ON THE GEOMETRY OF DEEP BAYESIAN ACTIVE LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

We present geometric Bayesian active learning by disagreements (GBALD), a framework that performs BALD on its geometric interpretation interacting with a deep learning model. There are two main components in GBALD: initial acquisitions based on core-set construction and model uncertainty estimation with those initial acquisitions. Our **key innovation** is to construct the core-set on an ellipsoid, not typical sphere, preventing its updates towards the boundary regions of the distributions. **Main improvements** over BALD are twofold: relieving sensitivity to uninformative prior and reducing redundant information of model uncertainty. To guarantee the improvements, our generalization analysis proves that, compared to typical Bayesian spherical interpretation, geodesic search with ellipsoid can derive a tighter lower error bound and achieve higher probability to obtain a nearly zero error. Experiments on acquisitions with several scenarios demonstrate that, yielding slight perturbations to noisy and repeated samples, GBALD further achieves significant accuracy improvements than BALD, BatchBALD and other baselines.

1 INTRODUCTION

Lack of training labels restricts the performance of deep neural networks (DNNs), though prices of GPU resources were falling fast. Recently, leveraging the abundance of unlabeled data has become a potential solution to relieve this bottleneck whereby expert knowledge is involved to annotate those unlabeled data. In such setting, the deep learning community introduced **active learning (AL)** (Gal et al., 2017) that, maximizing the model uncertainty (Ashukha et al., 2019; Lakshminarayanan et al., 2017) to acquire a set of highly informative or representative unlabeled data, and solicit experts' annotations. During this AL process, the learning model tries to achieve a desired accuracy using minimal data labeling. Recent shift of model uncertainty in many fields, such as Bayesian neural networks (Blundell et al., 2015), Monte-Carlo (MC) dropout (Gal & Ghahramani, 2016), and Bayesian core-set construction (Sener & Savarese, 2018), shows that, new scenarios arise from deep Bayesian AL (Pinsler et al., 2019; Kirsch et al., 2019).

Bayesian AL (Golovin et al., 2010; Jedoui et al., 2019) presents an expressive probabilistic interpretation on model uncertainty (Gal & Ghahramani, 2016). Theoretically, for a simple regression model such as linear, logistic, and probit, AL can derive their closed-forms on updating one sparse subset that maximally reduces the uncertainty of the posteriors over the regression parameters (Pinsler et al., 2019). However, for a DNN model, optimizing massive training parameters is not easily tractable. It is thus that Bayesian approximation provides alternatives including importance sampling (Doucet et al., 2000) and Frank-Wolfe optimization (Vavasis, 1992). With importance sampling, a typical approach is to express the information gain in terms of the predictive entropy over the model, and it is called Bayesian active learning by disagreements (BALD) (Houlsby et al., 2011).

BALD has two interpretations: model uncertainty estimation and core-set construction. To estimate the model uncertainty, a greedy strategy is applied to select those data that maximize the parameter disagreements between the current training model and its subsequent updates as (Gal et al., 2017). However, naively interacting with BALD using uninformative prior leads to unstable biased acquisitions (Gao et al., 2020). Moreover, the similarity or consistency of those acquisitions to the previous acquired samples, brings redundant information to the model and decelerates its training.

Core-set construction (Campbell & Broderick, 2018) avoids the greedy interaction to the model by capturing characteristics of the data distributions. By modeling the complete data posterior over the distributions of parameters, BALD can be deemed as a core-set construction process on a sphere

(Kirsch et al., 2019), which seamlessly solicits a compact subset to approximate the input data distribution, and efficiently mitigates the sensitivity to uninformative prior and redundant information.

From the view of geometry, updates of core-set construction is usually optimized with sphere geodesic as (Nie et al., 2013; Wang et al., 2019). Once the core-set is obtained, deep AL immediately seeks annotations from experts and starts the training. However, data points located at the boundary regions of the distribution, which are no longer distributed uniformly, cannot be highly-representative candidates for the core-set. Therefore, constructing the core-set on a sphere may not be the optimal choice for deep AL.

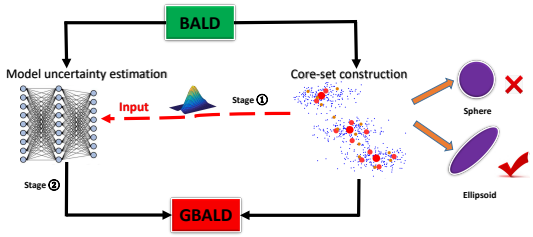

Figure 1: Illustration of two-stage GBALD framework. BALD has two types of interpretation: model uncertainty estimation and core-set construction where the deeper the color of the core-set element, the higher the representation; GBALD integrates them into a uniform framework. Stage ①: core-set construction is with an ellipsoid, not typical sphere, representing the original distribution to initialize the input features of DNN. Stage ②: model uncertainty estimation with those initial acquisitions then derives highly informative and representative samples for DNN.

This paper presents a novel AL framework, namely **Geometric BALD (GBALD)**, over the geometric interpretation of BALD that, interpreting BALD with core-set construction on an ellipsoid, initializes an effective representation to drive a DNN model. The goal is to seek for significant accuracy improvements against an uninformative prior and redundant information. Figure 1 describes this two-stage framework. In the first stage, geometric core-set construction on an ellipsoid initializes effective acquisitions to start a DNN model regardless of the uninformative prior. Taking the core-set as the input features, the next stage ranks the batch acquisitions of model uncertainty according to their geometric representativeness, and then solicits some highly-representative examples from the batch. With the representation constraints, the ranked acquisitions reduce the probability of sampling nearby samples of the previous acquisitions. It is thus the ranking rejects unnecessary redundant acquisitions. To guarantee the improvement, our generalization analysis shows that, the lower bound of generalization errors of AL with the ellipsoid is proven to be tighter than that of AL with the sphere. Achieving a nearly zero generalization error by AL with ellipsoid is also proven to have higher probability. Contributions of this paper are listed as follows.

- GBALD, a geometric Bayesian AL framework is presented to initialize the geometric interpretation of BALD on an ellipsoid, deriving the deep learning model into highly informative and representative acquisitions.
- To interpret the core-set construction on an ellipsoid, a set of geometric generalization bounds is presented, including error bounds of geodesic search with ellipsoid and sphere, and their probability bounds of achieving a nearly zero error. See Appendix B.

The rest of this paper is organized as follows. In Section 2, we first review the related work. Secondly, we elaborate BALD and GBALD in Sections 3 and 4, respectively. Experimental results are presented in Section 5. Finally, we conclude this paper in Section 6.

2 Related Work

Model uncertainty. In deep learning community, AL (Cohn et al., 1994) was introduced to improve the training of a DNN model by annotating unlabeled data, where the data which maximize the model uncertainty (Lakshminarayanan et al., 2017) are the primary acquisitions. For example, in ensemble deep learning (Ashukha et al., 2019), out-of-domain uncertainty estimation selects those data which do not follow the same distribution as the input training data; in-domain uncertainty draws the data from the original input distribution, producing reliable probability estimates. Gal & Ghahramani (2016) use MC dropout to estimate predictive uncertainty for approximating a Bayesian convolutional neural network. Lakshminarayanan et al. (2017) estimate predictive uncertainty using a proper scoring rule as the training criteria to fed a DNN.

Bayesian AL. Taking a Bayesian perspective (Golovin et al., 2010), AL can be deemed as minimizing the Bayesian posterior risk with multiple label acquisitions over the input unlabeled data. A potential informative approach is to reduce the uncertainty about the parameters using Shannon's entropy (Tang et al., 2002). This can be interpreted as seeking the acquisitions for which the Bayesian parameters

under the posterior disagree about the outcome the most, so this acquisition algorithm is referred to as Bayesian active learning by disagreement (BALD) (Houlsby et al., 2011).

Deep AL. Recently, deep Bayesian AL attracted the community. Gal et al. (2017) proposed to cooperate BALD with a DNN to improve the training. The unlabeled data which maximize the model uncertainty provide positive feedback. However, it needs to repeatedly update the model until the acquisition budget is exhausted. To improve the acquisition efficiency, batch sampling with BALD is applied as (Kirsch et al., 2019; Pinsler et al., 2019). In BatchBALD, Kirsch et al. (2019) developed a tractable approximation to the mutual information of one batch of unlabeled data and current model parameters. However, those uncertainty evaluations of Bayesian AL whether in single or batch acquisitions all take greedy strategies, which lead to computationally infeasible, or excursive parameter estimations. For deep Bayesian AL, being short of interactions to DNN can not maximally drive their model performance as (Pinsler et al., 2019; Sener & Savarese, 2018), etc.

3 BALD

BALD has two different interpretations: model uncertainty estimation and core-set construction. We simply introduce them in this section.

3.1 Model uncertainty estimation

We consider a discriminative model $p(y|x,\theta)$ parameterized by θ that maps $x \in \mathcal{X}$ into an output distribution over a set of $y \in \mathcal{Y}$. Given an initial labeled (training) set $\mathcal{D}_0 \in \mathcal{X} \times \mathcal{Y}$, the Bayesian inference over this parameterized model is to estimate the posterior $p(\theta|\mathcal{D}_0)$, i.e. estimate θ by repeatedly updating $\mathcal{D}_0$. AL adopts this setting from a Bayesian view.

With AL, the learner can choose unlabeled data from $\mathcal{D}_u = \{x_i\}_{j=1}^N \in \mathcal{X}$, to observe the outputs of the current model, maximizing the uncertainty of the model parameters. Houlsby et al. (2011) proposed a greedy strategy termed BALD to update $\mathcal{D}_0$ by estimating a desired data x^* that maximizes the decrease in expected posterior entropy:

$$x^* = \arg\max_{x \in \mathcal{D}_u} \mathrm{H}[\theta|\mathcal{D}_0] - \mathbb{E}_{y \sim p(y|x,\mathcal{D}_0)}\Big[\mathrm{H}[\theta|x,y,\mathcal{D}_0]\Big], \tag{1}$$

where the labeled and unlabeled sets are updated by $\mathcal{D}_0 = \mathcal{D}_0 \cup \{x^*, y^*\}, \mathcal{D}_u = \mathcal{D}_u \backslash x^*$, and y^* denotes the output of x^*. In deep AL, y^* can be annotated as a label from experts and θ yields a DNN model.

3.2 Core-set construction

Let $p(\theta|\mathcal{D}_0)$ be updated by its log posterior $\log p(\theta|\mathcal{D}_0, x^*)$, $y^* \in \{y_i\}_{i=1}^N$, assume the outputs are conditional independent of the inputs, i.e. $p(y^*|x^*, D_0) = \int_\theta p(y^*|x^*,\theta)p(\theta|D_0)\mathrm{d}\theta$, then we have the *complete data log posterior following* (Pinsler et al., 2019):

$$\begin{aligned}\mathbb{E}_{y^*}[\log p(\theta|\mathcal{D}_0, x^*, y^*)] &= \mathbb{E}_{y^*}[\log p(\theta|\mathcal{D}_0) + \log p(y^*|x^*,\theta) - \log p(y^*|x^*,\mathcal{D}_0)] \\ &= \log p(\theta|\mathcal{D}_0) + \mathbb{E}_{y^*}\Big[\log p(y^*|x^*,\theta) + \mathrm{H}[y^*|x^*,\mathcal{D}_0]\Big] \\ &= \log p(\theta|\mathcal{D}_0) + \sum_{i=1}^N \left(\mathbb{E}_{y_i}\Big[\log p(y_i|x_i,\theta) + \mathrm{H}[y_i|x_i,\mathcal{D}_0]\Big]\right).\end{aligned} \tag{2}$$

The key idea of core-set construction is to approximate the log posterior of Eq. (2) by a subset of $D'_u \subseteq D_u$ such that: $\mathbb{E}_{\mathcal{Y}_u}[\log p(\theta|\mathcal{D}_0, \mathcal{D}_u, \mathcal{Y}_u)] \approx \mathbb{E}_{\mathcal{Y}'_u}[\log p(\theta|\mathcal{D}_0, \mathcal{D}'_u, \mathcal{Y}'_u)]$, where $\mathcal{Y}_u$ and $\mathcal{Y}'_u$ denote the predictive labels of $\mathcal{D}_u$ and $\mathcal{D}'_u$ respectively by the Bayesian discriminative model, that is, $p(\mathcal{Y}_u|\mathcal{D}_u, D_0) = \int_\theta p(\mathcal{Y}_u|\mathcal{D}_u,\theta)p(\theta|D_0)\mathrm{d}\theta$, and $p(\mathcal{Y}'_u|\mathcal{D}'_u, D_0) = \int_\theta p(\mathcal{Y}'_u|\mathcal{D}'_u,\theta)p(\theta|D_0)\mathrm{d}\theta$. Here D'_u can be indicated by a core-set (Pinsler et al., 2019) that highly represents $\mathcal{D}_u$. Optimization tricks such as Frank-Wolfe optimization (Vavasis, 1992) then can be adopted to solve this problem.

Motivations. Eqs. (1) and (2) provide the Bayesian rules of BALD over model uncertainty and core-set construction respectively, which further attract the attention of the deep learning community. However, the two interpretations of BALD are limited by: 1) redundant information and 2) uninformative prior, where one major reason which causes these two issues is the poor initialization on the prior, i.e. $p(\mathcal{D}_0|\theta)$. For example, unbalanced label initialization on $\mathcal{D}_0$ usually leads to an uninformative prior, which further conducts the acquisitions of AL to select those unlabeled data from one or some fixed classes; highly-biased results with (Gao et al., 2020) redundant information are inevitable. Therefore, these two limitations affect each other.

4 GBALD

GBALD consists of two components: initial acquisitions based on core-set construction and model uncertainty estimation with those initial acquisitions.

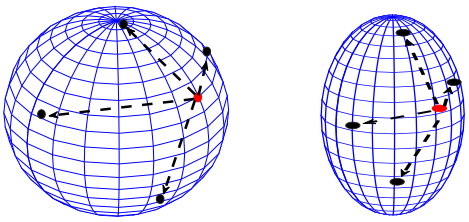

(a) Sphere geodesic (b) Ellipsoid geodesic

Figure 2: Optimizing BALD with sphere and ellipsoid geodesics. Ellipsoid geodesic rescales the sphere geodesic to prevent the updates of the core-set towards the boundary regions of the sphere where the characteristics of the distribution cannot be captured. Black points denote the feasible updates of the red points. Dash lines denote the geodesics.

4.1 Geometric interpretation of core-set

Modeling the complete data posterior over the parameter distribution can relieve the above two limitations of BALD. Typically, finding the acquisitions of AL is equivalent to approximating a core-set centered with spherical embeddings (Sener & Savarese, 2018). Let w_i be the sampling weight of x_i, $\|w_i\|_0 \leq N$, the core-set construction is to optimize:

$$\min_{w} \left\| \underbrace{\sum_{i=1}^{N} \mathbb{E}_{y_i}\Big[\log p(y_i|x_i,\theta) + \mathrm{H}[y_i|x_i,\mathcal{D}_0]\Big]}_{\mathcal{L}} - \underbrace{\sum_{i=1}^{N} w_i \mathbb{E}_{y_i}\Big[\log p(y_i|x_i,\theta) + \mathrm{H}[y_i|x_i,\mathcal{D}_0]\Big]}_{\mathcal{L}(w)} \right\|^2, \tag{3}$$

where $\mathcal{L}$ and $\mathcal{L}(w)$ denote the full and expected (weighted) log-likelihoods, respectively (Campbell & Broderick, 2018; 2019). Specifically, $\sum_{i=1}^{N} \mathrm{H}[y_i|x_i,\mathcal{D}_0] = -\sum_{y_i} p(y_i|x_i,\mathcal{D}_0)\log(p(y_i|x_i,\mathcal{D}_0)$, where $p(y_i|x_i,\mathcal{D}_0) = \int_\theta p(y_i|x_i,\theta)p(\theta|\mathcal{D}_0)\mathrm{d}\theta$. Note $\|\cdot\|$ denotes the ℓ^2 norm.

The approximation of Eq. (3) implicitly requires that the complete data log posterior of Eq. (2) w.r.t. $\mathcal{L}$ must be close to an expected posterior w.r.t. $\mathcal{L}(w)$ such that approximating a sparse subset for the original inputs by sphere geodesic search is feasible (see Figure 2(a)). Generally, solving this optimization is intractable due to cardinality constraint (Pinsler et al., 2019). Campbell & Broderick (2019) proposed to relax the constraint in Frank–Wolfe optimization, in which mapping $\mathcal{X}$ is usually performed in a Hilbert space (HS) with a bounded inner product operation. In this solution, the sphere embedded in the HS replaces the cardinality constraint with a polynomial constraint. However, the initialization on $\mathcal{D}_0$ affects the iterative approximation to $\mathcal{D}_u$ at the beginning of the geodesic search. Moreover, the posterior of $p(\theta|\mathcal{D}_0)$ is uninformative, if the initialized $\mathcal{D}_0$ is empty or not correct. Therefore, the typical Bayesian core-set construction of BALD cannot ideally fit an uninformative prior. The another geometric interpretation of core-set construction, such as k-centers (Sener & Savarese, 2018), is not restricted to this setting. We thus follow the construction of k-centers to find the core-set.

***k*-centers.** Sener & Savarese (2018) proposed a core-set representation approach for active deep learning based on k-centers. This approach can be adopted in core-set construction of BALD without the help of the discriminative model. Therefore, the uninformative prior has no further influence on the core-set. Typically, the k-centers approach uses a greedy strategy to search the data $\widetilde{x}$ whose nearest distance to elements of $\mathcal{D}_0$ is the maximal:

$$\widetilde{x} = \arg\max_{x_i \in \mathcal{D}_u} \min_{c_i \in \mathcal{D}_0} \|x_i - c_i\|, \tag{4}$$

then $\mathcal{D}_0$ is updated by $\mathcal{D}_0 \cup \{\widetilde{x}, \widetilde{y}\}$, $\mathcal{D}_u$ is updated by $\mathcal{D}_u \backslash \widetilde{x}$, where $\widetilde{y}$ denotes the output of $\widetilde{x}$. This max-min operation usually performs k times to construct the centers.

From the view of geometry, k-centers can be deemed as the core-set construction via spherical geodesic search (Bādoiu et al., 2002; Har-Peled & Mazumdar, 2004). Specifically, the max-min optimization guides $\mathcal{D}_0$ to be updated into one data, which draws the longest line segment from $x_i, \forall i$ across the sphere center. The iterative update on $\widetilde{x}$ is then along its unique diameter through the sphere center. However, this greedy optimization has large probability that yields the core-set to fall into boundary regions of the sphere, which cannot capture the characteristics of the distribution.

4.2 Initial acquisitions based on core-set construction

We present a novel greedy search which rescales the geodesic of a sphere into an ellipsoid following Eq. (4), in which the iterative update on the geodesic search is rescaled (see Figure 2(b)). We follow the importance sampling strategy to begin the search.

Initial prior on geometry. Initializing $p(\mathcal{D}_0|\theta)$ is performed with a group of internal spheres centered with $D_j, \forall j$, subjected to $D_j \in \mathcal{D}_0$, in which the geodesic between $\mathcal{D}_0$ and the unlabeled data is over those spheres. Since $\mathcal{D}_0$ is known, specification of θ plays the key role for initializing $p(\mathcal{D}_0|\theta)$. Given a radius R_0 for any observed internal sphere, $p(y_i|x_i,\theta)$ is firstly defined by

$$p(y_i|x_i,\theta) = \begin{cases} 1, & \exists j, \|x_i - D_j\| \leq R_0, \\ \max\left\{\dfrac{R_0}{\|x_i - D_j\|}\right\}, & \forall j, \|x_i - D_i\| > R_0, \end{cases} \tag{5}$$

thereby θ yields the parameter R_0. When the data is enclosed with a ball, the probability of Eq. (5) is 1. The data near the ball, is given a probability of $\max\left\{\frac{R_0}{\|x_i - D_j\|}\right\}$ constrained by $\min\|x_i - D_j\|, \forall j$, i.e. the probability is assigned by the nearest ball to x_i, which is centered with D_j. From Eq. (3), the information entropy of $y_i \sim \{y_1, y_2, ..., y_N\}$ over $x_i \sim \{x_1, x_2, ..., x_N\}$ can be expressed as the integral regarding $p(y_i|x_i,\theta)$:

$$\sum_{i=1}^{N} \mathrm{H}(y_i|x_i,\mathcal{D}_0) = -\sum_{i=1}^{N} \int_\theta p(y_i|x_i,\theta)p(\theta|D_0)d\theta \log\Big(\int_\theta p(y_i|x_i,\theta)p(\theta|D_0)\Big)d\theta, \tag{6}$$

which can be approximated by $-\sum_{i=1}^{N} p(y_i|x_i,\theta)\log\Big(p(y_i|x_i,\theta)\Big)$ following the details of Eq. (3). In short, this indicates an approximation to the entropy over the entire outputs on $\mathcal{D}_u$ that assumes the prior $p(D_0|\theta)$ w.r.t. $p(y_i|x_i,\theta)$ is already known from Eq. (5).

Max-min optimization. Recalling the max-min optimization trick of k-centers in the core-set construction of (Sener & Savarese, 2018), the minimizer of Eq. (3) can be divided into two parts: $\min_{x^*} \mathcal{L}$ and $\max_w \mathcal{L}(w)$, where $\mathcal{D}_0$ is updated by acquiring x^*. However, updates of $\mathcal{D}_0$ decide the minimizer of $\mathcal{L}$ with regard to the internal spheres centered with $D_i, \forall i$. Therefore, minimizing $\mathcal{L}$ should be constrained by an unbiased full likelihood over $\mathcal{X}$ to alleviate the potential biases from the initialization of $\mathcal{D}_0$. Let $\mathcal{L}_0$ denote the unbiased full likelihood over $\mathcal{X}$ that particularly stipulates $\mathcal{D}_0$ as the k-means centers written as $\mathcal{U}$ of $\mathcal{X}$ which jointly draw the input distribution. We define $\mathcal{L}_0 = |\sum_{i=1}^{N} \mathbb{E}_{y_i}[\log p(y_i|x_i,\theta) + \mathrm{H}[y_i|x_i,\mathcal{U}]]|$ to regulate $\mathcal{L}$, that is

$$\min_{x^*} \|\mathcal{L}_0 - \mathcal{L}\|^2, \ \text{ s.t. } \mathcal{D}_0 = \mathcal{D}_0 \cup \{x^*, y^*\}, \mathcal{D}_u = \mathcal{D}_u \backslash x^*. \tag{7}$$

The other sub optimizer is $\max_w \mathcal{L}(w)$. We present a greedy strategy following Eq. (1):

$$\begin{aligned} &\max_{1\leq i\leq N} \min_{w_i} \sum_{i=1}^{N} w_i \mathbb{E}_{y_i}[\log p(y_i|x_i,\theta) + \mathrm{H}[y_i|x_i,\mathcal{D}_0]] \\ &= \sum_{i=1}^{N} w_i \log p(y_i|x_i,\theta) - \sum_{i=1}^{N} w_i p(y_i|x_i,\theta)\log p(y_i|x_i,\theta), \end{aligned} \tag{8}$$

which can be further written as: $\sum_{i=1}^{N} w_i \log p(y_i|x_i,\theta)(1 - \log p(y_i|x_i,\theta))$. Let $w_i = 1, \forall i$ for unbiased estimation of the likelihood $\mathcal{L}(w)$, Eq. (8) can be simplified as

$$\max_{x_i \in \mathcal{D}_u} \min_{D_j \in \mathcal{D}_0} \log p(y_i|x_i,\theta), \tag{9}$$

where $p(y_i|x_i,\theta)$ follows Eq. (5). Combining Eqs. (7) and (9), the optimization of Eq. (3) is then transformed as

$$x^* = \arg\max_{x_i \in \mathcal{D}_u} \min_{D_i \in \mathcal{D}_0} \left\{\|\mathcal{L}_0 - \mathcal{L}\|^2 + \log p(y_i|x_i,\theta)\right\}, \tag{10}$$

where $\mathcal{D}_0$ is updated by acquiring x^*, i.e. $\mathcal{D}_0 = \mathcal{D}_0 \cup \{x^*, y^*\}$.

Geodesic line. For a metric geometry M, a geodesic line is a curve γ which projects its interval I to M: $I \to M$, maintaining everywhere locally a distance minimizer (Lou et al., 2020). Given a constant $\nu > 0$ such that for any $a, b \in I$ there exists a geodesic distance $d(\gamma(a),\gamma(b)) := \int_a^b \sqrt{g_{\gamma(t)}(\gamma'(t),\gamma'(t))}dt$, where $\gamma'(t)$ denotes the geodesic curvature, and g denotes the metric tensor over M. Here, we define $\gamma'(t) = 0$, then $g_{\gamma(t)}(0,0) = 1$ such that $d(\gamma(a),\gamma(b))$ can be generalized as a segment of a straight line: $d(\gamma(a),\gamma(b)) = \|a - b\|$.

Ellipsoid geodesic distance. For any observation points $p, q \in M$, if the spherical geodesic distance is defined as $d(\gamma(p),\gamma(q)) = \|p - q\|$. The affine projection obtains its ellipsoid interpretation: $d(\gamma(p),\gamma(q)) = \|\eta(p - q)\|$, where η denotes the affine factor subjected to $0 < \eta < 1$.

Optimizing with ellipsoid geodesic search. The max-min optimization of Eq. (10) is performed on an ellipsoid geometry to prevent the updates of core-set towards the boundary regions, where ellipsoid

geodesic line scales the original update on the sphere. Assume x_i is the previous acquisition and x^* is the next desired acquisition, the ellipsoid geodesic rescales the position of x^* as $x_e^* = x_i + \eta(x^* - x_i)$. Then, we update this position of x_e^* to its nearest neighbor x_j in the unlabeled data pool, i.e. $\arg\min_{x_j \in \mathcal{D}_u} \|x_j - x_e^*\|$, also can be written as

$$\arg\min_{x_j \in \mathcal{D}_u} \Big\| x_j - [x_i + \eta(x^* - x_i)] \Big\|. \tag{11}$$

To study the advantage of ellipsoid geodesic search, Appendix B presents our generalization analysis.

4.3 Model uncertainty estimation with core-set

GBALD starts the model uncertainty estimation with those initial core-set acquisitions, in which it introduces a ranking scheme to derive both informative and representative acquisitions.

Single acquisition. We follow (Gal et al., 2017) and use MC dropout to perform Bayesian inference on the model of the neural network. It then leads to ranking the informative acquisitions with batch sequences is with high efficiency. We first present the ranking criterion by rewriting Eq. (1) as batch returns:

$$\{x_1^*, x_2^*, ..., x_b^*\} = \underset{\{\hat{x}_1, \hat{x}_2, ..., \hat{x}_b\} \subseteq \mathcal{D}_u}{\arg\max} \; \mathrm{H}[\theta|\mathcal{D}_0] - \mathbb{E}_{\hat{y}_{1:b} \sim p(\hat{y}_{1:b}|\hat{x}_{1:b}, \mathcal{D}_0)} \Big[\mathrm{H}[\theta|\hat{x}_{1:b}, \hat{y}_{1:b}, \mathcal{D}_0] \Big], \tag{12}$$

where $\hat{x}_{1:b} = \{\hat{x}_1, \hat{x}_2, ..., \hat{x}_b\}$, $\hat{y}_{1:b} = \{\hat{y}_1, \hat{y}_2, ..., \hat{y}_b\}$, $\hat{y}_i$ denotes the output of $\hat{x}_i$. The informative acquisition x_t^* is then selected from the ranked batch acquisitions $\hat{x}_{1:b}$ due to the highest representation for the unlabeled data:

$$x_t^* = \underset{x_i^* \in \{x_1^*, x_2^*, ..., x_b^*\}}{\arg\max} \left\{ \max_{D_j \in \mathcal{D}_0} \; p(y_i|x_i^*, \theta) := \frac{R_0}{\|x_i^* - D_j\|} \right\}, \tag{13}$$

where t denotes the index of the final acquisition, subjected to $1 \leq t \leq b$. This also adopts the max-min optimization of k-centers in Eq. (4), i.e. $x_t^* = \arg\max_{x_i^* \in \{x_1^*, x_2^*, ..., x_b^*\}} \min_{D_j \in \mathcal{D}_0} \; \|x_i^* - D_j\|$.

Batch acquisitions. The greedy strategy of Eq. (13) can be written as a batch acquisitions by setting its output as a batch set, i.e.

$$\{x_{t_1}^*, ..., x_{t_{b'}}^*\} = \underset{x_{t_1:t_{b'}}^* \subseteq \{x_1^*, x_2^*, ..., x_b^*\}}{\arg\max} \; p(y_{t_1:t_{b'}}^* | x_{t_1:t_{b'}}^*, \theta), \tag{14}$$

where $x_{t_1:t_{b'}}^* = \{x_{t_1}^*, ..., x_{t_{b'}}^*\}$, $y_{t_1:t_{b'}}^* = \{y_{t_1}^*, ..., y_{t_{b'}}^*\}$, $y_{t_i}^*$ denotes the output of $x_{t_i}^*$, $1 \leq i \leq b'$, and $1 \leq b' \leq b$. This setting can be used to accelerate the acquisitions of AL in a large dataset. Appendix A presents the two-stage GBALD algorithm.

5 Experiments

In experiments, we start by showing how BALD degenerates its performance with uninformative prior and redundant information, and show that how our proposed GBALD relieves theses limitations.

Our experiments discuss three questions: 1) is GBALD using core-set of Eq. (11) competitive with uninformative prior? 2) can GBALD using ranking of Eq. (14) improve informative acquisitions of model uncertainty? and 3) can GBALD outperform state-of-the-art acquisition approaches? Following the experiment settings of (Gal et al., 2017; Kirsch et al., 2019), we use MC dropout to implement the Bayesian approximation of DNNs. Three benchmark datasets are selected: MNIST, SVHN and CIFAR10. More experiments are presented in Appendix C.

5.1 Uninformative priors

As discussed in the introduction, BALD is sensitive to an uninformative prior, i.e. $p(\mathcal{D}_0|\theta)$. We thus initialize $\mathcal{D}_0$ from a fixed class of the tested dataset to observe its acquisition performance. Figure 3 presents the prediction accuracies of BALD with an acquisition budget of 130 over the training set of MNIST, in which we randomly select 20 samples from the digit '0' and '1' to initialize $\mathcal{D}_0$, respectively. The classification model of AL follows a convolutional neural network with one block of

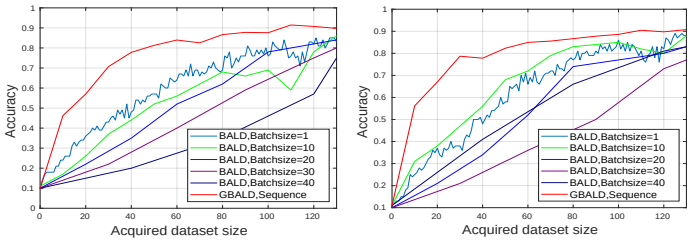

(a) Digit '0' (b) Digit '1'

Figure 3: Acquisitions with uninformative priors from digit '0' and '1'.

[convolution, dropout, max-pooling, relu], with 32, 3x3 convolution filters, 5x5 max pooling, and 0.5 dropout rate. In the AL loops, we use 2,000 MC dropout samples from the unlabeled data pool to fit the training of the network following (Kirsch et al., 2019).

The results show BALD can slowly accelerate the training model due to biased initial acquisitions, which cannot uniformly cover all the label categories. Moreover, the uninformative prior guides BALD to unstable acquisition results. As the shown in Figure 3(b), BALD with Bathsize = 10 shows better performance than that of Batchsize =1; while BALD in Figure 3(a) keeps stable performance. This is because the initial labeled data does not cover all classes and BALD with Batchsize =1 may further be misled to select those samples from one or a few fixed classes at the first acquisitions. However, Batchsize >1 may result in a random acquisition process that possibly covers more diverse labels at its first acquisitions. Another excursive result of BALD is that the increasing batch size cannot degenerate its acquisition performance in Figure 3(b). Specifically, Batchsize =10 $\succ$ Batchsize =1 $\succ$ Batchsize =20,40 $\succ$ Batchsize =30, where '>' denotes 'better' performance; Batchsize = 20 achieves similar results of Batchsize =40. **This undermines the acquisition policy of BALD: its performance would degenerate when the batch size increases.**

Different to BALD, the core-set construction of GBALD using Eq. (11) provides a **complete label matching against all classes**. Therefore, it outperforms BALD with the batch sizes of 1, 10, 20, 30, and 40. As the shown learning curves in Figure 3, GBALD with a batch size of 1 and sequence size of 10 (i.e. breakpoints of acquired size are 10, 20, ..., 130) achieves significantly higher accuracies than BALD using different batch sizes since BALD misguides the network updating using poor prior.

5.2 Improved Informative Acquisitions

Repeated or similar acquisitions delay the acceleration of the model training of BALD. Following the experiment settings of Section 5.1, we compare the best performance of BALD with a batch size of 1 and GBALD with different batch size parameters. Following Eq. (14), we set $b = \{3, 5, 7\}$ and b'=1, respectively, that means, we output the highest representative data from a batch of highly-informative acquisitions. Different settings on b and b' are used to observe the parameter perturbations of GBALD.

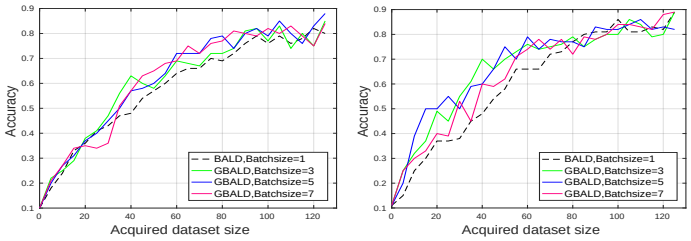

(a) Digit '0' (b) Digit '1'

Figure 4: GBALD outperforms BALD using ranked informative acquisitions which cooperate with representation constraints.

Training by the same parameterized CNN model as Section 5.1, Figure 4 presents the acquisition performance of parameterized BALD and GBALD. As the learning curves shown, BALD cannot accelerate the model as fast as GBALD due to the repeated information over the acquisitions. For GBALD, it ranks the batch acquisitions of the highly-informative samples and selects the highest representative ones. By employing this special ranking strategy, GBALD can **reduce the probability of sampling nearby data of the previous acquisitions**. It is thus GBALD significantly outperforms BALD, even if we progressively increase the ranked batch size b.

5.3 Active Acquisitions

GBALD using Eqs. (11) and (14) has been demonstrated to achieve successful improvements over BALD. We thus combine these two components into a uniform framework. Figure 5 reports the AL accuracies using different acquisition algorithms on the three image datasets. The selected baselines follow (Gal et al., 2017) including 1) maximizing the variation ratios (Var), 2) BALD, 3) maximizing the entropy (Entropy), 4) k-medoids, and one greedy 5) k-centers approach (Sener & Savarese, 2018). The network architecture is a three-layer MLP with three blocks of [convolution, dropout, max-pooling, relu], with 32, 64, and 128 3x3 convolution filters, 5x5 max pooling, and 0.5 dropout rate. In the AL loops, the MC dropout still randomly samples 2,000 data from the unlabeled data pool to approximate the training of the network architecture following (Kirsch et al., 2019). The initial labeled data of MNIST, SVHN and CIFAR-10 are 20, 1000, 1000 random samples from their full training sets. Details of baselines are presented in Appendix C.1.

The batch size of the compared baselines is 100, where GBALD ranks 300 acquisitions to select 100 data for the training, i.e. $b = 300, b' = 100$. As the learning curves shown in Figure 5, 1)

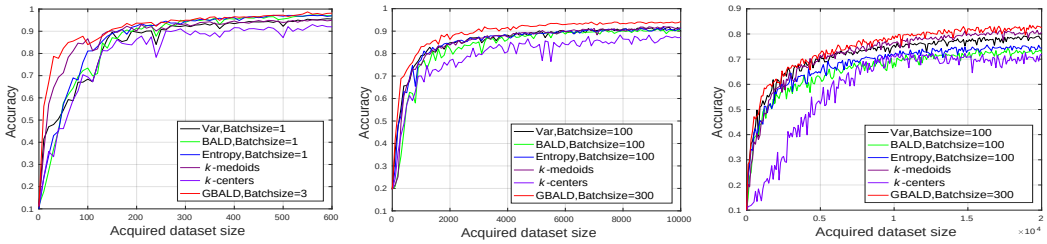

(a) MNIST (b) SVHN (c) CIFAR10

Figure 5: Active acquisitions on MNIST, SVHN, and CIFAR10 datasets.

Table 1: Mean±std of the test accuracies of the breakpoints of the learning curves on MNIST, SVHN, and CIFAR-10.

Datasets	Algorithms					
	Var	BALD	Entropy	k-medoids	k-centers	GBALD
MNIST	0.8419± 0.1721	0.8645±0.1909	0.8498±0.2098	0.8785±0.1433	0.8052±0.1838	**0.9106±0.1296**
SVHN	0.8535±0.1098	0.8510±0.1160	0.8294±0.1415	0.8498±0.1294	0.7909±0.1235	**0.8885±0.1054**
CIFAR-10	0.7122±0.1034	0.6760±0.1023	0.6536±0.1038	0.71837±0.1245	0.5890±0.1758	**0.7440±0.1087**

k-centers algorithm performs more poorly than other compared baselines because the representative optimization with sphere geodesic usually falls into the selection of boundary data; 2) Var, Entropy and BALD algorithms cannot accelerate the network model rapidly due to highly-skewed acquisitions towards few fixed classes at its first acquisitions (start states); 3) k-medoids approach does not interact with the neural network model while directly imports the clustering centers into its training set; 4) The accuracies of the acquisitions of GBALD achieve better performance at the beginning than the Var, Entropy and BALD approaches which fed the training set of the network model via acquisition loops. In short, **the network is improved faster after drawing the distribution characteristics** of the input dataset with sufficient labels. GBALD thus consists of the representative and informative acquisitions in its uniform framework. Advantages of these two acquisition paradigms are integrated and present higher accuracies than any single paradigm.

Table 2: Number of acquisitions on MNIST, SVHN and CIFAR10 until 70%, 80%, and 90% accuracies are reached.

Algorithms	Accuracies		
	70%	80%	90%
Var	140/1,700/5,700	150/2,200/>20,000	210/>10,000/>6,100
BALD	110/1,700 /8,800	120 /2,300/>20,000	190/7,100 / >20,000
Entropy	110/1,900/11,200	150/2,400/>20,000	200/8,600/>20,000
k-modoids	70/1,700/5,900	90/2,200/16,000	**170**/6,200 />20,000
k-centers	110/2,000/10,100	150/3,800/>20,000	280/>10,000/>20,000
GBALD	**50/1,400/4,800**	**70/1,900/12,200**	**170/3,900**/>20,000

Table 1 reports the mean±std values of the test accuracies of the breakpoints of the learning curves in Figure 5, where breakpoints of MNIST are $\{0, 10, 20, 30, ..., 600\}$, breakpoints of SVHN are $\{0, 100, 200, ..., 10000\}$, and breakpoints of CIFAR10 are $\{0, 100, 200, ..., 20000\}$. The results show that 1) GBALD achieves the highest average accuracies; k-medoids is ranked the second amongst the compared baselines; 2) k-centers has ranked the worst accuracies amongst these approaches; 3) the others, which iteratively update the training model are ranked at the middle including BALD, Var and Entropy algorithms. Table 2 shows the acquisition numbers of achieving the accuracies of 70%, 80%, and 90% on the three datasets. The three numbers of each cell are the acquisition numbers over MNIST, SVHN, and CIFAR10, respectively. The results show that GBALD can use fewer acquisitions to achieve a desired accuracy than the other algorithms.

6 Conclusion

We have introduced a novel Bayesian AL framework, GBALD, from the geometric perspective, which seamlessly incorporates representative (core-set) and informative (model uncertainty estimation) acquisitions to accelerate the training of a DNN model. Our GBALD yields significant improvements over BALD, flexibly resolving the limitations of an uninformative prior and redundant information by optimizing the acquisition on an ellipsoid. Generalization analysis has asserted that training with ellipsoid has tighter lower error bound and higher probability to achieve a zero error than training with a typical sphere. Compared to the representative or informative acquisition algorithms, experiments show that our GBALD spends much fewer acquisitions to accelerate the accuracy. Moreover, it keeps slighter accuracy reduction than other baselines against repeated and noisy acquisitions.

REFERENCES

Arsenii Ashukha, Alexander Lyzhov, Dmitry Molchanov, and Dmitry Vetrov. Pitfalls of in-domain uncertainty estimation and ensembling in deep learning. In *International Conference on Learning Representations*, 2019.

Mihai Bādoiu, Sariel Har-Peled, and Piotr Indyk. Approximate clustering via core-sets. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pp. 250–257, 2002.

Shai Ben-David and Ulrike Von Luxburg. Relating clustering stability to properties of cluster boundaries. In *21st Annual Conference on Learning Theory (COLT 2008)*, pp. 379–390. Omnipress, 2008.

Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pp. 1613–1622, 2015.

Trevor Campbell and Tamara Broderick. Bayesian coreset construction via greedy iterative geodesic ascent. In *International Conference on Machine Learning*, pp. 698–706, 2018.

Trevor Campbell and Tamara Broderick. Automated scalable bayesian inference via hilbert coresets. *The Journal of Machine Learning Research*, 20(1):551–588, 2019.

David Cohn, Les Atlas, and Richard Ladner. Improving generalization with active learning. *Machine learning*, 15(2):201–221, 1994.

Arnaud Doucet, Simon Godsill, and Christophe Andrieu. On sequential monte carlo sampling methods for bayesian filtering. *Statistics and computing*, 10(3):197–208, 2000.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059, 2016.

Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1183–1192. JMLR. org, 2017.

Mingfei Gao, Zizhao Zhang, Guo Yu, Sercan O Arik, Larry S Davis, and Tomas Pfister. Consistency-based semi-supervised active learning: Towards minimizing labeling cost. *ECCV*, 2020.

Daniel Golovin, Andreas Krause, and Debajyoti Ray. Near-optimal bayesian active learning with noisy observations. In *Advances in Neural Information Processing Systems*, pp. 766–774, 2010.

Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in neural information processing systems*, pp. 8527–8537, 2018.

Sariel Har-Peled and Soham Mazumdar. On coresets for k-means and k-median clustering. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pp. 291–300, 2004.

Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.

Khaled Jedoui, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Deep bayesian active learning for multiple correct outputs. *arXiv preprint arXiv:1912.01119*, 2019.

Andreas Kirsch, Joost van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. In *Advances in Neural Information Processing Systems*, pp. 7024–7035, 2019.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems*, pp. 6402–6413, 2017.

Aaron Lou, Isay Katsman, Qingxuan Jiang, Serge Belongie, Ser-Nam Lim, and Christopher De Sa. Differentiating through the frechet mean. *ICML*, 2020.

Feiping Nie, Hua Wang, Heng Huang, and Chris Ding. Early active learning via robust representation and structured sparsity. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.

Michael Osborne, Roman Garnett, Zoubin Ghahramani, David K Duvenaud, Stephen J Roberts, and Carl E Rasmussen. Active learning of model evidence using bayesian quadrature. In *Advances in neural information processing systems*, pp. 46–54, 2012.

Hae-Sang Park and Chi-Hyuck Jun. A simple and fast algorithm for k-medoids clustering. *Expert systems with applications*, 36(2):3336–3341, 2009.

Robert Pinsler, Jonathan Gordon, Eric Nalisnick, and José Miguel Hernández-Lobato. Bayesian batch active learning as sparse subset approximation. In *Advances in Neural Information Processing Systems*, pp. 6356–6367, 2019.

Nicholas Roy and Andrew McCallum. Toward optimal active learning through monte carlo estimation of error reduction. *ICML, Williamstown*, pp. 441–448, 2001.

Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL `https://openreview.net/forum?id=H1aIuk-RW`.

Min Tang, Xiaoqiang Luo, and Salim Roukos. Active learning for statistical natural language parsing. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pp. 120–127. Association for Computational Linguistics, 2002.

Stephen A Vavasis. Approximation algorithms for indefinite quadratic programming. *Mathematical Programming*, 57(1-3):279–311, 1992.

Zengmao Wang, Bo Du, Weiping Tu, Lefei Zhang, and Dacheng Tao. Incorporating distribution matching into uncertainty for multiple kernel active learning. *IEEE Transactions on Knowledge and Data Engineering*, 2019.