

---

# Deep Divergence Learning

---

Kubra Cilingir<sup>1</sup> Rachel Manzelli<sup>1</sup> Brian Kulis<sup>1</sup>

## Abstract

Classical linear metric learning methods have recently been extended along two distinct lines: deep metric learning methods for learning embeddings of the data using neural networks, and Bregman divergence learning approaches for extending learning Euclidean distances to more general divergence measures such as divergences over distributions. In this paper, we introduce *deep Bregman divergences*, which are based on learning and parameterizing functional Bregman divergences using neural networks, and which unify and extend these existing lines of work. We show in particular how deep metric learning formulations, kernel metric learning, Mahalanobis metric learning, and moment-matching functions for comparing distributions arise as special cases of these divergences in the symmetric setting. We then describe a deep learning framework for learning general functional Bregman divergences, and show in experiments that this method yields superior performance on benchmark datasets as compared to existing deep metric learning approaches. We also discuss novel applications, including a semi-supervised distributional clustering problem, and a new loss function for unsupervised data generation.

## 1. Introduction

The goal of metric learning is to use supervised data in order to learn a distance function (or more general divergence measure) that is tuned to the data and task at hand. Classical approaches to metric learning are generally focused on the linear regime, where one learns a linear mapping of the data and then applies the Euclidean distance in the mapped space for downstream tasks such as clustering, ranking, and

classification (Davis et al., 2007; Weinberger & Saul, 2009; Goldberger et al., 2004). These methods, known as Mahalanobis metric learning approaches, have been analyzed theoretically, are scalable, and usually involve convex optimization problems that can be solved globally (Kulis, 2013; Bellet et al., 2015).

Classical metric learning methods have been extended along various axes; two important directions are deep metric learning and Bregman divergence learning. Deep metric learning approaches replace the linear mapping learned in Mahalanobis metric learning methods with more general mappings that are learned via neural networks (Hoffer & Ailon, 2015; Chopra et al., 2005). On the other hand, Bregman divergence methods replace the squared Euclidean distance with arbitrary Bregman divergences (Bregman, 1967), and learn the underlying generating function of the Bregman divergence via piecewise linear approximators (Siahkamari et al., 2019) or convex combinations of existing basis functions (Wu et al., 2009). These two extensions of classical metric learning are complementary and disjoint. For instance, Bregman divergence approaches can be utilized in scenarios where one needs to compare distributions (the well-known KL-divergence arises as a special case), but the learning problems are not directly applicable to the deep learning setting. Similarly, deep metric learning methods still employ Euclidean distances, and are thus not directly amenable to problems where one needs to compare distributions.

In this paper, we introduce a framework for studying Bregman divergences that can naturally be learned in the deep setting. Figure 1 gives a high-level overview of our approach, which we term as *deep Bregman divergences*, in comparison to existing metric learning approaches. These divergences are based on functional Bregman divergences (Frigyik et al., 2008), which were introduced as an extension of classical Bregman divergences but with functional inputs instead of vector inputs. In this functional setting, the underlying Bregman divergence is parameterized by a convex functional whose input itself is a function.

We first perform an analysis for the symmetric divergence case. In this setting, we prove a result about the form for any functional Bregman divergence and observe that many existing metric learning models can be seen to arise from special

---

<sup>1</sup>Department of Electrical and Computer Engineering, Boston University, Boston, Massachusetts, USA. Correspondence to: Kubra Cilingir <kubra@bu.edu>, Rachel Manzelli <manzelli@bu.edu>, Brian Kulis <bkulis@bu.edu>.

cases of this form. These include deep learning methods, classical linear metric learning methods, and kernel metric learning. There are also special cases that include moment-matching functions, which yields connections to the Wasserstein distance (Arjovsky et al., 2017), maximum mean discrepancy (MMD), and kernel MMD (Gretton, 2012).

We then turn our attention to the strictly more general case, where the divergences need not be symmetric; the KL-divergence is a classical example of such an asymmetric Bregman divergence. In this setting, we describe a framework for learning an arbitrary deep Bregman divergence. Our approach is based on appropriately parameterizing the convex functional governing the underlying Bregman divergence with a neural network, and learning the resulting parameters of that network.

We describe several applications of our proposed deep Bregman divergence framework. First, we can extend existing deep metric learning formulations to learn more general deep Bregman divergences. Second, since our divergences can naturally be applied to compare distributions on data, another application is in unsupervised generative learning, where the goal is to minimize a learned distributional divergence between real and generated data. In particular, we discuss connections to GAN models and describe some novel algorithms for unsupervised data generation. Third, we describe a semi-supervised distributional clustering problem. Here, the problem is to cluster data where each data point is represented as a distribution—for example, a movie’s rating may be represented as a distribution over user scores—using training data where we know whether pairs of distributions should be clustered together or not.

In all three of the above settings, we show empirical results that highlight the benefits of our framework. In particular, we show that learning asymmetric divergences offers performance gains over existing symmetric models on benchmark data, and achieve state-of-the-art classification performance in some settings. We also show that our clustering algorithm outperforms existing baselines on a simple proof-of-concept dataset as well as several human activity sensor data sets, and that our data generation results suggest that there may be value in further developing and studying new learned distributional divergence measures. Our code is available at [https://github.com/kubrac/Deep\\_Bregman](https://github.com/kubrac/Deep_Bregman).

## 2. Related Work

Much of the early work on metric learning focused on the *linear* setting, often referred to as Mahalanobis metric learning. In this setting, the goal is to learn a global linear transformation of the data and apply standard distances such as the Euclidean distance on top of the learned transformation. This is often expressed as learning a distance function of the

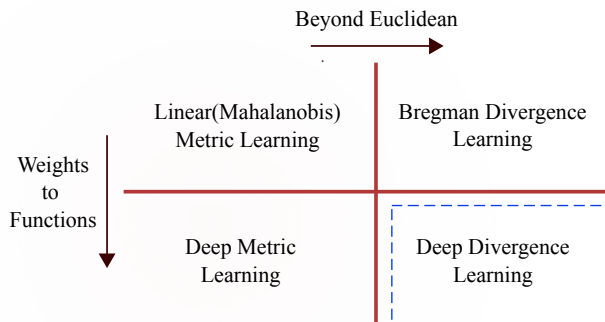


Figure 1. Overview of our framework in comparison to existing metric learning approaches. Deep Bregman divergences feature both the ability to learn divergences beyond Euclidean (such as divergences over distributions) while encompassing parameterizations that are amenable to deep learning architectures.

form  $d_A(\mathbf{x}, \mathbf{y}) = (\mathbf{x} - \mathbf{y})^T A (\mathbf{x} - \mathbf{y})$ , where  $A$  is a positive semi-definite matrix. This is equivalent to learning a linear transformation  $G$ , where  $A = G^T G$ , since  $d_A(\mathbf{x}, \mathbf{y}) = (\mathbf{x} - \mathbf{y})^T G^T G (\mathbf{x} - \mathbf{y}) = \|G\mathbf{x} - G\mathbf{y}\|_2^2$ . Examples of this approach to metric learning include MMC (Xing et al., 2003), MCML (Globerson & Roweis, 2005), LMNN (Weinberger & Saul, 2009), ITML (Davis et al., 2007), POLA (Shalev-Shwartz et al., 2004), LEGO (Jain et al., 2008), and others. See the surveys by Kulis (2013) and Bellet et al. (2015) for further references and details on some of these approaches. Note that one of the advantages of the linear approach is that one can often provide performance guarantees—for instance, a significant amount of work has gone into proving regret bounds in the online setting (Shalev-Shwartz et al., 2004), as well as generalization bounds (Bellet & Habrard, 2015; Cao et al., 2016) for some Mahalanobis metric learning models.

While linear methods are simpler and can often be analyzed theoretically, in practice it is often useful to learn other, non-linear, approaches to metric learning. For instance, one can show that many linear models can be appropriately adapted to run in kernel space (Jain et al., 2012; Chatpatanasiri et al., 2010). Another more recent approach to moving beyond linear metric learning is the Bregman divergence learning framework discussed in the introduction (Siahkamari et al., 2019; Wu et al., 2009). Here we move beyond learning Mahalanobis metrics, but instead focus on a strictly larger class of divergences that includes asymmetric divergences such as the KL-divergence, Itakura-Saito divergence, and others. These may be considered as non-linear approaches (since the resulting divergence does not involve linear transformations in general). The Bregman learning framework is thus more powerful than linear approaches but also remains well-principled: one can prove generalization bounds in this framework.

The third, and by far the most well-studied, approach to non-linear metric learning is known as deep metric learning, and involves learning a neural network to embed data into some new space, where standard distances such as Euclidean distance are used. If  $f$  is a function that maps an input  $\mathbf{x}$  to an embedding  $f(\mathbf{x})$ , then the resulting learned metric is typically  $\|f(\mathbf{x}) - f(\mathbf{y})\|_2^2$ . Several popular loss functions have been proposed to learn such a metric—the two main ones are the contrastive loss (Chopra et al., 2005) and the triplet loss (Hoffer & Ailon, 2015). Both utilize supervision (pairwise for the contrastive loss and relative constraints for the triplet loss) and use the learned distance  $\|f(\mathbf{x}) - f(\mathbf{y})\|_2^2$ . Moreover, there has been considerable follow-up work that explores how best to choose pairs or triples of points from a training set to achieve the best results (Hermans et al., 2017). There has also been work on deep metric learning using other losses, such as the angular loss (Wang et al., 2017) or the average precision for deep metric learning to rank (Cakir et al., 2019).

Our work also has ties to methods involving comparing distributions. Examples of such measures that are relevant to our work include the maximum mean discrepancy metric (also known as the integral probability metric) (Gretton, 2012), the kernel MMD, and the Wasserstein distance (Arjovsky et al., 2017). Several notions of divergences over distributions have been used for unsupervised data generation in GAN-type models, including the Jensen-Shannon divergence (Goodfellow et al., 2014), the Wasserstein distance (Arjovsky et al., 2017), and the MMD (Li et al., 2015; 2017).

### 3. Deep Bregman Divergences

We now turn our attention to functional Bregman divergences, the main tool for our learning problems. Our goal is two-fold: we first prove a result that characterizes the form for a symmetric functional Bregman divergence and show connections between this form and existing metric learning models. Second, we consider a parameterization for arbitrary functional Bregman divergences that will permit learning via neural networks.

#### 3.1. Bregman Divergences and Functional Bregman Divergences

A Bregman divergence is a generalized measure of distance between objects, parameterized by a strictly convex function  $\phi$  (Bregman, 1967). Let  $\phi : \Omega \rightarrow \mathbb{R}$ , where  $\Omega$  is a closed, convex set. The Bregman divergence with respect to  $\phi$  (for vector inputs) is defined as

$$D_\phi(\mathbf{x}, \mathbf{y}) = \phi(\mathbf{x}) - \phi(\mathbf{y}) - (\mathbf{x} - \mathbf{y})^T \nabla \phi(\mathbf{y}).$$

Note that the last term represents the derivative of  $\phi$  in the direction of  $\mathbf{x} - \mathbf{y}$ . Examples of Bregman divergences in-

clude the squared Euclidean distance, parameterized by  $\phi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$ ; the KL-divergence, parameterized by  $\phi(\mathbf{x}) = \sum_i x_i \log x_i$ ; and the Itakura-Saito distance, parameterized by  $\phi(\mathbf{x}) = -\sum_i \log x_i$ .

Bregman divergences arise in many settings in machine learning and related areas. In the study of exponential family distributions, there is a bijection between the class of regular Bregman divergences and regular exponential families (see Banerjee et al. (2005)). In optimization, Bregman divergences arise frequently; for instance, mirror descent utilizes Bregman divergences, and Bregman divergences were originally proposed as part of constrained optimization (Bregman, 1967). In the study of clustering, Bregman divergences offer a straightforward way to extend the k-means algorithm beyond the use of the squared Euclidean distance (Banerjee et al., 2005). A consequence of this is a way to cluster multivariate Gaussians in a k-means framework (Davis & Dhillon, 2006); we will use this algorithm as a baseline later in the paper.

More recently, Frigiyk et al. (2008) proposed and studied an extension to standard Bregman divergences called *functional Bregman divergences*, where instead of vector inputs, we compute a divergence between pairs of functions (or distributions). In this case, given two functions  $p$  and  $q$ , and a strictly convex functional  $\phi$  whose input space is a convex set of functions and whose output is in  $\mathbb{R}$ , the corresponding Bregman divergence is

$$D_\phi(p, q) = \phi(p) - \phi(q) - \int [p(x) - q(x)] \delta\phi(q)(x) dx.$$

Here  $\delta\phi(q)$  is the functional derivative of  $\phi$  at  $q$  and the integral term calculates this derivative in the direction of  $p - q$ .<sup>1</sup> An example of a functional Bregman divergence arises when we choose  $\phi(p) = \int p(x)^2 dx$ ; in this case, one can work out that the functional derivative of  $\phi$  at  $p$  is  $2p$  and that the resulting functional divergence is  $\int [p(x) - q(x)]^2 dx$ .

#### 3.2. The Symmetric Setting

Our first goal is to relate functional Bregman divergences back to concepts in metric learning and other related learning models. To do this, let us define a *symmetric* functional Bregman divergence as a functional Bregman divergence such that  $D_\phi(p, q) = D_\phi(q, p)$  for all  $p$  and  $q$ .

Our first result characterizes the form of an arbitrary symmetric functional Bregman divergence. This result can be stated as follows:

**Theorem 3.1.** *A functional Bregman divergence  $D_\phi(p, q)$  is a symmetric functional Bregman divergence if and only if*

<sup>1</sup>Note that Frigiyk et al. (2008) utilize the more general Fréchet derivative. Also, for simplicity, we limit ourselves to Riemann integrals unless otherwise noted. See Appendix for more details.

it has the following form:

$$D_\phi(p, q) = \iint (p(x) - q(x))(p(y) - q(y))\psi(x, y)dxdy,$$

where  $\psi(x, y)$  is some symmetric positive semi-definite function.

For instance, the example from above, where  $\phi(p) = \int p(x)^2 dx$ , can be seen as a special case where  $\psi(x, y) = 1$  if  $x = y$  and 0 otherwise. The proof of the theorem appears in the appendix. In essence, this result extends an analogous result known from the vector setting, which states that any symmetric Bregman divergence must be a Mahalanobis distance, namely  $D_\phi(\mathbf{x}, \mathbf{y}) = (\mathbf{x} - \mathbf{y})^T A(\mathbf{x} - \mathbf{y})$  for some positive semi-definite matrix  $A$  (Bauschke & Borwein, 2001).

Next we must show that, for particular choices of the symmetric positive semi-definite function  $\psi$ , as well as restrictions on  $p$  and  $q$ , the resulting divergence yields familiar forms.

**Deep Metric Learning and Moment-Matching.** Let us consider  $\psi(x, y) = f_W(x)^T f_W(y)$ , where  $f_W(x)$  is an embedding given by a neural network parameterized by weights  $W$ . This is clearly a positive semi-definite function, as it is an inner product between embedded data points. Further, assume  $p$  and  $q$  are distributions.

First, from Fubini’s theorem, observe that we can re-write the functional Bregman divergence in this special case as

$$D_\phi(p, q) = \|\mathbb{E}_p[f_W] - \mathbb{E}_q[f_W]\|^2.$$

This is a moment-matching type of metric. Note the similarity to the Wasserstein distance (Arjovsky et al., 2017) and the maximum mean discrepancy (Gretton, 2012). (In those cases, one further takes a supremum over the function  $f_W$ , which we would also do when performing optimization to learn  $f_W$ .)

This general form of the divergence is typically difficult to compute. We can consider the case when we have finite samples or, equivalently, we let  $p$  and  $q$  be given by empirical distributions over sets of points  $P$  and  $Q$ , respectively. In this case, the resulting divergence simplifies to

$$D_\phi(p, q) = \frac{1}{|P|} \sum_{x \in P} f_W(x) - \frac{1}{|Q|} \sum_{y \in Q} f_W(y) \quad ^2.$$

This yields a divergence measure between distributions  $p$  and  $q$  that matches the first moment, similar to how MMD operates.

To make connections to deep metric learning, consider the case where  $P$  and  $Q$  are of size one, namely Dirac delta

functions at points  $x$  and  $y$ , respectively. Then the divergence is simply

$$D_\phi(p, q) = \|f_W(x) - f_W(y)\|^2,$$

or just the squared Euclidean distance after embedding the data via a neural network. This form is precisely what nearly all deep metric learning methods employ: they learn a neural network to embed data, apply the (squared) Euclidean distance in the mapped space, and then apply a loss function such as a contrastive or triplet loss on top of this mapped distance (Chopra et al., 2005; Hoffer & Ailon, 2015).

**Linear Metric Learning.** If we replace the integral in the functional Bregman divergence with a Lebesgue integral (as it was defined in the original functional Bregman divergence paper), then use the counting measure for integration, the integral in the functional Bregman divergence simply becomes a sum over the elements in the measure space. In this case,  $\psi(x, y)$  is then replaced by a positive semi-definite matrix  $A$ , and function inputs to the divergence are replaced by vectors  $\mathbf{x}$  and  $\mathbf{y}$ . Then the resulting divergence is the usual Mahalanobis distance

$$D_\phi(\mathbf{x}, \mathbf{y}) = (\mathbf{x} - \mathbf{y})^T A(\mathbf{x} - \mathbf{y}).$$

Thus, we can recover the usual Mahalanobis metric used in linear metric learning under our framework.

**Kernel Metric Learning.** We can also recover familiar kernel forms of the preceding functions. In the case of a kernel function  $\psi(x, y) = \kappa(x, y)$ , the divergence recovers the moment-matching objective but with the norm induced by the kernel’s reproducing kernel Hilbert space, similar to kernel MMD (Gretton, 2012). Further, in the case of a kernel function  $\kappa(x, y) = g(x)^T A g(y)$ , where  $g(x)$  is an embedding to a reproducing kernel Hilbert space, and  $A$  is a positive-definite operator, the resulting divergence in the single-sample case yields the divergence studied for Mahalanobis metric learning in kernel space (Kulis et al., 2009).

A summary of some of the special cases described in this section appear in Table 3.2.

### 3.3. The General Setting

Next we consider the more general setting, i.e., when the functional divergence may not be symmetric. Here our goal is to introduce a parameterization of the functional divergences that are amenable to learning via neural networks. We term the resulting divergences as *deep Bregman divergences*.

A key insight of Siahkamari et al. (2019) was that one can approximate a strictly convex function arbitrarily well with a piecewise linear function. In particular, they chose to

## Deep Divergence Learning

Case	Integral Setting	$\psi(x, y)$	Inputs to $D_\phi$	$D_\phi$
Mahalanobis Distance	Lebesgue + Count. Meas.	$A \succeq 0$	Vectors $\mathbf{x}, \mathbf{y}$	$(\mathbf{x} - \mathbf{y})^T A (\mathbf{x} - \mathbf{y})$
Deep Metric Learning	Riemann	$f_W(\mathbf{x})^T f_W(\mathbf{y})$	Dirac Deltas at $\mathbf{x}, \mathbf{y}$	$\ f_W(\mathbf{x}) - f_W(\mathbf{y})\ ^2$
Moment Matching	Riemann	$f_W(\mathbf{x})^T f_W(\mathbf{y})$	Distributions $p, q$	$\ \mathbb{E}_p[f_W] - \mathbb{E}_q[f_W]\ ^2$

Table 1. Some of the special cases of  $D_\phi(p, q)$  for the symmetric divergence setting.

parameterize the generating function  $\phi$  of a vector Bregman divergence by the following max-affine function:

$$\phi(\mathbf{x}) = \max_c (\mathbf{x}^T \mathbf{w}_c + b_c).$$

Here  $c$  ranges from 1 to  $K$ , where  $K$  is the number of hyperplanes used to approximate the underlying strictly convex function. Such functions can be used to approximate any vector Bregman divergence arbitrarily well. Thus, learning a Bregman divergence amounts to learning the weights  $\mathbf{w}_i$  and biases  $b_i$  given appropriate supervision.

We can perform an analogous parameterization in the functional divergence setting. By generalizing the piecewise linear functions of Siahkamari et al, we can define a convex generating functional. The following theorem demonstrates that every convex generating functional can be expressed in terms of linear functionals, thus justifying our choice of parameterization:

**Theorem 3.2.** *Let  $\phi(p)$  be a convex generating functional corresponding to a functional Bregman divergence  $D_\phi$ . Then  $\phi$  can be formulated as*

$$\phi(p) = \sup_{(w, b_w) \in A} \int p(x)w(x)dx + b_w,$$

where  $A$  is a set of affine functionals in which each member is characterized by  $w$  and  $b_w$ .

For our parameterization, we replace supremum with maximum, and denote each function pair as  $(w_c, b_c)$  in a countable set of functionals  $A$ . See Appendix A.2 and B for the proof and details. In the case where  $p$  and  $q$  are distributions, we may write this more succinctly as  $\phi(p) = \max(\mathbb{E}_p[w_c] + b_c)$ , where the expectation is taken with respect to the subscript distribution  $p$ . Note that straightforward application of the calculus of variations reveals that the functional derivative of  $\phi(q)$  is simply  $w_{q^*}$ , where  $q^* = \operatorname{argmax}_c [\int q(x)w_c(x)dx + b_c]$ . Consequently, the functional Bregman divergence between  $p$  and  $q$  can be expressed as  $D_\phi(p, q) =$

$$\left( \int p(x)w_{p^*}(x)dx + b_{p^*} \right) - \left( \int p(x)w_{q^*}(x)dx + b_{q^*} \right). \quad (1)$$

For distributions, this is more succinctly  $D_\phi(p, q) = (\mathbb{E}_p[w_{p^*}] + b_{p^*}) - (\mathbb{E}_p[w_{q^*}] + b_{q^*})$ .

This parameterization of the functional  $\phi$  is now amenable to learning a functional divergence given data. In particular, we now parameterize a divergence by the corresponding weight functions  $w_1, \dots, w_K$  and biases  $b_1, \dots, b_K$ . If we assume that each of these weight functions are given by deep neural networks, then it becomes natural to set up learning problems where we aim to learn the underlying divergence given data. The resulting deep Bregman divergences will be shown to yield novel learning problems and strong empirical performance on benchmark metric learning tasks. In the next section we will detail our approach to extend deep metric learning to this setting.

## 4. Learning Problems and Applications

In the previous section, we saw in the symmetric setting how different choices of the functions related to a functional Bregman divergence yield existing forms, as well as how one may parameterize a general asymmetric functional Bregman divergence using deep neural networks. Now we connect the divergences discussed in the previous section to particular learning problems. In particular, we describe several novel applications and learning problems that arise from learning deep Bregman divergences.

### 4.1. From Deep Metric Learning to Deep Divergence Learning

Consider a learning problem where we aim to learn a deep divergence given supervised data. As with deep metric learning, we will consider the case when  $p$  and  $q$  are empirical distributions over single points  $x$  and  $y$ , respectively. We saw in the previous section that we will parameterize our deep divergence by weight functions  $w_1, \dots, w_K$  and biases  $b_1, \dots, b_K$ . To make things simpler, let us encompass all of these weight functions into a single large neural network with weights  $W$ . The network will have  $K$  different outputs, one for each weight function. Many possible architectures are possible to capture this type of network; we consider an architecture where several layers are shared in the network, and then the network branches into  $K$  subnetworks, each with its own independent set of weights. See Figure 2 for the network that we employ in our benchmark experiments.

Now, suppose we pass  $x$  through the network. Each subnetwork  $c$  produces a single output  $w_c(x) + b_c$ , and there are  $K$  total outputs, one per subnetwork. The index of the maximum output is  $p^*$ . Similarly, pass  $y$  through the network;




Figure 2. The general architecture we employ for deep Bregman divergences on image data. For  $K$  functionals, we produce  $K$  separate outputs, which are then used to compute the divergence over pairs of inputs.

the index of the maximum output across the  $K$  subnetworks is  $q^*$ . Then, by (1), the divergence is the difference between the output of  $x$  at  $p^*$  and the output of  $x$  at  $q^*$ . For instance, suppose that each of the  $K$  outputs corresponds to a different class. Then the divergence will be zero if both points achieve a maximum value for the same class (i.e., they are both classified into the same class). The divergence is non-zero if the points are assigned to different classes, and the divergence grows as the two outputs become more disparate.

One can now set up a divergence learning problem over pairs or triples of points under this framework. Suppose we are given triples of points  $(x, y, z)$ , where  $x$  should have a smaller divergence to  $y$  than to  $z$ . One can easily apply existing loss deep metric learning loss functions with the learned divergence in place of the usual squared Euclidean distance. The contrastive loss and triplet loss are the two most common ones with their definitions given below respectively:

$$\mathcal{L}_{cont}(x, y, z) = \frac{1}{2} [d(x, y) + \max\{m - d(x, z), 0\}]^2,$$

$$\mathcal{L}_{triplet}(x, y, z) = \max\{d(x, y) - d(z, x) + m, 0\},$$

where  $m$  is a margin value to impose a lower bound on the distance between the dissimilar pairs for the contrastive case and the difference between the relative distances for the triplet case. Typically, the distance measured in both loss functions,  $d$ , is the Euclidean distance; however, for our loss functions, we replace the distance measure by our learned deep Bregman divergence.

In experiments, we will compare existing deep metric learning approaches to the more general deep divergence learning problem considered here, and we will see that we obtain gains over the existing models on standard benchmarks.

## 4.2. Learning over Distributions

A key advantage to our framework is that we need not restrict ourselves only to divergences between single points.

As we saw earlier, we can also capture divergences between distributions of points that are similar to what is used for the MMD and the Wasserstein distance. Here we will discuss applications involving learning divergences over distributions.

**Data Generation.** Consider the problem encountered in many GAN applications: we aim to learn a generator for data such that we minimize some distributional divergence between the real and generated data distributions. In existing GAN literature, divergences considered include the Jensen-Shannon divergence (Goodfellow et al., 2014), MMD distance (Li et al., 2015; 2017), and the Wasserstein distance (Arjovsky et al., 2017).

Under the deep divergence framework, rather than employing a fixed divergence, we can learn one from data. In this setting, we consider two distributions  $p_{synth}$  and  $p_{real}$ , corresponding to distributions of generated and real data, respectively. Assume that  $p_{synth}$  is generated by passing randomly-generated input data through a generator  $g$ , as is standard with GAN models. As with GAN training, learning proceeds in an adversarial manner. We aim to learn a generator to minimize  $D_\phi(p_{synth}, p_{real})$ , while simultaneously we aim to learn weights of the underlying network parameterizing  $D_\phi$  to maximize  $D_\phi(p_{synth}, p_{real})$ . As with GANs, we alternate between gradient updates for these two objectives.

We note that, in practice, it is useful to restrict our attention to the case when  $K = 2$ , as it yields a particularly interpretable model. In this case, we can think of one of the two subnetworks as outputting a larger value on real data, while the other subnetwork as outputting a larger value on synthetic data. Thus, the network that parameterizes the divergence is analogous to the discriminator in a GAN model. When training the underlying weights of this network  $W$ , we can take pairs or triples of real and synthetic data and utilize a triplet or contrastive loss to encourage the output on the real data to be larger for one subnetwork and the output on the synthetic data to be larger for the other subnetwork. Similarly, when training the generator  $g$ , we use a loss that encourages real and synthetic data to both have the same maximal output.

## Semi-Supervised Distributional Clustering.

As another application of learning divergences over distributions, consider a scenario where instead of clustering a set of data points, we aim to cluster a set of distributions. In this setup, each distribution may correspond to an empirical distribution over a set of points—for instance, we may have a distribution of ratings for each item in an online store. The goal is: given a set of such distributions, to cluster the distributions together into a set of clusters.

Davis & Dhillon (2006) considered a version of this prob-




Figure 3. (Left) Plot of the means of the  $n = 500$  Gaussian distributions, color-coded by cluster identity. (Middle left) Plot of data after generating 50 points from each Gaussian. (Middle right) Embedding learned by our method using contrastive loss with a moment-matching function. (Right) Embedding learned by the baseline deep learning approach.

lem where each distribution was given by a multivariate Gaussian. Since the KL-divergence between multivariate Gaussians is itself a Bregman divergence, one can use properties of Bregman divergences to generalize the k-means algorithm to this setting. Here, we will consider a version of this problem that is both semi-supervised (so pairs of distributions that should or should not be clustered together are provided over a training set), and does not assume that each distribution is a multivariate Gaussian. Our approach also removes the implicit assumption that the means of the distributions are linearly separable for each cluster.

Analogous to Davis and Dhillon, given a functional Bregman divergence defined over distributions, one can apply a generalization of k-means to cluster the distributions. As shown by Frigyik et al. (2008), the mean minimizes the expected functional Bregman divergence over a set of distributions, analogous to the finite-dimensional case. Thus, k-means can be generalized to a setting where the squared Euclidean distance between vectors is replaced by the corresponding functional Bregman divergence over distributions.

If we represent each distribution by an empirical distribution over its underlying points, we can easily compute a parameterized functional Bregman divergence between pairs of distributions. In our experiments, we will consider in particular learning a symmetric divergence on supervised data using the moment-matching distance with a contrastive or triplet loss. Then, once we have learned the divergence from data, we replace the squared Euclidean distance in the k-means algorithm with the learned divergence to directly cluster data in the test set.

### 5. Experimental Results

We now empirically compare our proposed deep divergence framework to existing models. Due to space considerations, some further details and results are available in the supplementary material.

### 5.1. Clustering

**Clustering Multivariate Gaussian Distributions.** To begin, we consider a simple demonstration of the advantages of our approach on synthetic data for the semi-supervised distributional clustering problem. We generated  $n = 500$  training points, each assigned to one of three clusters. Each data point is represented by a multivariate Gaussian; the means of these Gaussians were uniformly sampled over rings of radius .2, .6, and 1 plus Gaussian noise, depending on the cluster identity, and the covariance of each Gaussian was .1 times the identity. See Figure 3 for a plot of sampled means, along with data after generating from these Gaussians. We also generated  $n = 200$  test points in the same manner.

We compare three approaches to cluster the data. Our first baseline is the method of Davis & Dhillon (2006), which is an unsupervised clustering algorithm designed specifically to cluster multivariate Gaussian distributions. The second baseline applies deep metric learning on all generated points from all the Gaussians; we apply contrastive and triplet losses separately and learn a 3-layer multilayer perceptron (MLP) over the data in each case. The number of units in each layer were set to 1000, 500, and 2, and standard ReLU activation was used. The third approach is our method; we apply the (empirical) moment-matching function from the symmetric setting, treating each distribution as its own data point, in conjunction with a contrastive and triplet losses to

Metrics	Baseline Method		Our Method		Davis & Dhillon	
	Triplet	Contrastive	Triplet	Contrastive		
RI	Mean	0.638	0.639	<b>0.997</b>	<b>0.999</b>	0.550
	Std	0.005	0.005	0.003	0.003	0.009
ARI	Mean	0.197	0.198	<b>0.993</b>	<b>0.997</b>	0.005
	Std	0.012	0.013	0.007	0.006	0.012

Table 2. Rand index and adjusted rand index scores for different clustering experiments, where the baseline method treats each training point independently.

Datasets			Baseline Method	Our Method	Davis & Dhillon
WHARF	RI	Mean	0.832	<b>0.887</b>	0.876
		Std	0.002	0.004	0.007
	ARI	Mean	0.098	<b>0.364</b>	0.327
		Std	0.006	0.022	0.026
MHEALTH	RI	Mean	0.849	<b>0.860</b>	0.664
		Std	0.005	0.007	0.006
	ARI	Mean	0.106	<b>0.149</b>	0.023
		Std	0.008	0.018	0.001
WISDM	RI	Mean	0.894	<b>0.907</b>	0.900
		Std	0.004	0.003	0.003
	ARI	Mean	0.086	<b>0.127</b>	0.089
		Std	0.005	0.014	0.009

Table 3. Rand index and adjusted rand index scores for different clustering experiments performed on real data, where the baseline method treats each training point independently.

learn a 3-layer MLP with the same settings as the baseline MLP. On the test set, we use the learned divergence in place of the squared Euclidean distance in a k-means algorithm for both the second and third method.

We compute the rand index and adjusted rand index scores on the test set in each case, averaged over 10 runs for each of the three methods. The results are given in Table 2. The Davis & Dhillon method cannot cluster the multivariate Gaussians, as their method is restricted to linear separability of the means. The baseline deep metric learning method fails due to the overlap of the generated data across clusters, whereas the distributional divergence approach is able to perfectly cluster the test data in most runs. We can also visualize the embeddings learned by the second and third method, where we see that our learned embeddings capture the correct cluster structure, as pictured in Figure 3.

**Human Activity Recognition Task.** We next perform further experiments on human activity recognition using time-varying sensor data, in order to demonstrate the capabilities of our distributional clustering method on real data. We use WHARF (Bruno et al., 2014), MHEALTH (Banos et al., 2014; 2015), and WISDM (Weiss et al., 2019) datasets in our initial experiments. These datasets are collections of multimodal body sensor recordings as test subjects perform different activities of daily living (ADL), including but not limited to sitting, standing, eating, walking, and jogging.

The experimental setup is the same as the previous task, where we compute the rand index and adjusted rand index scores on the test set in each experiment, averaged over 10 runs for each of the three methods. Results are given in Table 3. We note that only experiments using contrastive loss are reported here; though our distributional loss formula has the potential to be applied directly here, we leave this as future work.

We provide a visualization of the embeddings learned by

Datasets	Euclidean		Deep Bregman	
	Triplet	Contrastive	Triplet	Contrastive
MNIST	99.50	<b>99.63</b>	<b>99.61</b>	99.56
Fashion MNIST	93.24	93.57	<b>94.90</b>	<b>94.00</b>
SVHN	92.58	<b>94.88</b>	<b>94.03</b>	94.12
Cifar10	77.00	79.40	<b>81.40</b>	<b>80.80</b>
STL10	59.97	<b>63.10</b>	<b>62.64</b>	60.91

Table 4. K-nn classification accuracy results on the given datasets (without data augmentation or using learned features). The bold values indicate the best triplet loss (Bregman versus Euclidean) and contrastive loss (Bregman versus Euclidean) results.

our method and the baseline method, where we see that our learned embeddings capture the correct cluster structure, in the appendix.

### 5.2. Deep Metric Learning Comparisons

Next we consider comparisons between our general deep divergence learning framework and existing deep metric learning models on standard benchmarks, to demonstrate that our approach’s flexibility yields improved performance on several datasets and tasks.

We compare standard deep metric learning approaches to our proposed approach on the four benchmark datasets used in the original triplet loss paper (Hoffer & Ailon, 2015)—MNIST, Cifar10, SVHN, and STL10—as well as Fashion MNIST. We use the same basic architecture for the deep Bregman divergence network as shown in Figure 2; for the Euclidean case we do not employ separate subnetworks in the dense layers. We treat several architecture choices as hyperparameters and validate over these hyperparameters using Bayesian optimization (tuned separately for each dataset); Table 5 lists the hyperparameters that we search over, along with the ranges of values considered.

We consider separately both triplet loss and contrastive loss, and report in bold the best values for each loss. For the triplet loss, we consider all triplets in a batch when computing the loss. We perform no data augmentation. Results

Model hyperparams		Training hyperparams	
layers	2 - 5	margin	0.1 - 2.0
conv filters	16 - 128	epochs	10 - 40
conv kernels	3 - 9	learning rate	$10^{-5}$ - $10^{-1}$
conv biases	T / F	batch size	32-128
poolings	T / F	optimizer	adam / sgd / rms
batchnorms	T / F	K in k-nn	5 - 10
dense units	50 - 300	normalization	T / F

Table 5. Hyperparameter intervals used for tuning. First 100 iterations are used to narrow down the space, then 200 more iterations are run for each benchmark. T: True, F: False.






Figure 4. Real and generated sample batches from CelebA (Top row) and MNIST (Bottom row) datasets.

are shown in Table 4, where we see small but significant gains in classification accuracy for the Bregman method as compared to the standard deep metric learning approach, particularly in the triplet loss case. On Fashion MNIST, we outperform the current state-of-the-art for no data augmentation (94.23% from Assunção et al. (2018)), even though we are not directly training a classifier. We also note that we would expect further gains in performance with more sophisticated architectures (e.g., ResNets and other more recent architectures), perhaps yielding near state-of-the-art performance on more datasets; however, the main goal of this comparison is not to achieve state-of-the-art performance but rather to present a fair comparison between the Bregman and Euclidean approaches on standard benchmarks.

### 5.3. Unsupervised Data Generation

Finally, we consider some qualitative results where we show that our approach can be used for generating data with similar performance to GANs. We consider the problem discussed earlier, namely where we train a deep divergence model to minimize a learned divergence between real and synthetic data. We apply our approach on 28x28 MNIST and CELEBA datasets, as is standard for GAN applications. We adjust the strides to adapt the network for different input sizes. We keep model structures close to standard in order to show the effectiveness of the divergence formula we introduced. We use a generator consisting of 4 deconvolutional layers and a discriminator (i.e., the network parameterizing the deep Bregman divergence) with 4 convolutional layers, with a dropout rate of 0.5 in between the layers as well as lrelu and tanh activations. In the discriminator network, the convolutional layers are followed by two 2-layer subnetworks (again similar to Figure 2, where  $K = 2$  in this case). For the discriminator, we use the contrastive loss with a margin of 0.4, whereas the generator directly attempts to minimize deep Bregman divergence between the real and generated images. More hyperparameter details are given in the appendix.

Some randomly chosen results are presented in Figure 4, where we see that the distribution divergence learned by

our method is able to generate realistic-looking images with no labeled supervision. We note that further theoretical analysis and experimentation of these methods is required to determine situations where our loss functions may be more desirable than existing GAN approaches.

## 6. Conclusions

In this paper, we examined a novel generalization of both Bregman divergence learning and deep metric learning, which we call deep divergence learning. This framework offers several appealing advantages: it unifies a number of existing ideas in metric learning under a single framework, it suggests a way to extend deep metric learning beyond the Euclidean setting, and it naturally yields learning problems involving divergences over distributions. Empirically we have seen advantages of our approach compared to existing deep metric learning methods.

## 7. Acknowledgements

This research was supported by NSF CAREER Award 1559558.

## References

- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN, 2017. arXiv:1701.07875.
- Assunção, F., Lourenço, N., Machado, P., and Ribeiro, B. Denser: Deep evolutionary network structured representation. *arXiv preprint arXiv:1801.01563*, 2018.
- Banerjee, A., Merugu, S., Dhillon, I. S., and Ghosh, J. Clustering with Bregman divergences. *Journal of Machine Learning Research*, 6:1705–1749, 2005.
- Banos, O., Garcia, R., Holgado-Terriza, J. A., Damas, M., Pomares, H., Rojas, I., Saez, A., and Villalonga, C. mhealthdroid: a novel framework for agile development of mobile health applications. In *International workshop on ambient assisted living*, pp. 91–98. Springer, 2014.
- Banos, O., Villalonga, C., Garcia, R., Saez, A., Damas, M., Holgado-Terriza, J. A., Lee, S., Pomares, H., and Rojas, I. Design, implementation and validation of a novel open framework for agile development of mobile health applications. *Biomedical engineering online*, 14(2):S6, 2015.
- Bauschke, H. H. and Borwein, J. M. Joint and separate convexity of the Bregman distance. *Studies in Computational Mathematics*, 8:23–36, 2001.
- Bellet, A. and Habrard, A. Robustness and generalization for metric learning. *Neurocomputing*, 151:259–267, 2015.

- Bellet, A., Habrard, A., and Sebban, M. Metric learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 9(1):1–151, 2015.
- Bregman, L. M. The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3): 200–217, 1967.
- Bruno, B., Mastrogiovanni, F., and Sgorbissa, A. A public domain dataset for adl recognition using wrist-placed accelerometers. In *the 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pp. 738–743. IEEE, 2014.
- Cakir, F., He, K., Xide, X., Kulis, B., and Sclaroff, S. Deep metric learning to rank. In *Computer Vision and Pattern Recognition*, 2019.
- Cao, Q., Guo, Z.-C., and Ying, Y. Generalization bounds for metric and similarity learning. *Machine Learning*, 102(1):115–132, 2016.
- Chatpatanasiri, R., Korsrilabutr, T., Tangchanachaianan, P., and Kijssirikul, B. A new kernelization framework for Mahalanobis distance learning algorithms. *Neurocomputing*, 73(10–12):1570–1579, 2010.
- Chopra, S., Hadsell, R., and LeCun, Y. Learning a similarity metric discriminatively, with application to face verification. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005.
- Davis, J. and Dhillon, I. S. Differential entropic clustering of multivariate Gaussians. In *Advances in Neural Information Processing Systems (NIPS)*, 2006.
- Davis, J., Kulis, B., Jain, P., Sra, S., and Dhillon, I. Information-theoretic metric learning. In *Proc. 24th International Conference on Machine Learning (ICML)*, 2007.
- Frigyik, B. A., Srivastava, S., and Gupta, M. R. Functional Bregman divergences and Bayesian estimation of distributions. *IEEE Transactions on Information Theory*, 54(11):5130–5139, 2008.
- Globerson, A. and Roweis, S. Metric learning by collapsing classes. In *Advances in Neural Information Processing Systems (NIPS)*, 2005.
- Goldberger, J., Roweis, S., Hinton, G., and Salakhutdinov, R. Neighbourhood components analysis. In *Advances in Neural Information Processing Systems (NIPS)*, 2004.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- Gretton, A. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13:723–773, 2012.
- Hermans, A., Beyer, L., and Leibe, B. In Defense of the Triplet Loss for Person Re-Identification. *arXiv preprint arXiv:1703.07737*, 2017.
- Hoffer, E. and Ailon, N. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, pp. 84–92. Springer, 2015.
- Jain, P., Kulis, B., Dhillon, I., and Grauman, K. Online metric learning and fast similarity search. In *Advances in Neural Information Processing Systems (NIPS)*, 2008.
- Jain, P., Kulis, B., Davis, J., and Dhillon, I. Metric and kernel learning using a linear transformation. *Journal of Machine Learning Research*, 13:519–547, 2012.
- Kulis, B. Metric learning: A survey. *Foundations and Trends R in Machine Learning*, 5(4):287–364, 2013.
- Kulis, B., Sustik, M., and Dhillon, I. Low-rank kernel learning with Bregman matrix divergences. *Journal of Machine Learning Research*, 10:341–376, 2009.
- Li, C., Chang, W., Cheng, Y., Yang, Y., and Póczos, B. MMD-GAN: Towards deeper understanding of moment matching network. In *Neural Information Processing Systems*, 2017.
- Li, Y., Swersky, K., and Zemel, R. Generative moment matching networks. In *International Conference on Machine Learning*, 2015.
- Shalev-Shwartz, S., Singer, Y., and Ng, A. Y. Online and batch learning of pseudo-metrics. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 94. ACM, 2004.
- Siahkamari, A., Saligrama, V., Castanon, D., and Kulis, B. Learning Bregman divergences, 2019. arXiv:1905.11545.
- Wang, J., Zhou, F., Wen, S., Liu, X., and Lin, Y. Deep metric learning with angular loss. In *International Conference on Computer Vision*, 2017.
- Weinberger, K. Q. and Saul, L. K. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10:207–244, 2009.
- Weiss, G. M., Yoneda, K., and Hayajneh, T. Smartphone and smartwatch-based biometrics using activities of daily living. *IEEE Access*, 7:133190–133202, 2019.

Wu, L., Jin, R., Hoi, S. C., Zhu, J., and Yu, N. Learning Bregman distance functions and its application for semi-supervised clustering. In *Advances in neural information processing systems*, pp. 2089–2097, 2009.

Xing, E. P., Jordan, M. I., Russell, S. J., and Ng, A. Y. Distance metric learning with application to clustering with side-information. In *Advances in neural information processing systems*, pp. 521–528, 2003.