quad (2)$$

These two constructions offer a practical design choice for Algorithm 2 while providing the desired properties. For instance, By choosing an appropriate  $\ell$ , both  $\mathcal{M}_\ell$  and  $\mathcal{M}_\ell^*$  can capture specific information with desired degree of freedom.

**Efficient Projection.** For sub-manifolds  $\mathcal{M}_\ell$  designed for many-body approximation with  $B$  many non-fixed indexes (i.e.,  $\ell$ -body parameters), the projection can be efficiently computed via formulating the projection as a convex program over  $B$  variables that can be solved via gradient descent in polynomial time (in terms of  $B$ ). We note that the projection in the statistical manifold offers other desirable theoretical guarantees, such as minimizing the KL-divergence (corresponding to energy minimization) and uniqueness of the projection. Moreover, in the case of many-body approximation, it admits an efficient algorithmic implementation, since the gradient of the corresponding convex program has a closed-form, which makes the optimization extremely efficient. We refer the reader to Appendix A.2 for a detailed discussion.

## 5 EXPERIMENTS

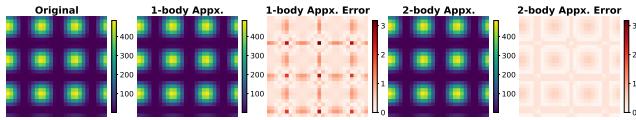
The structure of this section is as follows. Firstly, Section 5.1 demonstrates the energy aspect of the proposed framework, validating our motivation. Then, Section 5.2 illustrates the proposed pseudo-nonlinear data augmentation on real-world visual datasets. Finally, we apply the pseudo-nonlinear data augmentation on downstream tasks in Section 5.3.

### 5.1 LOG-LINEAR MODEL WITH ENERGY INTERACTION

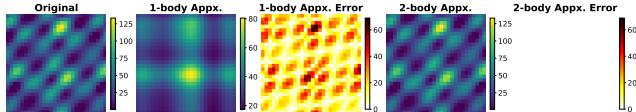
**Setup.** We demonstrate our proposed log-linear model on the posets framework with simple synthetic data under controlled feature interactions. The goal of this section is to illustrate both the energy-modeling perspective of our proposed framework, as well as how designing base sub-manifolds  $\mathcal{B}$  allows us to capture higher-order mode interactions among data.

In particular, we model both the 2-dimensional intensity data and 1-dimensional time-series data as a 2<sup>nd</sup>-order tensor as described in Example 4.1, where we apply reshaping to time-series data to capture higher-order mode dependencies across the temporal dimension.

Figure 4 consists of two main sections, (a) and (b), each showing a 2x5 grid of heatmaps. The columns are labeled 'Original', '1-body Appx.', '1-body Appx. Error', '2-body Appx.', and '2-body Appx. Error'. Each heatmap represents a 2D intensity distribution on a 10x10 grid. The color scale for the original data and approximations ranges from 0 to 400 for (a) and 0 to 125 for (b). The error heatmaps show the difference between the original data and the approximations, with a color scale ranging from 0 to 3 for (a) and 0 to 60 for (b). In section (a), the data is separable, showing two distinct clusters of high intensity. The 1-body approximations capture these clusters but fail to capture the interaction between them, while the 2-body approximations capture the interaction more accurately. In section (b), the data is non-separable, showing a more complex, interconnected pattern of high intensity. The 1-body approximations fail to capture this complexity, while the 2-body approximations capture it more accurately.



(a) Weak mode interaction with separable data.



(b) Strong mode interaction with non-separable data.

Figure 4: Mode interaction and approximation.

Following a similar setup, we conduct a case study in Figure 5 to illustrate the advantages of our method's induced geometry compared to the vanilla ambient space geometry. In particular, we **interpolate** between separable and non-separable data within both the *base sub-manifold* (Figure 5(a)) and *ambient space* (Figure 5(b)), and measure the *interaction energy*: the discrepancy between the

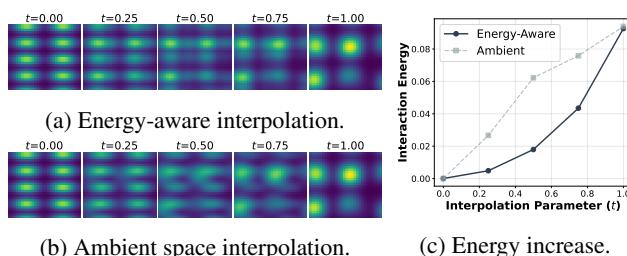


Figure 5: Interpolation energy in different geometries.

1-body approximation 2-body approximation of the interpolated data. The result is shown in Figure 5(c), where we see that energy-aware base sub-manifold consistently induces a lower interaction energy compared to the ambient geometry, indicating that our method, when operating on top of which, utilizes less energy compared to its ambient space geometry counterpart.

Finally, in Figure 6, we consider time series data with temporal dependencies emerging from higher-order interactions. A similar trend as Figure 4 emerges: the latent space geometry derived from a simple base sub-manifold results in a larger error at the place where higher-order mode interactions are presented (i.e., high frequency chirp pattern). In contrast, an appropriate base sub-manifold can capture such a mode interaction perfectly.

## 5.2 CONTROLLABILITY WITH CHOICES OF SUB-MANIFOLDS

**Setup.** As discussed in Section 4.4, constructing the sub-manifold carefully provides the essential controllability. We demonstrate this with MNIST (LeCun, 1998) and CIFAR (Krizhevsky & Hinton, 2009), where we first apply the log-linear model on posets for positive tensors (Example 4.1) by

<sup>2</sup>Equivalently, separable data is *linear* in log space where  $\log I(i, j) = \log f(i) + \log g(j)$ , while non-separable data has non-linear mode interactions in log space.

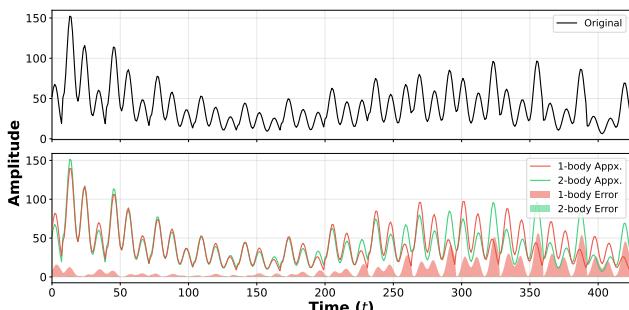


Figure 6: Approximation of temporal-dependent time-series.

normalizing features to be positive and reshaping the features as a tensor of suitable dimensions, then apply our proposed pseudo-nonlinear data augmentation afterwards. In particular: 1.) for MNIST, we let  $\mathcal{B} = \mathcal{M}_1$  and  $\mathcal{D} = \mathcal{M}_1^*$ , which correspond to preserving *shape* information; 2.) for CIFAR, we first carefully *reshape* the colored images to higher-order tensors, and let  $\mathcal{B} = \mathcal{M}_5$  and  $\mathcal{D} = \mathcal{M}_4^*$ , which correspond to preserving *fine-grained collective shape and color* information.

**Experiment.** Figures 7(a) and 7(d) show the results of the forward projection, while Figures 7(b) and 7(e) show the results of the backward projection of the latent representations sampled from the kernel density model  $M$  fitted on the results ( $\theta$ -coordinates) of Figures 7(a) and 7(d), respectively.

For MNIST, we see that the augmentation results (Figure 7(b)) with backward projection successfully reconstruct the digit structures, indicating that the essential *shape* information is indeed preserved and separated in the latent space  $\mathcal{B}$  to provide non-trivial neighbor information for constructing a sufficiently good  $\mathcal{D}$  for backward projection. Note that the local data sub-manifold  $\mathcal{D}$  has a dimension of 767, indicating a high degree of freedom for backward projection.

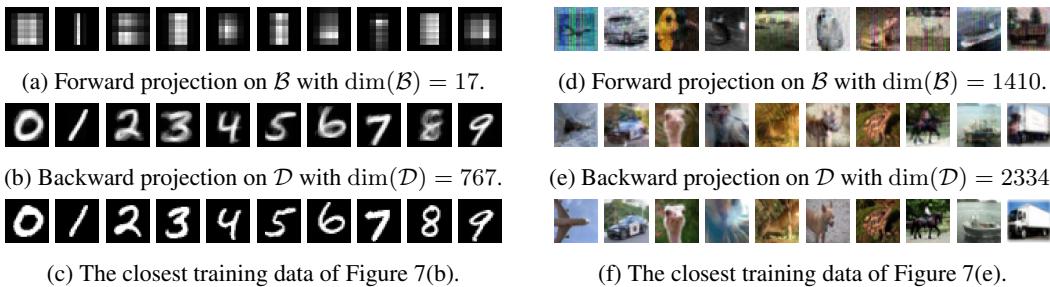


Figure 7: Results of MNIST (Left) and CIFAR-10 (Right) for Algorithm 2. The first row shows some representative latent representations from the dataset, while the second row shows the backward projection from an *augmented* latent representation.

More interesting results for CIFAR-10 are shown in Figures 7(d) to 7(f). By our proposed projection-based augmentation method, the *fine-grained shape and color* information is preserved. For instance, the third image, ostrich, successfully preserves the fine-grained shape and color relationship (e.g., colors for eyes and beak, and small pink flowers in the background), while the crude shape-to-color information is lost (e.g., colors for the background without shape details shift noticeably). The same trend can be observed consistently, validating the proposed method’s efficacy.

In practice, by carefully reshaping the data into higher-order tensors such that some modes of the tensors contain the essential relationship between features that one wishes to control, with many-body approximation, it is possible to construct suitable sub-manifolds that preserve the chosen information, providing a controllable augmentation of the original data via simple projection operations.

### 5.3 CLASSIFICATION PERFORMANCE

**Setup.** We apply our proposed pseudo-nonlinear data augmentation algorithm on downstream classification tasks across different datasets with various modalities, including image (MNIST and CIFAR-10), audio (Speech Commands (Warden, 2018)), and tabular data (Connectionist Bench (Sejnowski & Gorman, 1988), Taiwanese Bankruptcy (Journal, 2020), and Wine Quality (Cortez et al., 2009)). For each dataset, we train a classifier on both the original training set and an augmented training set, where the augmented portion corresponds to 20% of the original training size and is generated from the training data. The classifiers used are ResNet-18 (He et al., 2016) for CIFAR-10, M5 (Dai et al., 2017) for SpeechCommands, and a simple MLP for the remaining datasets. All of the models are evaluated on 20 randomly bootstrapped test subsets, each containing 50% of the original test data. Further details of the training setup are provided in Appendix C.1.

**Experiment.** Several baselines are compared against our proposed method, including both learning-based and learning-free methods, including: 1.) pseudo-nonlinear (**PNL**, ours), 2.) autoencoder-based augmentation (**AE**) (Kingma & Welling, 2014; Chadebec et al., 2022),<sup>3</sup> 3.) mixup

<sup>3</sup>More recent learning-based baselines, e.g., diffusion-based augmentation (Trabucco et al., 2024), are often computationally infeasible at our scale and can only be applied in a “few-shot” setting.

486

487

Table 1: Test accuracy of classifiers trained on different datasets.

Training Set	Dataset					
	MNIST	CIFAR-10	Speech Commands	Connectionist Bench	Taiwanese Bankruptcy	Wine Quality
<b>OG</b>	97.98 $\pm$ 0.19%	88.57 $\pm$ 0.57%	84.48 $\pm$ 0.50%	88.10 $\pm$ 8.58%	96.54 $\pm$ 0.56%	55.00 $\pm$ 1.69%
<b>OG<sup>STD</sup></b>	97.98 $\pm$ 0.24%	89.89 $\pm$ 0.44%	82.98 $\pm$ 0.50%	85.24 $\pm$ 7.66%	96.17 $\pm$ 0.57%	57.85 $\pm$ 1.81%
<b>OG<sup>PNL</sup></b>	97.91 $\pm$ 0.21%	88.07 $\pm$ 0.46%	84.35 $\pm$ 0.37%	93.81 $\pm$ 4.54%	96.53 $\pm$ 0.47%	59.03 $\pm$ 1.74%
<b>OG<sup>AE</sup></b>	97.97 $\pm$ 0.25%	88.36 $\pm$ 0.46%	83.13 $\pm$ 0.32%	82.86 $\pm$ 7.59%	95.92 $\pm$ 0.62%	57.23 $\pm$ 1.67%
<b>OG<sup>MU</sup></b>	96.45 $\pm$ 0.23%	86.60 $\pm$ 0.49%	81.85 $\pm$ 0.61%	89.29 $\pm$ 4.97%	96.55 $\pm$ 0.68%	57.76 $\pm$ 1.67%
<b>OG<sup>MMU</sup></b>	97.52 $\pm$ 0.30%	88.02 $\pm$ 0.39%	83.06 $\pm$ 0.54%	91.19 $\pm$ 5.06%	96.44 $\pm$ 0.53%	58.70 $\pm$ 1.74%

(MU) (Zhang et al., 2018) 4.) manifold mixup (MMU) (Verma et al., 2019), and 5.) standard augmentation (STD). For images, STD includes standard techniques such as random cropping, flipping, rotations, and affine transformations. For speech, STD combines random volume scaling, time stretching, MelSpectrogram conversion, frequency masking, and time masking (Park et al., 2019). For other data types, STD is implemented as Gaussian noise perturbations.

Results are summarized in Table 1, where we denote the original dataset as OG, and the augmented dataset using a method AG as OG<sup>AG</sup>. In most cases, the classifier trained on the augmented data achieves a better prediction accuracy compared to the one trained only on OG. One exception is the dataset associated with image modality, where all other baselines perform worse than OG and OG<sup>STD</sup>. A potential explanation is that the baseline augmentation algorithms for other baselines act more like regularizers, while OG<sup>STD</sup> explicitly forces the classifier to learn the robust visual features associated with modality-specific transformation, promoting properties such as rotation, location, and color invariance. Nevertheless, in general, PNL consistently outperforms other baselines on all datasets, demonstrating a competitive performance across modalities.

Importantly, we highlight the stability, a crucial but often overlooked goal in data augmentation. This is best illustrated by the Connectionist Bench dataset, which contains only 208 data points and 60 features, posing a significant challenge for stable training and testing for generalization. From Table 1, the standard deviations in accuracy on this dataset are noticeably higher for OG, OG<sup>STD</sup>, and OG<sup>AE</sup>. In contrast, our method achieves a substantially lower standard deviation of 4.54%, indicating improved consistency across bootstrap testing runs. Other data augmentation baselines, such as OG<sup>MU</sup> and OG<sup>MMU</sup>, also achieve a substantial improvement. Still, our method achieves the lowest standard deviation among all, and this trend is consistently presented across all datasets.

#### 5.4 ADDITIONAL EXPERIMENTS

We conduct a series of additional ablation studies and experiments in Appendices C.2 to C.5. Specifically, Appendices C.2 and C.3 assess the robustness of our proposed method and the impact of augmentation on downstream task performance. In contrast, Appendix C.4 provides a justification for the necessity of forward projection. Finally, Appendix C.5 explores the effect of different latent space design choices on augmentation outcomes, offering insights into how these design decisions can be leveraged to better control the augmentation process.

## 6 CONCLUSION

We introduced the *pseudo-nonlinear data augmentation* framework, which leverages information geometry and energy-based models to provide a **learning-free**, **efficient**, and **controllable** augmentation method. Our approach, grounded in the log-linear model on posets, endows data with a rich information-geometric structure induced from the designed energy potential, facilitating both geometric reasoning and principled algorithm design. A key component is the backward projection algorithm, which reverses dimension reduction in a geometrically intuitive and data-centric manner.

Through extensive experiments, we demonstrated the effectiveness of our method across diverse modalities and datasets. In particular, it enables scalable augmentation for general data types while offering controllability via the design of 1.) the base sub-manifold  $\mathcal{B}$ , 2.) the local data sub-manifold  $\mathcal{D}$ , and 3.) the poset structure  $\Omega$ . Empirically, our framework outperforms both learning-based and learning-free data augmentation baselines, even on common modalities such as images and audios.

540 ETHICS STATEMENT  
541

542 We have carefully reviewed and adhered to the ICLR Code of Ethics. This work does not involve  
543 human subjects, personally identifiable information, or sensitive data. All datasets used in our ex-  
544 periments are publicly available and commonly used in the community. We have followed standard  
545 practices for dataset preprocessing and model evaluation to avoid introducing unfair bias. To the  
546 best of our knowledge, our methodology and findings do not pose potential harm, nor do they raise  
547 concerns regarding privacy, security, discrimination, or misuse.

548  
549 REPRODUCIBILITY STATEMENT  
550

551 We are committed to ensuring the reproducibility of our results. The main paper provides a detailed  
552 description of the proposed method and the experimental setup. All hyperparameters, training proce-  
553 dures, and evaluation metrics are documented in Appendix C.1. To further facilitate reproducibility,  
554 we provide the source code and instructions for reproducing all experiments.

555 REFERENCES  
556

558 Olusola Oluwakemi Abayomi-Alli, Robertas Damaševičius, Michał Wiczorek, and Marcin  
559 Woźniak. Data augmentation using principal component resampling for image recognition by  
560 deep learning. In *Artificial Intelligence and Soft Computing: 19th International Conference,  
561 ICAISC 2020, Zakopane, Poland, October 12-14, 2020, Proceedings, Part II 19*, pp. 39–48.  
562 Springer, 2020.

563 Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–  
564 276, 1998.

566 Shun-ichi Amari. *Information Geometry and Its Applications*, volume 194 of *Applied Mathematical  
567 Sciences*. Springer Japan, 2016. ISBN 978-4-431-55977-1 978-4-431-55978-8. doi: 10.1007/  
568 978-4-431-55978-8.

569 Shun-ichi Amari and Hiroshi Nagaoka. *Methods of information geometry*, volume 191. American  
570 Mathematical Soc., 2000.

572 A Antoniou. Data augmentation generative adversarial networks. *arXiv preprint arXiv:1711.04340*,  
573 2017.

574 Nihat Ay, Jürgen Jost, Hông Vân Lê, and Lorenz Schwachhöfer. *Information Geometry*, volume 64  
575 of *Ergebnisse Der Mathematik Und Ihrer Grenzgebiete 34*. Springer International Publishing,  
576 2017. ISBN 978-3-319-56477-7 978-3-319-56478-4. doi: 10.1007/978-3-319-56478-4.

578 Lev M Bregman. The relaxation method of finding the common point of convex sets and its applica-  
579 tion to the solution of problems in convex programming. *USSR computational mathematics and  
580 mathematical physics*, 7(3):200–217, 1967.

581 Clément Chadebec, Elina Thibeau-Sutre, Ninon Burgos, and Stéphanie Allassonnière. Data aug-  
582mentation in high dimensional low sample size setting using a geometry-based variational au-  
583toencoder. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):2879–2896,  
584 2022.

586 Ronald R Coifman and Stéphane Lafon. Geometric harmonics: a novel tool for multiscale out-of-  
587 sample extension of empirical functions. *Applied and Computational Harmonic Analysis*, 21(1):  
588 31–52, 2006.

589 Paulo Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis. Wine Quality. UCI Machine Learning  
590 Repository, 2009. DOI: <https://doi.org/10.24432/C56S3T>.

592 Wei Dai, Chia Dai, Shuhui Qu, Juncheng Li, and Samarjit Das. Very deep convolutional neural  
593 networks for raw waveforms. In *2017 IEEE international conference on acoustics, speech and  
signal processing (ICASSP)*, pp. 421–425. IEEE, 2017.

594 Zhou Derun and Mahito Sugiyama. Optimal submanifold structure in log-linear models. In *The 41st*  
 595 *Conference on Uncertainty in Artificial Intelligence*, 2025. URL <https://openreview.net/forum?id=kVC4QV6Q4b>.

596

597 Andrés F Duque, Sacha Morin, Guy Wolf, and Kevin Moon. Extendable and invertible manifold  
 598 learning with geometry regularized autoencoders. In *2020 IEEE International Conference on Big*  
 599 *Data (Big Data)*, pp. 5027–5036. IEEE, 2020.

600

601 Steven Y. Feng, Varun Gangal, Jason Wei, Sarah Chandar, Soroush Vosoughi, Teruko Mitamura,  
 602 and Eduard Hovy. A survey of data augmentation approaches for NLP. In Chengqing Zong, Fei  
 603 Xia, Wenjie Li, and Roberto Navigli (eds.), *Findings of the Association for Computational Lin-*  
 604 *guistics: ACL-IJCNLP 2021*, pp. 968–988, Online, August 2021. Association for Computational  
 605 Linguistics. doi: 10.18653/v1/2021.findings-acl.84.

606

607 Zhengyang Geng, Mingyang Deng, Xingjian Bai, J Zico Kolter, and Kaiming He. Mean flows for  
 608 one-step generative modeling. *arXiv preprint arXiv:2505.13447*, 2025.

609

610 Kazu Ghalamkari, Mahito Sugiyama, and Yoshinobu Kawahara. Many-body approximation for  
 611 non-negative tensors. *Advances in Neural Information Processing Systems*, 36, 2024.

612

613 Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino  
 614 Pedreschi. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):1–42, 2018.

615

616 Henry Han, Wentian Li, Jiacun Wang, Guimin Qin, and Xianya Qin. Enhance explainability of  
 617 manifold learning. *Neurocomputing*, 500:877–895, 2022.

618

619 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
 620 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
 621 770–778, 2016.

622

623 Alex Hernández-García and Peter König. Further advantages of data augmentation on convolu-  
 624 tional neural networks. In *Artificial Neural Networks and Machine Learning–ICANN 2018: 27th*  
*625 International Conference on Artificial Neural Networks, Rhodes, Greece, October 4–7, 2018, Pro-  
 ceedings, Part I* 27, pp. 95–103. Springer, 2018.

626

627 Geoffrey E Hinton and Sam Roweis. Stochastic neighbor embedding. In S. Becker, S. Thrun, and  
 628 K. Obermayer (eds.), *Advances in Neural Information Processing Systems*, volume 15. MIT Press,  
 629 2002.

630

631 Taiwan Economic Journal. Taiwanese Bankruptcy Prediction. UCI Machine Learning Repository,  
 632 2020. DOI: <https://doi.org/10.24432/C5004D>.

633

634 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *2nd International Con-  
 ference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Confer-  
 ence Track Proceedings*, 2014.

635

636 Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Mas-  
 ter's thesis, Department of Computer Science, University of Toronto*, 2009.

637

638 Yann LeCun. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.

639

640 John M Lee. *Smooth manifolds*. Springer, 2012.

641

642 Adrien Marie Legendre. *Mémoire sur l'intégration de quelques équations aux différences partielles*.  
 Imprimerie royale, 1787.

643

644 Bohan Li, Yutai Hou, and Wanxiang Che. Data augmentation approaches in natural language pro-  
 645 cessing: A survey. *Ai Open*, 3:71–90, 2022a.

646

647 Chao Li, Kelu Yao, Jin Wang, Boyu Diao, Yongjun Xu, and Quanshi Zhang. Interpretable generative  
 648 adversarial networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(2):1280–  
 1288, Jun. 2022b. doi: 10.1609/aaai.v36i2.20015.

648 Kiran Maharana, Surajit Mondal, and Bhushankumar Nemade. A review: Data pre-processing and  
 649 data augmentation techniques. *Global Transitions Proceedings*, 3(1):91–99, 2022.  
 650

651 L. McInnes, J. Healy, and J. Melville. UMAP: Uniform Manifold Approximation and Projection for  
 652 Dimension Reduction. *ArXiv e-prints*, February 2018.  
 653

654 Marina Meilă and Hanyu Zhang. Manifold learning: What, how, and why. *Annual Review of  
 655 Statistics and Its Application*, 11, 2024.  
 656

657 Alhassan Mumuni and Fuseini Mumuni. Data augmentation: A comprehensive survey of modern  
 658 approaches. *Array*, 16:100258, 2022.  
 659

660 Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and  
 661 Quoc V. Le. SpecAugment: A simple data augmentation method for automatic speech recognition.  
 662 *Interspeech 2019*, Sep 2019. doi: 10.21437/interspeech.2019-2680. URL <http://dx.doi.org/10.21437/Interspeech.2019-2680>.  
 663

664 Sylvestre-Alvise Rebuffi, Sven Gowal, Dan Andrei Calian, Florian Stimberg, Olivia Wiles, and  
 665 Timothy A Mann. Data augmentation can improve robustness. *Advances in Neural Information  
 666 Processing Systems*, 34:29935–29948, 2021.  
 667

668 Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint  
 669 arXiv:1609.04747*, 2016.  
 670

671 Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and  
 672 use interpretable models instead. *Nature machine intelligence*, 1(5):206–215, 2019.  
 673

674 Terry Sejnowski and R. Gorman. Connectionist Bench (Sonar, Mines vs. Rocks). UCI Machine  
 675 Learning Repository, 1988. DOI: <https://doi.org/10.24432/C5T01Q>.  
 676

677 Rui She, Pingyi Fan, Xiao-Yang Liu, and Xiaodong Wang. Interpretable generative adversarial  
 678 networks with exponential function. *IEEE Transactions on Signal Processing*, 69:3854–3867,  
 679 2021.  
 680

681 Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning.  
 682 *Journal of big data*, 6(1):1–48, 2019.  
 683

684 Connor Shorten, Taghi M Khoshgoftaar, and Borko Furht. Text data augmentation for deep learning.  
 685 *Journal of big Data*, 8(1):101, 2021.  
 686

687 Nikolay Metodiev Sirakov, Tahsin Shahnewaz, and Arie Nakhmani. Training data augmentation  
 688 with data distilled by principal component analysis. *Electronics*, 13(2):282, 2024.  
 689

690 Gilbert W Stewart. On the early history of the singular value decomposition. *SIAM review*, 35(4):  
 691 551–566, 1993.  
 692

693 Mahito Sugiyama, Hiroyuki Nakahara, and Koji Tsuda. Information decomposition on structured  
 694 space. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pp. 575–579, 2016.  
 695 doi: 10.1109/ISIT.2016.7541364.  
 696

697 Mahito Sugiyama, Hiroyuki Nakahara, and Koji Tsuda. Tensor Balancing on Statistical Manifold. In  
 698 *Proceedings of the 34th International Conference on Machine Learning*, pp. 3270–3279. PMLR,  
 699 2017.  
 700

701 Mahito Sugiyama, Hiroyuki Nakahara, and Koji Tsuda. Legendre decomposition for tensors. *Ad-  
 702 vances in Neural Information Processing Systems*, 31, 2018.  
 703

704 Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A global geometric framework for  
 705 nonlinear dimension reduction. *Science*, 290(5500):2319–2323, 2000. doi: 10.1126/science.290.  
 706 5500.2319.  
 707

708 Brandon Trabucco, Kyle Doherty, Max A Gurinas, and Ruslan Salakhutdinov. Effective data aug-  
 709 mentation with diffusion models. In *The Twelfth International Conference on Learning Repre-  
 710 sentations*, 2024. URL <https://openreview.net/forum?id=ZWzUA9zeAg>.

702 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine*  
 703 *learning research*, 9(11), 2008.

704

705 Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-  
 706 Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states.  
 707 In *International conference on machine learning*, pp. 6438–6447. PMLR, 2019.

708

709 Max Vladymyrov and Miguel Á Carreira-Perpiñán. Locally linear landmarks for large-scale mani-  
 710 fold learning. In *Machine Learning and Knowledge Discovery in Databases: European Confer-  
 711 ence, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III*  
 712 13, pp. 256–271. Springer, 2013.

713 Jason Wang, Luis Perez, et al. The effectiveness of data augmentation in image classification using  
 714 deep learning. *Convolutional Neural Networks Vis. Recognit*, 11(2017):1–8, 2017.

715 Pete Warden. Speech commands: A dataset for limited-vocabulary speech recognition. *arXiv*  
 716 *preprint arXiv:1804.03209*, 2018.

717

718 Christopher Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines.  
 719 *Advances in neural information processing systems*, 13, 2000.

720 Svante Wold, Kim Esbensen, and Paul Geladi. Principal component analysis. *Chemometrics and*  
 721 *intelligent laboratory systems*, 2(1-3):37–52, 1987.

722

723 Sebastien C Wong, Adam Gatt, Victor Stamatescu, and Mark D McDonnell. Understanding data  
 724 augmentation for classification: when to warp? In *2016 international conference on digital image*  
 725 *computing: techniques and applications (DICTA)*, pp. 1–6. IEEE, 2016.

726

727 Jianwen Xie, Yang Lu, Song-Chun Zhu, and Yingnian Wu. A theory of generative convnet. In  
 728 *International conference on machine learning*, pp. 2635–2644. PMLR, 2016.

729 Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical  
 730 risk minimization. In *International Conference on Learning Representations*, 2018.

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 A PROJECTION THEORY IN INFORMATION GEOMETRY  
757

758 We will assume some familiarity with the basic terminologies for manifold (Lee, 2012, Chapter 1,  
759 4). In particular, in this section, we explain the main concepts of information geometry used in  
760 this study, including natural parameters, expectation parameters, model flatness, and convexity of  
761 optimization. In the following, we consider only discrete probability distributions for simplicity and  
762 refer to Amari (2016) for more general cases.

763  
764 A.1  $(\theta, \eta)$ -COORDINATE AND GEODESICS  
765

766 Consider  $\mathcal{S}$  as the space of discrete probability distributions, which is a non-Euclidean space with the  
767 Fisher information matrix  $G$  as the metric. This metric measures the distance between two points,  
768 i.e., discrete probability distributions, in  $\mathcal{S}$ . In Euclidean space, the shortest path between two points  
769 is a straight line, while in a non-Euclidean space, such a shortest path is called a *geodesic*. In the  
770 space  $\mathcal{S}$ , two kinds of geodesics can be introduced: *e*-geodesics and *m*-geodesics. For two points  
771  $p_1, p_2 \in \mathcal{S}$ , *e*- and *m*-geodesics are defined as

772  $\{r_t \mid \log r_t = (1-t) \log p_1 + t \log p_2 - \phi(t), 0 \leq t \leq 1\}, \quad \{r_t \mid r_t = (1-t)p_1 + tp_2, 0 \leq t \leq 1\},$   
773 respectively, where  $\phi(t)$  is a normalization factor to keep  $r_t$  to be a distribution.

774 We can parameterize distributions  $p \in \mathcal{S}$  by parameters known as *natural parameters*. In Sec-  
775 tion 3.2, we have described the relationship between a distribution  $p$  and a natural parameter vector  
776  $\theta \in \mathbb{R}^{D-1}$  for a discrete probability distribution over a sample space of  $D$  elements in the log-linear  
777 model. The natural parameter  $\theta$  serves as a coordinate system of  $\mathcal{S}$ , since any distribution in  $\mathcal{S}$  is  
778 specified by determining  $\theta$ . Furthermore, we can also specify a distribution  $p$  by its expectation  
779 parameter vector  $\eta \in \mathbb{R}^{D-1}$ , which corresponds to expected values of the distribution and an alter-  
780 native coordinate system of  $\mathcal{S}$ . More explicitly, the definition of the expectation parameter  $\eta$  is  
781 defined as  $\eta_x = \sum_{y \geq x} p(y)$  for  $x \in \Omega$ , and  $\eta_\perp = 1$ , where  $p(x)$  is the probability mass function  
782 of  $p$  over the sample set  $\Omega$ , which is assumed to be a poset. The  $\theta$ -coordinates and  $\eta$ -coordinates  
783 are orthogonal with each other, which means that the Fisher information matrix  $G$  has the following  
784 property,  $G_{u,v} = \partial \eta_u / \partial \theta_v$ , and  $(G^{-1})_{u,v} = \partial \theta_u / \partial \eta_v$ . *e*- and *m*-geodesics can also be described  
785 using these parameters as follows:

$$786 \{ \theta_t \mid \theta_t = (1-t)\theta^{p_1} + t\theta^{p_2}, 0 \leq t \leq 1 \}, \quad \{ \eta_t \mid \eta_t = (1-t)\eta^{p_1} + t\eta^{p_2}, 0 \leq t \leq 1 \},$$

787 where  $\theta^p$  and  $\eta^p$  are  $\theta$ - and  $\eta$ -coordinate of a distribution  $p \in \mathcal{S}$ .

789 A.2 FLATNESS, PROJECTION, AND ITS OPTIMIZATION  
790

791 A subspace is called *e-flat* when any *e*-geodesic connecting two points in a subspace is included  
792 in the subspace. The vertical descent of an *m*-geodesic from a point  $p \in \mathcal{S}$  onto *e*-flat subspace  
793  $\mathcal{B}_e$  is called *m*-projection. Similarly, *e*-projection is obtained when we replace all *e* with *m* and  
794 *m* with *e*. The flatness of subspaces guarantees the uniqueness of the projection destination. The  
795 projection destination  $\bar{p}$  or  $\tilde{p}$  obtained by *m*- or *e*-projection onto  $\mathcal{B}_e$  or  $\mathcal{B}_m$  minimizes the following  
796 KL divergence

$$797 \bar{p} = \arg \min_{q \in \mathcal{B}_e} D_{\text{KL}}(p, q), \quad \tilde{p} = \arg \min_{q \in \mathcal{B}_m} D_{\text{KL}}(q, p),$$

798 where the KL divergence from discrete distributions  $p \in \mathcal{S}$  to  $q \in \mathcal{S}$  is given as

$$800 D_{\text{KL}}(p, q) = \sum_{x \in \Omega} p(x) \log \frac{p(x)}{q(x)}, \quad (3)$$

803 where  $p(x)$  and  $q(x)$  are the probability mass functions of  $p$  and  $q$ , respectively. A subspace with  
804 some of its natural parameters fixed at 0 is *e*-flat (Amari, 2016, Chapter 2), which is obvious from  
805 the definition of *e*-flatness. More generally, any subspace  $\mathcal{B}$  resulting from linear constraints on  
806 the natural parameter is *e*-flat. Similarly, any subspace  $\mathcal{B}$  resulting from linear constraints on the  
807 expectation parameter is *m*-flat. When a space is *e*-flat and *m*-flat at the same time, we say that the  
808 space is *dually-flat*. The set of discrete probability distributions  $\mathcal{S}$  is dually-flat.

809 Both *e*- and *m*-flatness guarantee that the cost functions to be optimized in Eq. (3) are convex.  
Therefore, *m*- and *e*-projection onto an *e*- or *m*-flat subspace can be implemented by a gradient

method using a second-order gradient. This second-order gradient method is known as the *natural gradient method* (Amari, 1998). The Fisher information matrix  $G$  appears by second-order differentiation of the KL divergence. For instance, given  $p$  and an  $e$ -flat subspace  $\mathcal{B}_e$ , the optimization problem  $\bar{p} = \arg \min_{q \in \mathcal{B}_e} D_{\text{KL}}(p, q)$  can be efficiently solved via gradient descent with second-order derivative by the update rule  $\theta_{t+1} = \theta_t - G^{-1}(\eta_t - \eta^p)$ , where  $G \in \mathbb{R}^{D \times D}$  is the Hessian matrix, and  $\partial D_{\text{KL}}(P, Q)/\partial \theta = \eta - \eta^p$  is the derivative of the KL divergence. The updated natural parameters  $\theta_{t+1}$  can then be used to construct  $q_{t+1} \in \mathcal{B}_e$  that is closer to the destination  $\bar{p}$  along with the  $e$ -geodesic from  $q_t$  to  $\bar{p}$ . By repeating this process until convergence, we can always find the global optimal solution. A similar algorithm can be implemented for the other case, i.e.,  $\tilde{p} = \arg \min_{q \in \mathcal{B}_m} D_{\text{KL}}(q, p)$ .

We make some remarks on the optimization of many-body approximation (Ghalamkari et al., 2024) that we omit in Section 4.4, which is a specific case of the above discussion.

**Example A.1** (Many-body approximation). For  $\mathcal{B}_e = \mathcal{M}_\ell$  defined in Eq.(1):<sup>4</sup>

1. Convexity and uniqueness: *The solution of the many-body approximation is always unique, and the objective function of the many-body approximation is convex (Ghalamkari et al., 2024, Theorem 1). In particular, the many-body approximation is a maximum likelihood estimation that approximates a non-negative tensor, which is regarded as an empirical distribution, by an extended Boltzmann machine without hidden variables.*
2. Computational complexity: *The computational complexity of the many-body approximation for  $\mathcal{B}_e = \mathcal{M}_\ell$  with  $B$  many non-fixed indexes (i.e.,  $\ell$ -body parameters) is  $O(T|B|^3)$ , where  $T$  is the number of iterations of the optimization. This is because the overall complexity is dominated by the update of  $\theta$ , which includes matrix inversion of  $G$ , and the complexity of computing the inverse of an  $n \times n$  matrix is  $O(n^3)$ . As an alternative, one can appeal to first-order methods such as gradient descent, which will then reduce the complexity to  $O(T|B|)$ , where  $|B|$  corresponds to computing the gradient.*  
*Note that this complexity can be reduced if one reshapes tensors so that the size of each mode becomes small. We explore this idea further in Appendix C.5.*

## B OMITTED DETAILS FROM SECTION 4

We provide the pseudocode for the proposed algorithm in Algorithms 1 and 2.

---

### Algorithm 1: Backward projection

---

**Data:** A representation  $w^* \in \mathcal{B}$ ,  $\varphi$ -embedded dataset  $\{z'_i\}_{i=1}^n$  with projection  $\{w_i\}_{i=1}^n$  on  $\mathcal{B}$ ,  $k \in \mathbb{N}$   
**Result:** Backward projected data  $z'^*$

```

 $N \leftarrow \text{Nearest-Neighbor}(k, w^*, \{w_i\}_{i=1}^n)$ 
 $\mathcal{D} \leftarrow \text{Construct-Sub-Manifold}(\{z'_i\}_{i \in N})$ 
 $z'^* \leftarrow \text{Proj}(w^*, \mathcal{D})$ 
return  $z'^*$ 

```

---

<sup>4</sup>One can consider  $\mathcal{B}_m$  with Eq.(1) being defined w.r.t. the  $\eta$ -coordinate system, and similar remarks hold.

864

**Algorithm 2:** Pseudo-non-linear data augmentation

865

**Data:** A dataset  $\{z_i\}_{i=1}^n$ , embedding  $\varphi: \Omega_{\mathbb{R}} \rightarrow \mathcal{S}$ ,  $k \in \mathbb{N}$ , flat base sub-manifold  $\mathcal{B}$ , size  $m \in \mathbb{N}$ 

866

**Result:** A generated dataset  $\{z_j^*\}_{j=1}^m$  of size  $m$ 

867

```

for  $i = 1, \dots, n$  do // Encoding
   $z'_i \leftarrow \varphi(z_i)$ 
   $w_i \leftarrow \text{Proj}_{\mathcal{B}}(z'_i, \mathcal{B})$  //  $w = \text{Enc}(z) = \text{Proj}_{\mathcal{B}} \circ \varphi(z)$ 
for  $j = 1, \dots, m$  do // Augmenting
   $w_i^* \leftarrow \text{Augment}(\{w_i\}_{i=1}^n, \mathcal{B})$ 
for  $j = 1, \dots, m$  do // Decoding
   $z_j'^* \leftarrow \text{Back-Proj}(w_j^*, \{z'_i\}_{i=1}^n, \{w_i\}_{i=1}^n, k)$  // Algorithm 1
   $z_j^* \leftarrow \varphi^{-1}(z_j'^*)$  //  $z^* = \text{Dec}(w^*) = \varphi^{-1} \circ \text{Proj}_{\mathcal{B}}^{-1}(w^*)$ 
return  $\{z_j^*\}_{j=1}^m$ 

```

868

869

**C OMITTED DETAILS FROM SECTION 5**

870

871

In Appendix C.1, we provide the details of experimental setup omitted in Section 5, and Appendices C.3 to C.5 consists of additional experiments.

872

873

**C.1 DETAILS OF EXPERIMENTAL SETUP**

874

875

In this section, we provide the details of each dataset and other experimental setups with relevant explanations.

876

877

**Datasets.** First, we summarize the details of each dataset in Table 2 and relevant parameters for applying the log-linear model on posets and Algorithm 2.

878

879

Table 2: Summary of each dataset and the corresponding setups of Algorithm 2.

	Dataset					
	MNIST	CIFAR-10	Speech Commands	Connectionist Bench	Taiwanese Bankruptcy	Wine Quality
<b>Train Size</b>	60,000	60,000	84,848	166	5,455	5,197
<b>Test Size</b>	10,000	10,000	4,890	42	1,364	1,300
<b>Augment Size</b>	10,000	10,000	7,000	32	1,090	1,036
<b>Class</b>	10	10	35	2	2	7
<b>Feature</b>	784	3,072	16,000 $\searrow$ 4,000	60	95	11
<b>Poset <math>\Omega</math></b>	$\mathbb{R}_{\geq 0}^{7^2 \times 2^4}$	$\mathbb{R}_{\geq 0}^{2^{10} \times 3}$	$\mathbb{R}_{\geq 0}^{2^5 \times 5^3}$	$\mathbb{R}_{\geq 0}^{2^2 \times 3 \times 5}$	$\mathbb{R}_{\geq 0}^{5 \times 19}$	$\mathbb{R}_{\geq 0}^{2^2 \times 3}$
<b>Base <math>\mathcal{B}</math> (dim)</b>	$\bar{\mathcal{M}}_1^{(17)}$	$\bar{\mathcal{M}}_5^{(1,410)}$	$\bar{\mathcal{M}}_2^{(136)}$	$\bar{\mathcal{M}}_1^{(9)}$	$\bar{\mathcal{M}}_1^{(23)}$	$\bar{\mathcal{M}}_2^{(10)}$
<b>Local Data <math>\mathcal{D}</math> (dim)</b>	$\mathcal{M}_1^{*(767)}$	$\mathcal{M}_4^{*(2,334)}$	$\mathcal{M}_3^{*(3,430)}$	$\mathcal{M}_2^{*(30)}$	$\mathcal{M}_1^{*(72)}$	$\mathcal{M}_1^{*(7)}$
<b>Bandwidth</b>	0.05	0.05	0.05	0.05	0.05	0.05
<b>Neighbor <math>k</math></b>	8	3	3	2	5	10

890

Note that MNIST holds a CC BY-SA 3.0 license, CIFAR-10 is released with a MIT license, and Speech Commands is released with a CC BY 4.0 license. Finally, all the UCI datasets (Connectionist Bench, Taiwanese Bankruptcy, and Wine Quality) are licensed under CC-BY 4.0. We now break each group down and explain it in detail:

891

1. The first group consists of basic dataset statistics. The first three datasets (MNIST, CIFAR-10, and Speech Commands) come with a default train/test split; for the last three UCI datasets, since there is no default train-test split, we take 80% of the whole dataset as the training set, and the remaining 20% as the test dataset. **Augment Size** reports the size of the augmented data, which is roughly 20% of **Train Size**, off by some rounding errors since we assume we augment the same amount of data for each class.

892

In all cases, the full training set is used to train the classifier when evaluating the classification performance in Section 5.3, and also to train our data augmentation baseline (i.e., autoencoder) for comparison. The only exception is that for MNIST and Speech Commands, we choose an equal number of samples for every class when doing our pseudo-non-linear data augmentation for implementation convenience.

918 Finally, due to the extremely high dimensionality of Speech Commands (16000), we down-  
 919 sampled each data to 4000 dimensions in the entire experiment due to the computational con-  
 920 straint: solving 16000-dimension convex programs is infeasible in terms of the memory require-  
 921 ment.

922 2. The second group consists of the poset structure we impose on each dataset when applying the  
 923 log-linear model for positive tensors. Since our ultimate goal is to utilize the many-body ap-  
 924 proximation (Eqs. (1) and (2)), by reshaping the feature vector into a high-order tensor, a finer  
 925 hierarchy of projection can be obtained. Hence, in all experiments, we reshape the feature w.r.t.  
 926 the prime-number factorization of the number of features. For instance, an MNIST image is  
 927 in  $\mathbb{R}_{\geq 0}^{28 \times 28}$ , and we reshape it into a tensor of shape (7, 2, 2, 7, 2, 2), giving it a 6<sup>th</sup>-order tensor  
 928 structure  $\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2}$ . For notation convenience, in **Poset**  $\Omega$ , we overload  $\mathbb{R}_{\geq 0}^D$  to indicate the  
 929 natural poset structure introduced in Example 4.1, and compress the repeated prime factors in the  
 930 exponent. We note that for the Wine Quality dataset, since the feature dimension is originally 11,  
 931 which is a prime, we artificially add 1 dimension by padding 0's, so we get a non-trivial prime  
 932 factorization.

933 Next, **Base  $\mathcal{B}$**  and **Local Data  $\mathcal{D}$**  report the corresponding construction of the base sub-manifold  
 934 and the local data sub-manifold using either Eq.(1) or Eq.(2) w.r.t. the given poset structure of that  
 935 dataset. The number in the parentheses reports the corresponding dimension of the constructed  
 936 sub-manifolds.

937 3. Finally, **Bandwidth** reports the bandwidth parameter we used when fitting the kernel density  
 938 model  $M$  in the generating step, and **Neighbor  $k$**  reports the number of nearest neighbors used  
 939 in Algorithm 1.

940 **Classification Performance.** In Section 5.3, we evaluate the augmentation methods by training  
 941 classifiers and evaluating their prediction accuracy. For MNIST and three UCI datasets, we consider  
 942 the simple 2-layer MLP, while we use the ResNet-18 for the other two datasets. All models are  
 943 trained by Stochastic Gradient Descent (SGD) (Ruder, 2016) with learning rate 0.1, momentum 0.9,  
 944 and weight decay  $5 \times 10^{-4}$ . A step size scheduler is utilized to reduce the learning rate by a factor  
 945 of 0.1 every 30 epochs until convergence.

946 For the two baseline data augmentation methods:

947 

- 948 • **Standard method (STD):** The standard image augmentation methods include random horizontal  
 949 flipping and cropping. For other modalities, the standard augmentation method corresponds to  
 950 adding Gaussian noise, which is sampled from  $\mathcal{N}(0, \sigma^2)$ , where  $\sigma$  is chosen to be 1/4 of the  
 951 minimum standard deviations over each feature dimension to make sure that the noise level is  
 952 reasonable.
- 953 • **Autoencoder (AE):** For UCI datasets, we consider a simple two-layer MLP encoder-decoder ar-  
 954 chitecture, where the dimension of the hidden layer is the same as  $\dim(\mathcal{B})$  indicated in Table 2,  
 955 with a bottleneck dimension being 3. For the image datasets, both the encoder and the decoder are  
 956 based on convolutional layers. The encoder uses a series of convolutional layers with increasing  
 957 feature map sizes to progressively downsample the input image, while the decoder mirrors this  
 958 structure with transposed convolutional layers to reconstruct the image from the latent representa-  
 959 tion.

960 In all experiments, the autoencoder is trained with the Adam optimizer with a learning rate  $10^{-3}$   
 961 under the mean squared error until convergence.

## 962 C.2 IMPACT OF THE SIZE OF AUGMENTATION

963 In this section, we conduct additional experiments on the impact of the ratio between the augmented  
 964 dataset size on the performance. Specifically, we consider the two image datasets (MNIST and  
 965 CIFAR-10) and the audio dataset (Speech Commands), and vary the size of the original size of the  
 966 augmented dataset, which is 20% of the original dataset size. The results are presented in Tables 3  
 967 to 5.

972  
973  
974  
975  
976  
977  
978  
979

Table 3: Impacts on the sizes of the augmentation dataset for MNIST.

AG	25%	50%	75%	100%
<b>None</b>	98.14%	98.08%	98.10%	98.12%
<b>STD</b>	92.08%	92.00%	92.02%	92.16%
<b>AE</b>	97.91%	97.86%	97.90%	97.97%
<b>PNL</b>	98.01%	98.07%	97.93%	97.55%

980  
981  
982  
983  
984  
985  
986

Table 4: Impacts on the sizes of the augmentation dataset for CIFAR-10.

AG	25%	50%	75%	100%
<b>None</b>	89.06%	89.32%	88.63%	88.83%
<b>STD</b>	90.42%	92.00%	90.60%	90.92%
<b>AE</b>	88.56%	88.76%	88.42%	86.49%
<b>PNL</b>	88.72%	88.32%	88.53%	88.67%

994  
995  
996  
997  
998  
999  
1000

Table 5: Impacts on the sizes of the augmentation dataset for Speech Commands.

AG	25%	50%	75%	100%
<b>None</b>	83.68%	83.26%	84.72%	84.16%
<b>STD</b>	84.34%	83.84%	84.98%	84.25%
<b>AE</b>	82.30%	83.49%	84.72%	85.21%
<b>PNL</b>	84.20%	84.61%	84.56%	84.43%

1008  
1009  
1010  
1011

### C.3 SENSITIVITY AND ROBUSTNESS

1012  
1013  
1014  
1015  
1016  
1017

We examine our proposed method’s robustness and sensitivity of the *bandwidth* used when fitting the kernel density model, and also the *number k of the nearest neighbors* used in Algorithm 1. For an easier visual inspection, we use MNIST in this section.

1018  
1019  
1020  
1021  
1022

**Bandwidth of Kernel Density Estimation Model.** Consider varying the bandwidth we use when fitting the kernel density model, ranging among  $\{0.01, 0.05, 0.1, 0.2, 0.5\}$ . The results are shown in Figure 8. We observe that Algorithm 2 is robust under different bandwidths when working with the kernel density estimation model in the generating step.

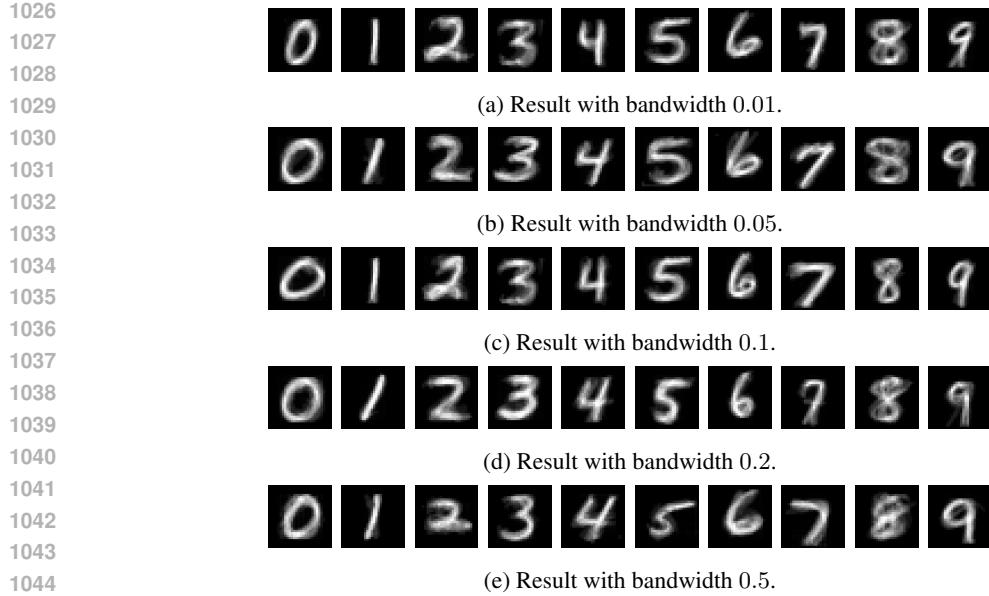
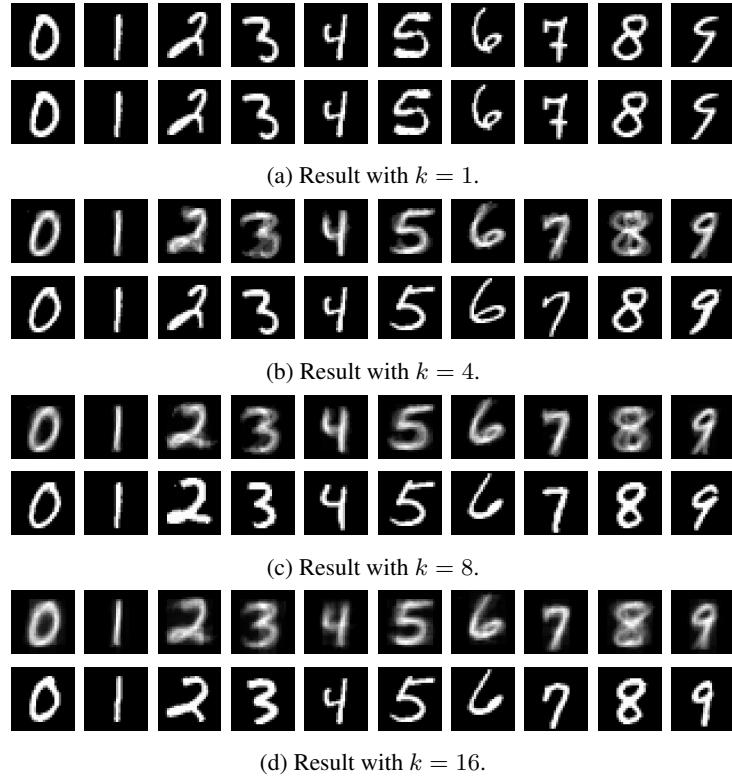


Figure 8: Augmented data via Algorithm 2 with different kernel density estimation bandwidths.

**Number of Nearest Neighbors.** Next, we consider ranging  $k$  among  $\{1, 4, 8, 16\}$ . The results are shown in Figure 9. Observe that when  $k$  is small, e.g., 1, the result of Algorithm 2 tends to overfit since the local sub-manifold  $\mathcal{D}$  in Algorithm 1 is defined using only the nearest neighbor. When  $k$  goes up, a non-trivial augmentation emerges, robust across different  $k$ 's.

Figure 9: (Top) Augmented data via Algorithm 2 with different  $k$ 's for Algorithm 1. (Bottom) The closest training data.

1080  
1081

## C.4 NECESSITY OF DIMENSION REDUCTION

1082  
1083  
1084

We demonstrate that dimension reduction, a key building block of our proposed method based on the intuition we have from autoencoder-like models, is necessary for Algorithm 2 to work. For an easier visual inspection, we use MNIST in this section.

1085  
1086  
1087

**Direct Fitting.** Naive perturbation-based data augmentation methods fall short of high-dimensional data due to the sparsity of the data. Figure 10 shows the results of directly fitting a kernel density estimation model on MNIST with 1000 samples.

1088  
1089  
1090  
1091  
1092  
10931094  
1095  
1096

Figure 10: (Top) Augmented data via directly fitting a kernel density estimation model with a bandwidth 30. (Bottom) The closest training data.

1097  
1098  
1099

Observe that even with a large bandwidth (30) to introduce variability, we only see a meaningless noisy perturbation on one of the exact training samples, indicating overfitting.

1100  
1101  
1102  
1103  
1104  
1105  
1106

**Local Data Sub-manifold.** A potential problem related to the necessity of dimension reduction is that, if  $\mathcal{D}$  captures too much local information about the data (i.e., with low dimension), backward projecting a random latent representation  $w^* \in \mathcal{B}$  might already suffice to augment the data in a non-trivial way, without the need for knowing the latent representations of the training dataset. To this end, consider sampling uniformly random latent representations within the empirical range we observed from the latent representations of the training data and perform Algorithm 1. The results are shown in Figure 11.

1107



(a) Result with  $k = 1$ .

1108  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119

(b) Result with  $k = 4$ .

1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130

(c) Result with  $k = 8$ .

1131  
1132  
1133

(d) Result with  $k = 16$ .

Figure 11: (Top) Augmented data on random latent representations via Algorithm 2 with different  $k$ 's for Algorithm 1. (Bottom) The closest training data.

1134 For  $k = 1$ , Figure 11(a) shows that, similar to Figure 9(a), it is possible to overfit one of the  
 1135 training data (i.e., the nearest neighbor of the randomly sampled latent representation). This is not  
 1136 surprising since the base sub-manifold is only of dimension 17, as the random latent representation  
 1137 is sufficiently close to one of the representations of the training data in  $\mathcal{B}$ , their backward projection  
 1138 result should not deviate too much. Furthermore, we observe the *fading effect*, which intuitively  
 1139 corresponds to *misspecification of the energy*, indicating that the sampled latent representation is  
 1140 fundamentally different from that of the dataset.

1141 As  $k$  increases, the benefit of getting informative and meaningful latent representations from the  
 1142 original dataset becomes clear. Specifically, we start to see *degeneration*: from unclear overlappings  
 1143 to collapsing (i.e., only a few pixels are showing). Intuitively speaking, it is because the random  
 1144 latent representation's nearest neighbors appear to be significantly different, hence failing to provide  
 1145 a consistent local data sub-manifold. For instance, in the extreme case when  $k = 16$ , the local data  
 1146 sub-manifold is completely not informative, resulting in collapsing. Overall, without dimension  
 1147 reduction, we will lose the reference of *realistic latent representations* provided by the original  
 1148 dataset, which leads to bad performance once we are beyond the trivial overfitting regime.

### 1149 C.5 CHOICES OF TENSOR STRUCTURE AND CONSTRUCTION OF SUB-MANIFOLDS

1150 In Section 5.2, we consider varying  $\ell$  for  $\mathcal{B} = \mathcal{M}_\ell$  with the tensor structure  $\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2}$  on  
 1151 the MNIST dataset. In this section, we further vary the tensor structure as well: in particular, we  
 1152 consider the tensor structure of the MNIST image being  $\mathbb{R}_{\geq 0}^{28 \times 28}$ ,  $\mathbb{R}_{\geq 0}^{7 \times 4 \times 7 \times 4}$ , and  $\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2}$ .

1153 **Remark C.1.** For notation convenience, we also write their corresponding poset structures as  
 1154  $\mathbb{R}_{\geq 0}^{28 \times 28}$ ,  $\mathbb{R}_{\geq 0}^{7 \times 4 \times 7 \times 4}$ , and  $\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2}$ , and further write the many-body approximation sub-  
 1155 manifold (Eq.(1)) as  $\mathcal{M}_\ell(\Omega)$  and its dual (Eq.(2)) as  $\mathcal{M}_\ell^*(N, \Omega)$  for a particular poset  $\Omega$  to empha-  
 1156 size the dependency.

1157 Finally, we consider ranging  $\ell$  from 1 to at most 4 where we neglect the degenerate case: for instance,  
 1158 in the case of  $\mathbb{R}_{\geq 0}^{28 \times 28}$ ,  $\mathcal{M}_2(\mathbb{R}_{\geq 0}^{28 \times 28}) = \mathcal{S}$  as there are only two modes for a matrix, therefore degen-  
 1159 erates to direct fitting which is not of interest (see Appendix C.4). Note that throughout this section,  
 1160 we fix the default local data sub-manifold to be  $\mathcal{D} = \mathcal{M}_1^*(N, \mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2})$  for consistency.

1161 The results for the finest structure,  $\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2}$ , are shown in Figure 12. As  $\ell$  grows, the forward  
 1162 projection results (*Top*) preserve the structure of the data better, subsequently improving the quality  
 1163 of the augmented data (*Bottom*). Similar trends can be found in the case of  $\mathbb{R}_{\geq 0}^{7 \times 4 \times 7 \times 4}$ , as shown in  
 1164 Figure 13.

1165 If we look at the results when using the original matrix structure  $\mathbb{R}_{\geq 0}^{28 \times 28}$  (Figure 14), some interest-  
 1166 ing comparisons can be made. Firstly, if we compare the augmentation results for  $\mathcal{B} = \mathcal{M}_1(\mathbb{R}_{\geq 0}^{28 \times 28})$   
 1167 (Figure 14 (*Bottom*)) with the finer structures counterparts, e.g., (Figure 12(a) (*Bottom*)) for  
 1168  $\mathcal{M}_1(\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2})$ , one can observe that the results are worse. However, the former requires  
 1169 more dimension ( $\dim(\mathcal{M}_1(\mathbb{R}_{\geq 0}^{28 \times 28})) = 55 > 17 = \dim(\mathcal{M}_1(\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2}))$ ) for the base sub-  
 1170 manifold. Similarly, the augmentation results with  $\mathcal{B} = \mathcal{M}_1(\mathbb{R}_{\geq 0}^{7 \times 4 \times 7 \times 4})$  (Figure 13(a) (*Bottom*))  
 1171 also achieve better performance with a lower base sub-manifold dimension.

1172  
 1173  
 1174  
 1175

1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187

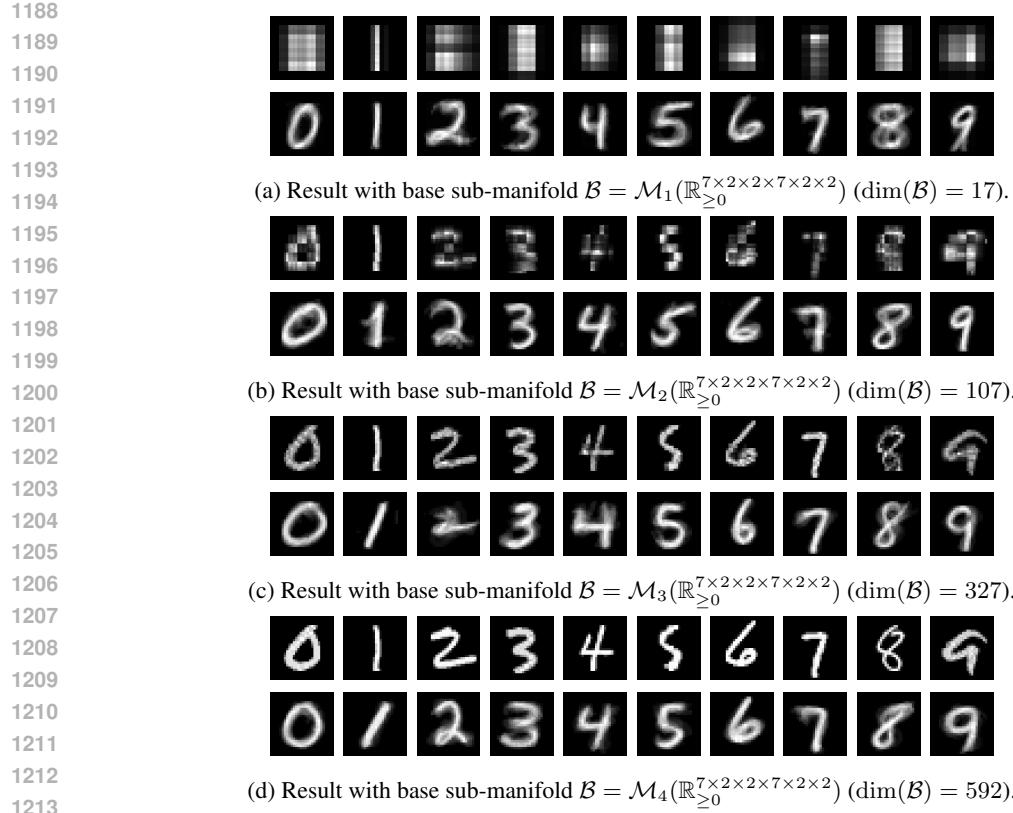


Figure 12: (Top) Forward projection on  $\mathcal{B} = \mathcal{M}_\ell(\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2})$ . (Bottom) Backward projection on  $\mathcal{D} = \mathcal{M}_1^*(\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2})$ .

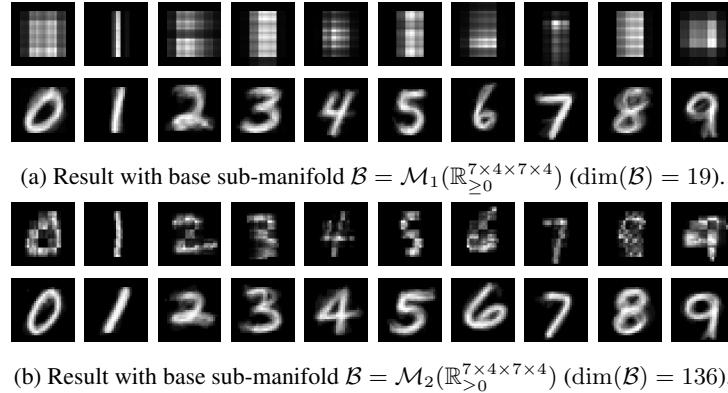


Figure 13: (Top) Forward projection on  $\mathcal{B} = \mathcal{M}_\ell(\mathbb{R}_{\geq 0}^{7 \times 4 \times 7 \times 4})$ . (Bottom) Backward projection on  $\mathcal{D} = \mathcal{M}_1^*(\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2})$ .

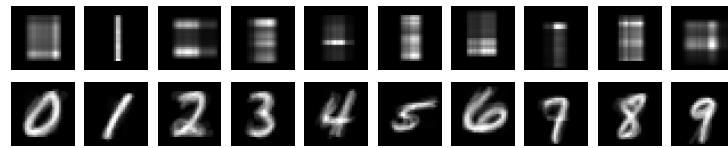


Figure 14: (Top) Forward projection on  $\mathcal{B} = \mathcal{M}_\ell(\mathbb{R}_{\geq 0}^{28 \times 28})$ . (Bottom) Backward projection on  $\mathcal{D} = \mathcal{M}_1^*(\mathbb{R}_{\geq 0}^{7 \times 2 \times 2 \times 7 \times 2 \times 2})$ .