Black-box Pseudodata Variational Inference

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recent advances in coreset methods have shown that a selection of representative datapoints can replace massive volumes of data for Bayesian inference, preserving the relevant statistical information and significantly accelerating subsequent downstream tasks. Existing variational coreset constructions rely on either selecting subsets of the observed datapoints, or jointly performing approximate inference and optimizing pseudodata in the observed space akin to inducing points methods in Gaussian Processes. So far, both approaches are limited by complexities in evaluating their objectives for general purpose models, and require generating samples from a typically intractable posterior over the coreset throughout inference and testing. In this work, we present a black-box variational inference algorithm for coresets that overcomes these constraints and enables principled application of variational coresets to intractable models, such as Bayesian neural networks. We apply our techniques to supervised learning problems, and compare it with existing approaches in the literature for data summarization and inference.

1 Introduction

Recent work [25] has demonstrated that Bayesian posteriors on large scale datasets can be sparsely represented via surrogate densities defined through a set of learnable weighted pseudo-observations in data space. In a similar vein with inducing points in Gaussian Processes [38], these target to represent parsimoniously the sufficient statistics of the observed dataset that are required for inference in the model, and learned via a variational scheme termed Pseudodata Sparse Variational Inference (PSVI). Replacing the original dataset via the reduced set of pseudo-observations unlocks the potential of scaling up downstream learning tasks, compressing for efficient storage and visualisation, and performing rapid model exploration. See Fig. 1 for an introductory visualization. However, in the general case the achievable quality of data summarizations for posterior approximations is often limited by the intractability of the original model, hindering the attainable convergence guarantees of coreset constructions.

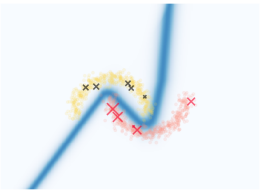

Figure 1: Binary classification example with mean-field. Crosses mark the weighted pseudodata, blue shades predictive uncertainty.

In an effort to make PSVI available to general black-box models, we make the following contributions:

1. we derive a principled interpretation of variational coresets as rich variational families,
2. we modify the objective to obtain a black-box version of it using importance sampling and a nested variational inference step,
3. utilizing this objective, we propose BB PSVI and BB Sparse VI, two novel suites of variational inference algorithms for black-box probabilistic models,
4. we empirically compare to prior Bayesian and non-Bayesian coreset techniques, and
5. we showcase these techniques on previously infeasible models, Bayesian neural networks.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

2 Background

2.1 Variational Inference

Consider a modeling problem where we are given a standard model:

$$p(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{\theta}) = p(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta})p(\boldsymbol{\theta}). \tag{1}$$

Given a dataset $\mathcal{D} = \{(\boldsymbol{x}, \boldsymbol{y})\} \in (\mathcal{X} \times \mathcal{Y})^N$, the task of inference and maximization of the marginal likelihood $p(\boldsymbol{y}|\boldsymbol{x})$ entails estimation of the posterior distribution over model parameters $\boldsymbol{\theta}$, $p(\boldsymbol{\theta}|\boldsymbol{x}, \boldsymbol{y})$. In general, the posterior $p(\boldsymbol{\theta}|\boldsymbol{x}, \boldsymbol{y})$ is not tractable. A common solution is to resort to approximate inference, for instance via variational inference (VI). In VI, we assume that we can approximate the true posterior within a *variational family* $q(\boldsymbol{\theta}; \boldsymbol{\lambda})$ with free variational parameters $\boldsymbol{\lambda}$, by minimizing the Kullback-Leibler (KL) divergence between the approximate and the true posterior $\mathrm{D}_{\mathrm{KL}}\left(q(\boldsymbol{\theta}; \boldsymbol{\lambda})||p(\boldsymbol{\theta}|\boldsymbol{x}, \boldsymbol{y})\right)$, or equivalently by maximizing the Evidence Lower Bound (ELBO):

$$\log p(\boldsymbol{y}|\boldsymbol{x}) \geq \text{ELBO}(\boldsymbol{\lambda}) = \mathbb{E}_{q(\boldsymbol{\theta};\boldsymbol{\lambda})}\left[\log \frac{p(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta})p(\boldsymbol{\theta})}{q(\boldsymbol{\theta}; \boldsymbol{\lambda})}\right]. \tag{2}$$

2.2 (Bayesian) coresets and pseudocoresets

Bayesian coreset constructions aim to provide a highly automated data austere approach to the problem of inference, given a dataset $\mathcal{D}$ and a probabilistic model Eq. (1). Agnostic to the particular inference method used, Bayesian coresets shrink in a model-specific way the original dataset to a much sparser weighted dataset representation $\{(\boldsymbol{v}, \boldsymbol{x}, \boldsymbol{y})_{n=1}^N\} \in (\mathbb{R}_{\geq 0} \times \mathcal{X} \times \mathcal{Y})^N$, where $||\boldsymbol{v}||_0 = M \ll N$, that maximizes the pasimony of observations considered over inference, while allowing close approximations of the full data posterior, i.e. $p(\boldsymbol{\theta}|\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{v}) \approx p(\boldsymbol{\theta}|\boldsymbol{x}, \boldsymbol{y})$. Initial Bayesian coreset constructions formulated the problem as yielding a sparsity-constrained approximation of the true data likelihood function uniformly over the parameter space. This perspective has admitted solutions via convexifications relying on importance sampling and conditional gradient methods [19, 10], greedy methods for geodesic ascent [9], as well as non-convex approaches employing iterative hard thresholding [43]. Recently, Sparse VI [8] reformulated the problem as performing directly VI within a sparse exponential family of approximate posteriors defined on the coreset datapoints.

The idea of using a set of learnable weighted *pseudodata* (or *inducing points*) $(\boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}) := \{(\boldsymbol{v}_m, \boldsymbol{u}_m, \boldsymbol{z}_m)_{m=1}^M\} \in (\mathbb{R}_{\geq 0} \times \mathcal{X} \times \mathcal{Y})^M$ to summarize a dataset for the purpose of posterior inference was introduced in PSVI [25], as the following objective:

$$\boldsymbol{v}^\star, \boldsymbol{u}^\star, \boldsymbol{z}^\star = \underset{\boldsymbol{v},\boldsymbol{u},\boldsymbol{z}}{\arg\min}\, \mathrm{D}_{\mathrm{KL}}\left(p(\boldsymbol{\theta}|\boldsymbol{u}, \boldsymbol{z}, \boldsymbol{v})||p(\boldsymbol{\theta}|\boldsymbol{x}, \boldsymbol{y})\right). \tag{3}$$

This objective facilitates learning weighted pseudodata such that the exact posterior over pseudodata $p(\boldsymbol{\theta}|\boldsymbol{u}, \boldsymbol{z}, \boldsymbol{v})$ approximates the exact posterior over the true dataset $p(\boldsymbol{\theta}|\boldsymbol{x}, \boldsymbol{y})$, and was demonstrated on a variety of models as a generally applicable objective. However, in [25] it was assumed that this objective is tractable as demonstrated in experiments on models admitting analytical posteriors, or via the use of heuristics to sample from the coreset posterior, as is the case in a logistic regression employing a Laplace approximation.

Sparse VI Sparse VI [8] targets to incrementally minimize the rhs of Eq. (3) via fixing the coreset points locations and optimizing their weights $\boldsymbol{v}$. In its original construction, the coreset points are initialized to the empty set and $\boldsymbol{v}$ to $\mathbf{0}$. In the general case of an intractable model, the coreset posterior is approximately computed via fitting a Laplace approximation $q_{\text{Lap}}(\boldsymbol{\theta}; \boldsymbol{v})$ on the weighted datapoints of the coreset $(\boldsymbol{v}, \boldsymbol{x}, \boldsymbol{y})$. The next point selection step involves: (*i*) computing approximations of the centered log-likelihood function of each datapoint $(\boldsymbol{x}_i, y_i)$ for model parameters $\boldsymbol{\theta}_s$ sampled from the current coreset posterior, via an MC estimate $\tilde{\mathbf{f}}(\boldsymbol{x}_n, y_n, \boldsymbol{\theta}_s) := \left(\log p(y_n|\boldsymbol{x}_n, \boldsymbol{\theta}_s) - \frac{1}{S}\sum_{s'=1}^S \log p(y_n|\boldsymbol{x}_n, \boldsymbol{\theta}_{s'})\right) \in \mathbb{R}^S$, $\boldsymbol{\theta}_{s'} \sim q_{\text{Lap}}(\boldsymbol{\theta}_{s'}; \boldsymbol{v})$, (*ii*) computing correlations with the residual error of the total data log-likelihood $\boldsymbol{r}(\boldsymbol{v}) = \mathbf{f}^T(1 - \boldsymbol{v})$, and (*iii*) making a greedy selection of the point that maximizes the correlation between the two: $n^\star \in \underset{n\in[N]}{\arg\max} \begin{cases} |\text{Corr}(\tilde{\mathbf{f}}, \boldsymbol{r})| & v_n > 0 \\ \text{Corr}(\tilde{\mathbf{f}}, \boldsymbol{r}) & v_n = 0 \end{cases}$, where $\text{Corr}(\tilde{\mathbf{f}}, \boldsymbol{r}) := \text{diag}\left[\frac{1}{S}\sum_s \tilde{\mathbf{f}}_s\tilde{\mathbf{f}}_s^T\right]^{-\frac{1}{2}}\left(\frac{1}{S}\sum \tilde{\mathbf{f}}_s\boldsymbol{r}_s^T\right)$. Subsequently, the vector $\boldsymbol{v}$ gets optimized via minimizing the KL divergence between the coreset

and the true posterior. For the variational family of coresets, under the assumption that $\boldsymbol{\theta}$ is sampled from the *true* coreset posterior, the formula of the gradient of the objective in Eq. (3) can be derived, and approximated via resampling from the coreset over the course of $\boldsymbol{v}$ optimization. At the end of the optimization, the extracted points and weights are given as input to an approximate inference algorithm—commonly a Laplace approximation—which provides the final estimate of the posterior.

3 Black-box coreset Variational Inference

In order to derive a general PSVI objective, we observe that we first need to posit a variational family involving the pseudodata appropriately. Let's assume the (intractable) variational family $q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z}) := p(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})$, which has the following particular structure: it is parametrized as the posterior over $\boldsymbol{\theta}$ as estimated given inducing points $\boldsymbol{u}$, $\boldsymbol{z}$. In this case, $q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})$ can be thought of as a *variational program* (see [34]) with parameters $\{\boldsymbol{u},\boldsymbol{z}\}$ and expressed as:

$$q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z}) := p(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z}) = \frac{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\boldsymbol{z}|\boldsymbol{u})}, \tag{4}$$

where the marginal likelihood term $p(\boldsymbol{z}|\boldsymbol{u})$ is typically intractable.

Utilizing this for the formulation of a pseudo-coreset variational inference algorithm leads to the following lower bound on the marginal likelihood of the *true* data:

$$\text{ELBO}_{\text{PSVI}}(\boldsymbol{u},\boldsymbol{z}) = \mathbb{E}_{q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})}\left[\log \frac{p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})}\right]. \tag{5}$$

Performing VI involves maximizing Eq. (5) with respect to the parameters $\{\boldsymbol{u},\boldsymbol{z}\}$.

Scrunitizing this objective reveals that the posterior $p(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})$ is used in two locations: **first**, as the *sampling* distribution for $\boldsymbol{\theta}$, second to evaluate log-densities of samples (*scoring*). As such, without further assumptions on the inference algorithm, this objective can only be evaluated if the posterior $p(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})$ can be readily evaluated in closed form and sampled from directly, causing this expression to remain intractable but for the simplest of cases; **second**, the same quantity exists inside the logarithm as a *quantity to evaluate*, revealing the need to make it computable for any sample $\boldsymbol{\theta}$.

In Section 3.1 we will derive a black-box objective for PSVI to overcome the limitation to tractable posteriors by adressing these two points separately and composing these solutions to a joint black box objective, which we will use for deriving black box versions of PSVI and Sparse-VI.

In the appendix we also derive a Gaussian Process version of our algorithm as a special case where the posterior is tractable to formally link this work to inducing points for GPs.

3.1 PSVI for intractable posteriors

In Sec. 3 we identified two key problems to evaluate Eq. (5) involving $p(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})$. To tackle the first, we devise a strategy to utilize a tractable distribution $q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z};\boldsymbol{\psi})$ paired with importance sampling [30] to be able to *sample* from that distribution. We assume that we can draw K samples $\boldsymbol{\theta}_k \sim q(\boldsymbol{\theta};\boldsymbol{\psi})$,[1] however we cannot plug these into Eq. (5) directly as they do not follow the exact distribution. Yet, we can leverage a self-normalized importance sampling (IS) [30] correction to obtain the desired approximate *samples* from the coreset posterior. Assigning the samples unnormalized weights:

$$w_k = \frac{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta}_k)p(\boldsymbol{\theta}_k)}{q(\boldsymbol{\theta}_k;\boldsymbol{\psi})}, \tag{6}$$

for an appropriately selected q such that support$(p) \subseteq$ support(q), we denote the corresponding normalized weights $\tilde{w}_k = w_k / \sum_j w_j$. Using $q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})$ from Eq. (4), we can approximate Eq. (5) as:

$$\begin{aligned}\text{ELBO}_{\text{PSVI-IS}}(\boldsymbol{u},\boldsymbol{z}) &= \int_{\boldsymbol{\theta}} q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z}) \log \frac{p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta})p(\boldsymbol{\theta})}{q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})} d\boldsymbol{\theta} \qquad (7)\\ &= \int_{\boldsymbol{\theta}} \frac{q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})}{q(\boldsymbol{\theta};\boldsymbol{\psi})} q(\boldsymbol{\theta};\boldsymbol{\psi}) \log \frac{p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta})p(\boldsymbol{\theta})}{q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})} d\boldsymbol{\theta} \approx \sum_{\boldsymbol{\theta}_i \sim q(\boldsymbol{\theta};\boldsymbol{\psi})} \tilde{w}_i \log \frac{p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta}_i)p(\boldsymbol{\theta}_i)}{q(\boldsymbol{\theta}_i|\boldsymbol{u},\boldsymbol{z})}.\end{aligned}$$

[1] We will also denote this sampling as $\boldsymbol{\theta} \sim q$ for brevity in the following.

After having tackled sampling, we shift our attention to the second problem, which was about *evaluating $p(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})$ inside the logarithm*. Applying Jensen's inequality and Eq. (4), we get an evidence lower bound of the inducing points (IP):

$$\log p(\boldsymbol{z}|\boldsymbol{u}) \geq \text{ELBO}_{\text{IP}}(\boldsymbol{\psi}) = \mathop{\mathbb{E}}_{q(\boldsymbol{\theta};\boldsymbol{\psi})} \log \frac{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta})p(\boldsymbol{\theta})}{q(\boldsymbol{\theta};\boldsymbol{\psi})}. \tag{8}$$

Thus, substituting $q(\boldsymbol{\theta}_i|\boldsymbol{u},\boldsymbol{z})$, $i = 1,\ldots,M$, per Eq. (4) in Eq. (7), we can derive the following multi-sample importance sampling lower bound on the evidence of the *observed* data:

$$\text{ELBO}_{\text{PSVI-IS}}(\boldsymbol{u},\boldsymbol{z}) \approx \sum_{\boldsymbol{\theta}_i\sim q} \tilde{w}_i \left[\log \frac{p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta}_i)p(\boldsymbol{z}|\boldsymbol{u})}{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta}_i)}\right] = \sum_{\boldsymbol{\theta}_i\sim q} \tilde{w}_i \log \frac{p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta}_i)}{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta}_i)} + \log p(\boldsymbol{z}|\boldsymbol{u}).$$

We now realize that by lower bounding $\log p(\boldsymbol{z}|\boldsymbol{u})$ using $\text{ELBO}_{\text{IP}}(\boldsymbol{\psi})$ we are jointly deriving a lower bound to $\text{ELBO}_{\text{PSVI-IS}}(\boldsymbol{u},\boldsymbol{z})$, which we denote as $\text{ELBO}_{\text{PSVI-IS-BB}}(\boldsymbol{u},\boldsymbol{z},\boldsymbol{\psi})$. We remind the reader at this point that this ensures the following relationship holds rigorously:

$$\log p(\boldsymbol{y}|\boldsymbol{x}) \geq \text{ELBO}_{\text{PSVI-IS}}(\boldsymbol{u},\boldsymbol{z}) \geq \text{ELBO}_{\text{PSVI-IS-BB}}(\boldsymbol{u},\boldsymbol{z},\boldsymbol{\psi}),$$

which means that the following multi-sample black-box objective, $\text{ELBO}_{\text{PSVI-IS-BB}}$, is a valid variational inference objective for the original problem under which we can tractably *score* $\boldsymbol{u},\boldsymbol{z}$:

$$\begin{aligned}
&\text{ELBO}_{\text{PSVI-IS}}(\boldsymbol{u},\boldsymbol{z}) \geq \text{ELBO}_{\text{PSVI-IS-BB}}(\boldsymbol{u},\boldsymbol{z},\boldsymbol{\psi}) \\
&= \sum_{\boldsymbol{\theta}_i\sim q} \tilde{w}_i \log \frac{p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta}_i)}{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta}_i)} + \mathop{\mathbb{E}}_{\boldsymbol{\theta}\sim q}\left[\log \frac{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta})p(\boldsymbol{\theta})}{q(\boldsymbol{\theta};\boldsymbol{\psi})}\right] \\
&\approx \sum_{\boldsymbol{\theta}_i\sim q} \left[\tilde{w}_i \log \frac{p(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta}_i)}{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta}_i)} + \frac{1}{K}\log \frac{p(\boldsymbol{z}|\boldsymbol{u},\boldsymbol{\theta}_i)p(\boldsymbol{\theta}_i)}{q(\boldsymbol{\theta}_i;\boldsymbol{\psi})}\right].
\end{aligned} \tag{9}$$

Maximizing the objective of Eq. (9) can directly serve the purposes of black-box VI, as it lifts the requirement for sampling from the exact posterior on the summarizing data, proposing instead importance sampling with the use of $q(\boldsymbol{\theta};\boldsymbol{\psi})$ constrained on the summary. Crucially, as the $\text{ELBO}_{\text{PSVI-IS-BB}}$ mixes information from the true and inducing datapoints, the assumed probabilistic structure of the variational distribution should be enforced over the optimization of Eq. (9), namely that ψ is being inferred only on the basis of the inducing points via the variational program of Eq. (8), whilst $(\boldsymbol{u},\boldsymbol{z})$ get optimized according to Eq. (9), so that the information of the true dataset gets distilled in a way that increases the evidence lower bound. Moreover, replacing our pseudodata dependent importance weights with a uniform distribution $\tilde{w}_k = 1/K$, $k = 1,\ldots,K$, or using a single-sample objective (i.e. in the degenerate case when $K = 1$), allows us to recover the classical ELBO of Eq. (2), cancelling the pseudodata likelihood terms in the variational objective.[2] Iteratively maximizing the $\text{ELBO}_{\text{PSVI-IS-BB}}$ enables *adapting* the importance sampling scheme of Eq. (6) to the intractable coreset posterior, automating the selection of improved proposals q for the importance weighting throughout inference, within the confines of the chosen variational family.

Finally, to compute expectations under the intractable true posterior using samples from the learned coreset posterior, assuming $q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z}) \approx p(\boldsymbol{\theta}|\boldsymbol{x},\boldsymbol{y})$, we can sample from q and correct again using the above defined importance sampling scheme as follows:

$$\mathop{\mathbb{E}}_{p(\boldsymbol{\theta}|\boldsymbol{x},\boldsymbol{y})}[f(\boldsymbol{\theta})] = \int f(\boldsymbol{\theta})p(\boldsymbol{\theta}|\boldsymbol{x},\boldsymbol{y})d\boldsymbol{\theta} \approx \int f(\boldsymbol{\theta})\frac{q(\boldsymbol{\theta}|\boldsymbol{u},\boldsymbol{z})}{q(\boldsymbol{\theta};\boldsymbol{\psi})}q(\boldsymbol{\theta};\boldsymbol{\psi})d\boldsymbol{\theta} \approx \sum_{\boldsymbol{\theta}_i\sim q(\boldsymbol{\theta};\boldsymbol{\psi})} \tilde{w}_i f(\boldsymbol{\theta}_i). \tag{10}$$

3.2 Incremental and Batch Sparse BBVI

The estimator we derive can also be used to design a black box version of the Sparse VI algorithm introduced in [8] (see Sec. 2.2), by the insight that only the variational parameters for locations are not updated and we can apply various choices to the overall algorithmic flow. The incremental construction scheme originally can become black-box via modifying how the approximate posterior on the coreset data gets computed, and introducing a generalized objective for the optimization of model evidence involving the coreset parameters. In this construction (*Sparse BBVI*), we can

[2] The nested variational program Eq. (8) still allows propagation of non-zero gradients w.r.t. the pseudodata.

posit a standard variational family $q(\boldsymbol{\theta};\boldsymbol{\psi})$ (e.g. mean-field variational distributions) and maximize the corresponding ELBO computed on the weighted datapoints of the coreset. Using the extracted variational approximation, we can draw samples from the coreset and correct for the model. Next, over the greedy selection step, we can use the samples and importance weights to compute the centered log-likelihood vectors $\tilde{\mathbf{f}}(\boldsymbol{x}_n, y_n, \boldsymbol{\theta}_s) := \log p(y_n|\boldsymbol{x}_n, \boldsymbol{\theta}_s) - \sum_{s'} w_{s'} \log p(y_n|\boldsymbol{x}_n, \boldsymbol{\theta}_{s'})$, and select the next datapoint via the greedy correlation maximization criterion. Subsequently, to refine the weight vector towards the true data posterior, we can maximize wrt $\boldsymbol{v}, \boldsymbol{\psi}$ an extension of the PSVI ELBO per Eq. (9), where the coreset data log-likelihood is multiplied by the weight vector $\boldsymbol{v}$ as in Eq. (14):

$$\text{ELBO}_{\text{Sparse-BBVI}}(\boldsymbol{v}, \boldsymbol{\psi}) = \sum_{\boldsymbol{\theta}_i \sim q} \left[\tilde{w}_i \log \frac{p(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta}_i)}{p(\boldsymbol{y}|\boldsymbol{v}, \boldsymbol{x}, \boldsymbol{\theta}_i)} + \frac{1}{K} \log \frac{p(\boldsymbol{y}|\boldsymbol{v}, \boldsymbol{x}, \boldsymbol{\theta}_i) p(\boldsymbol{\theta}_i)}{q(\boldsymbol{\theta}_i; \boldsymbol{\psi})} \right]. \tag{11}$$

The importance sampling scheme for the posterior samples is defined identically to Eq. (6), after substituting the likelihood term with $p(\boldsymbol{y}|\boldsymbol{v}, \boldsymbol{x}, \boldsymbol{\theta}_i)$ for each sample $\boldsymbol{\theta}_i \sim q(\boldsymbol{\theta}_i; \boldsymbol{\psi})$. At each iteration of the algorithm, this variational objective (which, following similar reasoning with the previous section, can be shown to be a lower bound of the evidence), can be maximized with respect to the coreset and variational parameters $\boldsymbol{v}, \boldsymbol{\psi}$. Moreover, we can omit the incremental inclusion of points to the coreset, and consider a *batch* version of this construction, where we initialise at a random subset of the original dataset, keep fixed coreset point locations, and optimize only the weights attached to the coreset support using Eq. (11) with non-negativity constraints. Finally, we can use the variational approximation fitted on the coreset to compute predictive posteriors on unseen datapoints correcting for our importance sampling.

We also define a variant of the algorithm without incremental selection through a simple *pruning* strategy based on the weighting assigned to the coreset points when optimizing the ELBO of Eq. (11). We can start from a large coreset size and, after training, reduce the size of the summary by keeping the top-K datapoints and re-initialising.

3.3 Variational Families: (Un)weighted inducing points

When studying the objective in Eq. 5 it quickly becomes evident that the exact amount of chosen pseudo-data (coreset) defines the available evidence for that particular posterior construction and as such is a quantity of interest. We will consider this quantity *pseudo-evidence*.

For notational brevity so far we have considered the variational parameters to represent pseudodata directly, while rescaling using the fixed data compression ratio to maintain invariance of the total pseudo-evidence to the coreset size. We will be calling this setting *unweighted pseudodata*.

One choice we can make now is to ensure that the available pseudo-evidence is invariant to the chosen size of the coreset set and potentially equal to the regular dataset. To achieve that, we consider setting:

$$q(\boldsymbol{\theta}|\boldsymbol{u}, \boldsymbol{z}) := p(\boldsymbol{\theta}|\boldsymbol{u}, \boldsymbol{z}) = \frac{p(\boldsymbol{z}|\boldsymbol{u}, \boldsymbol{\theta})^{N/M} p(\boldsymbol{\theta})}{p(\boldsymbol{z}|\boldsymbol{u})} = \frac{\prod^{M} p(\boldsymbol{z}_m|\boldsymbol{u}_m, \boldsymbol{\theta})^{N/M} p(\boldsymbol{\theta})}{p(\boldsymbol{z}|\boldsymbol{u})}, \tag{12}$$

with $p(\boldsymbol{z}|\boldsymbol{u})$ similarly reweighted via the selected observational *data compression ratio* N/M, i.e. for continuous parameter spaces $p(\boldsymbol{z}|\boldsymbol{u}) := \int p(\boldsymbol{z}|\boldsymbol{\theta}, \boldsymbol{u})^{N/M} p(\boldsymbol{\theta}) d\boldsymbol{\theta}$.

We may also consider a scenario where we have more generally *weighted pseudodata*, which involves an additional learnable parameter $\boldsymbol{v}_i \geq 0$ per datapoint $\boldsymbol{u}_i$ and permits the following interpretation: $\boldsymbol{v}_i \sim \text{Mult}(\boldsymbol{v})$. Sampling repeatedly from this distribution yields a collection of inducing points with their frequencies governed by their respective weights:

$$q(\boldsymbol{\theta}|\boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}) := p(\boldsymbol{\theta}|\boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}) = \frac{p(\boldsymbol{\theta}) \prod p(\boldsymbol{z}_i|\boldsymbol{u}_i, \boldsymbol{\theta}_i)^{\boldsymbol{v}_i}}{p(\boldsymbol{z}|\boldsymbol{v}, \boldsymbol{u})}. \tag{13}$$

This allows us to introduce the following simplifying notation for weighted log-likelihoods, which will be used in the continuation:

$$\log p(z_i|v_i, \boldsymbol{u}_i, \boldsymbol{\theta}_i) := v_i \log p(z_i|\boldsymbol{u}_i, \boldsymbol{\theta}_i). \tag{14}$$

We can now also combine the ideas of learning the weights of pseudodata and controlling the amount of pseudoevidence by posing parametrizations where v_i is non-negative and sums to 1

(e.g. $v_i := \texttt{softmax}(\beta_i)$, $i = 1, \ldots, M$), and adding global variational parameters to optimize the magnitude of total evidence over the pseudodata (e.g. $\log p(\boldsymbol{z}|\boldsymbol{v}, \boldsymbol{u}, \boldsymbol{\theta}, \alpha) := \alpha \boldsymbol{v}^T \log p(\boldsymbol{z}|\boldsymbol{u}, \boldsymbol{\theta})$), where α can be set by hand (i.e. to N/M) or be a learnable quantity. We will explore various such choices for coreset parametrizations in our experiments and the effects they have on learning.

3.4 Nested optimization

When using a nested optimizer for the maximization of the $\text{ELBO}_{\text{PSVI-IS-BB}}$, rearranging the terms of the objective enables us to place our method in the generic cardinality-constrained bilevel optimization framework for data distillation and coreset constructions proposed in [21, 5]. VI maximizing the objective of Eq. (5) inside a parametric family $q(\boldsymbol{\theta}; \boldsymbol{\psi})$ can be reformulated as a bilevel optimization problem in the following way:

$$\boldsymbol{v}^\star, \boldsymbol{u}^\star, \boldsymbol{z}^\star = \underset{\boldsymbol{v},\boldsymbol{u},\boldsymbol{z}}{\arg\min}\, L(\boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}, \boldsymbol{\psi}^\star(\boldsymbol{u}, \boldsymbol{z}); \boldsymbol{x}, \boldsymbol{y}) \qquad \text{s.t.} \quad \boldsymbol{\psi}^\star = \underset{\boldsymbol{\psi}}{\arg\min}\, \ell(\boldsymbol{\psi}, \boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}), \tag{15}$$

where

$$L(\boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}, \boldsymbol{\psi}(\boldsymbol{u}, \boldsymbol{z})) = -\text{ELBO}_{\text{PSVI-IS}}(\boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}, \boldsymbol{\psi}) \tag{16}$$

and

$$\ell(\boldsymbol{\psi}, \boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}) = -\text{ELBO}(\boldsymbol{\psi}, \boldsymbol{v}, \boldsymbol{u}, \boldsymbol{z}) := \mathbb{E}_{\boldsymbol{\theta}_i \sim q} \log \frac{p(v_i, \boldsymbol{z}|\boldsymbol{u}, \boldsymbol{\theta}_i) p(\boldsymbol{\theta}_i)}{q(\boldsymbol{\theta}_i; \boldsymbol{\psi})}. \tag{17}$$

The bilevel structure allow us some freedom in treating our optimization problem. Two principled approaches are:

- Iterative differentiation [24]: We trace the gradient computation over the optimization of the variational parameters $\boldsymbol{\psi}$ when solving the inner problem, and use this information at the computation of the outer objective gradient wrt the coreset parameters.
- Approximate hypergradients [32, 21, 18]: We can compute approximations of the outer gradient using the information of an approximate solution in the inner problem, without fully tracing the gradient throughout the inner optimization trajectory.

Our method of choice in the experimental evaluation will be iterative differentiation, while we include some additional results using hypergradients in the supplementary material.

4 Experiments

In this section we evaluate the performance of our inference framework in intractable models and compare against standard variational inference methods and earlier Bayesian coreset constructions, as well as black-box extensions of existing variational coresets that rely on our generalized ELBO Eq. (9).

As a running baseline with unrestricted access to the training data, we use the standard *mean-field* VI with a diagonal covariance matrix. To capture the implications of enforcing data cardinality constraints, we also construct a *random coreset* using a randomly selected data subset with fixed data locations and fixed likelihood multiplicities correcting for the dataset size reduction, which we then use to optimize a mean-field approximate posterior. As earlier work relied on Laplace approximations on coreset data as a black-box approach for sampling for intractable coreset posteriors, we also experiment with Laplace approximations on random data subsets as an extra baseline.

4.1 Logistic regression

In the first experiment we perform inference on logistic regression models fitting 3 publicly available binary classification datasets [41, 14] with sizes ranging between $10K$ and $150K$ datapoints, and 10 and 128 dimensions. At all models we posit normal priors $\boldsymbol{\theta} \sim \mathcal{N}(0, I)$ and consider mean-field variational approximations with diagonal covariance. In this and following experiments, all presented metrics are averaged across 3 independent trials for each method, and we show the means over the trials along with the corresponding standard errors.

Impact of variational parameters We show predictive metrics of log-likelihood and accuracy on the test set inTables 1 and 2. We observe that coreset methods relying on optimizing pseudodata

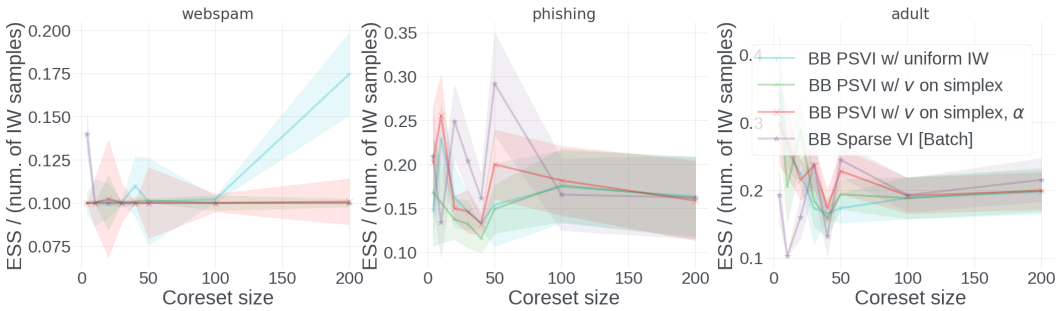

Figure 2: Normalized effective sample size for predictions on test data using 10 Monte Carlo samples from the coreset posterior over increasing coreset size.

Table 1: Test accuracy on 3 logistic regression datasets. ($\star$) denotes use of softmax parameterisation for the weights $\boldsymbol{v} = \texttt{softmax}(\beta)$.

		webspam			phishing			adult		
	M	10	40	100	10	40	100	10	40	100
BB PSVI	$\boldsymbol{v}(\star)$	$\mathbf{92.2} \pm 0.1$	$\mathbf{92.4} \pm 0.1$	$\mathbf{92.3} \pm 0.0$	90.1 ± 0.5	$\mathbf{91.2} \pm 1.0$	90.4 ± 0.0	82.2 ± 0.9	$\mathbf{83.5} \pm 0.2$	83.5 ± 0.0
	$\boldsymbol{v}(\star)$, α	$\mathbf{92.2} \pm 0.2$	92.2 ± 0.1	$\mathbf{92.3} \pm 0.1$	90.7 ± 0.5	90.7 ± 0.9	90.4 ± 0.0	$\mathbf{82.6} \pm 0.9$	$\mathbf{83.5} \pm 0.0$	$\mathbf{83.6} \pm 0.0$
	$\boldsymbol{v}$	$\mathbf{92.2} \pm 0.1$	92.3 ± 0.1	92.2 ± 0.1	90.3 ± 0.8	89.8 ± 0.2	90.3 ± 0.1	80.3 ± 0.8	$\mathbf{83.5} \pm 0.0$	83.5 ± 0.0
	$\boldsymbol{v}{=}\frac{N}{M}\mathbf{1}$	90.9 ± 0.3	91.9 ± 0.1	92.2 ± 0.5	90.2 ± 0.3	88.9 ± 0.7	$\mathbf{90.5} \pm 0.0$	81.0 ± 0.4	$\mathbf{83.5} \pm 0.1$	83.5 ± 0.5
	unif. IW, ($\star$)	92.1 ± 0.1	92.3 ± 0.1	92.2 ± 0.1	$\mathbf{90.8} \pm 1.0$	90.6 ± 0.9	90.4 ± 0.0	82.1 ± 0.8	$\mathbf{83.5} \pm 0.1$	83.5 ± 0.0
	w/o IW, ($\star$)	90.1 ± 0.6	92.0 ± 0.1	$\mathbf{92.3} \pm 0.1$	90.1 ± 1.4	89.2 ± 1.2	90.4 ± 0.2	77.8 ± 1.5	$\mathbf{83.5} \pm 0.0$	83.5 ± 0.1
	$\boldsymbol{v}{=}\mathbf{1}$	61.1 ± 1.8	61.2 ± 4.7	61.2 ± 2.2	85.8 ± 10.4	86.3 ± 11.1	87.3 ± 11.5	75.0 ± 3.7	75.1 ± 1.6	75.8 ± 1.2
Sparse BBVI	Batch	80.1 ± 0.9	86.9 ± 0.2	88.9 ± 0.4	87.6 ± 0.8	90.7 ± 0.3	90.2 ± 1.2	78.1 ± 1.1	82.6 ± 0.4	83.5 ± 0.7
	Incremental	61.2 ± 0.6	65.6 ± 1.2	64.4 ± 7.1	88.8 ± 2.1	88.9 ± 3.7	88.1 ± 2.2	81.3 ± 1.0	79.2 ± 3.1	78.2 ± 1.3
PSVI [25]		83.9 ± 0.7	88.8 ± 0.4	90.3 ± 0.3	88.3 ± 1.7	89.0 ± 0.9	90.1 ± 0.1	82.5 ± 0.4	$\mathbf{83.5} \pm 0.0$	83.5 ± 0.0
Sparse VI [8]		64.8 ± 0.0	71.4 ± 0.0	72.4 ± 0.0	80.8 ± 0.0	82.9 ± 0.0	83.5 ± 0.0	65.6 ± 0.0	64.7 ± 0.0	64.3 ± 0.0
Subset Laplace		65.6 ± 0.0	66.6 ± 0.0	67.0 ± 0.0	69.1 ± 0.0	77.6 ± 0.0	79.6 ± 0.0	75.3 ± 0.0	78.2 ± 0.0	79.3 ± 0.0
Rand. Coreset		71.3 ± 2.9	83.6 ± 1.2	86.3 ± 0.6	81.4 ± 2.7	86.1 ± 3.8	84.4 ± 1.7	75.9 ± 1.1	76.2 ± 0.3	79.5 ± 0.5
Full MFVI				$\mathbf{92.7} \pm 0$			$\mathbf{90.4} \pm 0.1$			$\mathbf{83.5} \pm 0$

have the capacity to approximate quickly inference on the full dataset regardless of dimension, and reach the performance of unrestricted mean-field VI with the use of less than 50 weighted points. In contrast, the approximations that rely on existing datapoints are limited by data dimensionality and need to acquire a larger support to reach the performance of full-data VI (with effects being more evident in webspam, which has a data dimensionality of 128). Including variational parameterisations for optimal scaling of the coreset likelihood aids inference, offering faster convergence. Omitting to correct for dataset reduction ($\boldsymbol{v} = 1$) has detrimental effects on predictive performance. As opposed to earlier constructions limited by heuristics on coreset posterior approximation throughout optimization, our variational framework enables approximating the full data posterior within the assumed variational family. Evaluation of predictive metrics and computation time requirements across a wider range of coreset sizes is included in the Supplement.

Impact of importance weighting In Fig. 2 we plot the normalized effective sample size (ESS) of our importance weighting scheme for 10 Monte Carlo samples computed on test data. We observe that the posteriors can benefit from IW, achieving non-trivial ESS, which tend to converge to close values for large coreset sizes. This finding applies at the testing phase also on ablations of the PSVI variational training that replace the IW scheme with uniform sampling.

4.2 Bayesian Neural Networks

In this section we present inference results on Bayesian neural networks (BNNs), a model class that previous work on Bayesian coresets did not consider due to the absence of a black box variational estimator. In the first part we perform inference via black-box PSVI on 2-dimensional synthetic data using single-hidden layer BNNs, while in the latter part we evaluate the performance of our methods on compressing the MNIST dataset using the LeNet architecture [20].

Simulated datasets We generate two synthetic datasets with size $1k$ datapoints, corresponding to noisy samples from a half-moon shaped 2-class dataset, and a mixture of 4 clusters of data each

Table 2: Test negative log. likelihood on 3 logistic regression datasets.

		webspam			phishing			adult		
	M	10	40	100	10	40	100	10	40	100
BB PSVI	$\boldsymbol{v}(\star)$	21.7 ± 0.3	21.9 ± 0.2	21.2 ± 0.2	36.1 ± 0.8	34.8 ± 3.8	26.0 ± 0.1	$\mathbf{36.1} \pm 0.8$	34.8 ± 3.8	$\mathbf{26.0} \pm 0.1$
	$\boldsymbol{v}(\star)$, α	22.1 ± 0.5	21.6 ± 0.2	21.0 ± 0.0	33.6 ± 0.6	34.2 ± 3.1	$\mathbf{25.9} \pm 0.1$	50.2 ± 3.2	34.0 ± 0.0	33.9 ± 0.0
	$\boldsymbol{v}$	21.9 ± 0.4	21.9 ± 0.1	21.6 ± 0.3	40.3 ± 2.1	48.9 ± 6.2	26.0 ± 0.1	78.0 ± 6.2	34.0 ± 0.1	34.0 ± 0.1
	$\boldsymbol{v}=\frac{N}{M}\mathbf{1}$	26.4 ± 1.1	23.6 ± 0.3	22.3 ± 1.3	40.8 ± 0.8	57.7 ± 10.9	$\mathbf{25.9} \pm 0.1$	79.5 ± 14.5	$\mathbf{33.9} \pm 0.1$	33.9 ± 1.4
	unif. IW, $(\star)$	$\mathbf{21.2} \pm 0.2$	$\mathbf{20.8} \pm 0.1$	$\mathbf{20.8} \pm 0.1$	$\mathbf{31.8} \pm 1.8$	$\mathbf{31.3} \pm 2.7$	$\mathbf{25.9} \pm 0.1$	48.1 ± 2.2	$\mathbf{33.9} \pm 0.1$	33.9 ± 0.1
	w/o IW, $(\star)$	29.8 ± 2.1	21.8 ± 0.4	21.2 ± 0.3	43.8 ± 5.5	35.8 ± 6.3	26.9 ± 0.5	74.9 ± 10.2	34.2 ± 0.1	34.0 ± 0.3
	$\boldsymbol{v}=\mathbf{1}$	69.4 ± 3.2	78.4 ± 4.9	70.4 ± 4.7	37.5 ± 24.4	34.7 ± 23.6	38.3 ± 24.8	57.7 ± 7.1	53.1 ± 3.6	49.4 ± 3.0
Sparse BBVI	Batch	65.0 ± 8.6	39.1 ± 1.6	35.9 ± 1.2	40.2 ± 2.9	35.2 ± 0.9	25.9 ± 2.1	70.8 ± 4.4	48.3 ± 1.6	33.9 ± 10.2
	Incremental	66.8 ± 4.2	65.8 ± 1.6	81.6 ± 15.1	33.6 ± 3.6	34.8 ± 3.1	29.4 ± 2.9	41.5 ± 0.6	42.8 ± 4.7	46.6 ± 2.5
PSVI [25]		81.9 ± 8.4	42.4 ± 0.6	37.2 ± 0.9	32.5 ± 3.1	51.2 ± 16.9	26.0 ± 0.1	40.7 ± 2.4	34.0 ± 0.1	34.0 ± 0.0
Sparse VI [8]		69.0 ± 0.0	62.9 ± 0.0	61.1 ± 0.0	63.1 ± 0.0	58.5 ± 0.0	58.6 ± 0.0	57.3 ± 0.0	56.4 ± 0.0	52.8 ± 0.0
Subset Laplace		62.0 ± 0.0	58.1 ± 0.0	57.7 ± 0.0	61.9 ± 0.0	46.9 ± 0.0	44.1 ± 0.0	66.5 ± 0.0	51.1 ± 0.0	48.8 ± 0.0
Rand. Coreset		175.1 ± 25.3	76.9 ± 15.5	83.1 ± 5.8	122.0 ± 13.4	156.0 ± 53.0	47.6 ± 3.4	251.4 ± 32.1	178.5 ± 39.5	51.8 ± 6.7
Full MFVI				$\mathbf{19.8} \pm 0.0$			$\mathbf{26.1} \pm 0.1$			$\mathbf{33.9} \pm 0.0$

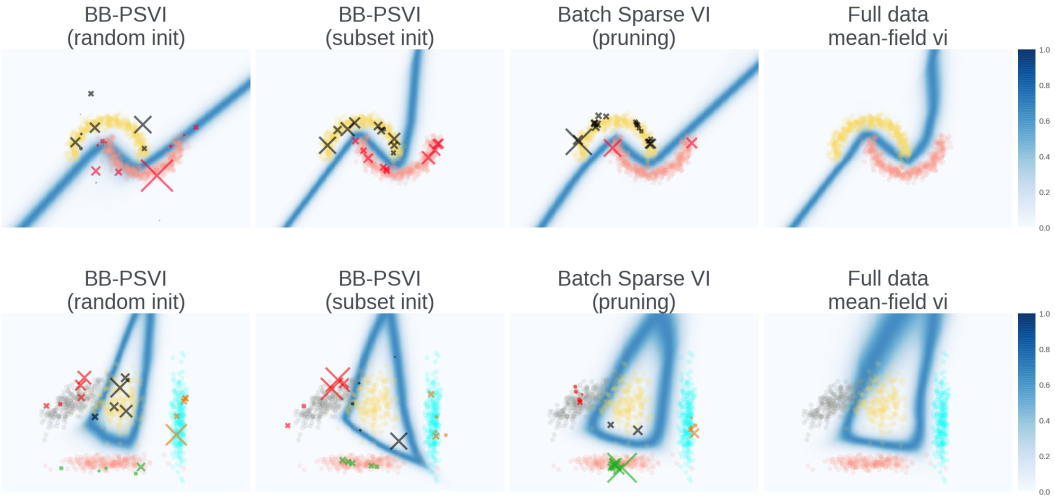

Figure 3: Predictive posterior entropy on simulated datasets. Coreset points are visualised in crosses, with size proportional to their weights.

belonging to a different class, and use a layer with 20 and 50 units respectively. To evaluate the representation ability of the pseudocoresets we consider two initialization schemes: we initialise the pseudo locations on a random subset equally split across categories, and a random initialization using a Gaussian centered on the means of the empirical distributions. In Fig. 3 we visualize the inferred predictive posterior entropy. Regardless of the initialization, BB PSVI has the capacity to optimize the locations of the coreset support according to the statistics of the true data distribution, yielding correct decision boundaries. Importantly, the coreset support consists itself of separable summarizing pseudo points, while the likelihood reweighting variational parameters enable emphasizing on critical regions of the landscape, assigning larger weights to pseudopoints lying close to the boundaries among the classes. Moreover, our pruning scheme which makes use of batch Sparse BBVI is able to gradually compress a large coreset of 250 datapoints to a compact performant subset of 20 datapoints arranged in critical locations for learning the Bayesian posterior.

MNIST In this part we assess the approximation quality of large-scale dataset compression for BNNs via coresets. We compare the predictive performance of black-box PSVI against standard mean-field VI, random coresets and frequentist methods relying on learnable synthetic data, namely dataset distillation w/ and w/o learnable soft labels [42, 37], and data condensation [44]. We can see that BB PSVI consistently outperforms the random coreset baseline for coreset sizes in the orders of few tens or hundreds, until the random coreset eventually meets its performance for sizes near 500 (Table 3). Moreover, at size 100 our construction provides better test accuracy compared to two frequentist approaches, enjoying the additional advantage of providing uncertainty estimates.

Table 3: Test accuracy on the MNIST dataset using NNs with the LeNet architecture for our black-box coreset constructions and other learning baselines.

	M	10	30	50	80	100	250	500
BB PSVI	$\boldsymbol{v}, \mathbf{u}, \mathbf{z}, \alpha$	$\mathbf{53.65} \pm 2.48$	$\mathbf{76.78} \pm 2.1$	$\mathbf{85.65} \pm 0.0$	89.43 ± 0.56	$\mathbf{91.28} \pm 0.74$	$\mathbf{93.85} \pm 0.0$	94.66 ± 0.25
	$\boldsymbol{v}, \mathbf{u}, \mathbf{z}$	52.97 ± 3.32	76.07 ± 2.06	85.05 ± 0.3	$\mathbf{89.87} \pm 0.95$	90.77 ± 0.92	93.51 ± 0.14	$\mathbf{94.73} \pm 0.12$
	$\boldsymbol{v}, \mathbf{u}, \alpha$	40.43 ± 0.66	63.15 ± 0.14	77.54 ± 1.32	84.52 ± 0.18	85.47 ± 0.21	90.98 ± 0.33	94.1 ± 0.27
	$\boldsymbol{v}, \mathbf{u}$	40.69 ± 0.84	62.8 ± 0.02	76.85 ± 0.86	84.52 ± 0.31	86.03 ± 0.6	91.3 ± 0.37	93.77 ± 0.24
Sparse VI [Batch]		43.7 ± 2.14	68.41 ± 0.6	77.21 ± 2.43	83.88 ± 0.44	85.65 ± 0.17	91.54 ± 0.16	94.37 ± 0.28
Dataset Condens. [44]		-	-	-	-	$\mathbf{93.9} \pm 0.6$	-	-
SLDD [37]		-	-	-	-	82.7 ± 2.8	-	-
Dataset Distill. [42]		-	-	-	-	79.5 ± 8.1	-	-
Random Coreset		42.17 ± 1.74	68.67 ± 3.21	78.75 ± 2.17	86.24 ± 0.25	87.44 ± 0.32	92.1 ± 0.65	95.11 ± 0.03
Full MFVI								99.13 ± 0.07

5 Related work

Coresets and data pruning beyond Bayesian inference. The paradigm of coresets emerged from seminal work in computational geometry [1, 16], and has since found broad use in data-intensive machine learning, including least mean square solvers [23], dimensionality reduction [17], clustering [15, 22, 2] and maximum likelihood estimation [3, 26]. Coreset ideas have also found recent applications in active learning [33, 6], continual learning [5, 4] and robustness [27]. Recent works in neural networks have investigated several heuristics for pruning datasets, including the misclassification events of datapoints over learning [40], or gradient scores in early stages of training [31].

Pseudodata for learning. Despite the multiple breakthroughs achieved by classical coreset methods, the common confine of selecting the representative data instances among the constituent datapoints might lead to large summary sizes in practice. Exempt from this limitation are methods that synthesize informative pseudo-examples for training that can summarize statistics lying away from the data manifold. These appear in the literature under various terms, including dataset distillation [24, 42, 21], condensation [44], dataset meta-learning [28] and inducing points [29]. Moreover, recent work [39] has demonstrated that incorporating learnable pseudo-inputs as hyperparameters of the model prior can enhance inference in variational auto-encoders. Learnable pseudodata that resemble the sufficient statistics of the original dataset have also been used in herding [11] and scalable GP inference methods [36, 38].

Variational inference, importance weighting, and Monte Carlo objectives. Importance weighting can provide tighter lower bounds of the data log-likelihood for training variational auto-encoders [7], resulting in richer latent representations [12]. In later work, [13] propose an importance weighting VI scheme for general purpose probabilistic inference. [35, 34] introduce variational programs as programmatic ways to specify a variational distribution.

6 Conclusion

We have developed a novel black-box construction for variational coresets that can be applied at scale to large datasets as well as richly structured models, and result in effective model-conditional dataset summarization. We overcome the intractability of previous inference algorithms with our proposed black box objectives, which we develop for both Sparse VI as well as PSVI, and propose a suite of algorithms utilizing them.

Through experiments on intractable models, we show that the inferred coresets form a useful summary of a dataset with respect to the learning problem at hand, and are capable to efficiently compress information for Bayesian posterior approximations. In future work, we plan to extend this inference toolbox for privacy purposes and non-traditional learning settings, e.g. continual learning.

References

[1] Pankaj K Agarwal, Sariel Har-Peled, Kasturi R Varadarajan, et al. Geometric approximation via coresets. *Combinatorial and computational geometry*, 2005.

[2] Olivier Bachem, Mario Lucic, and Andreas Krause. Coresets for nonparametric estimation - the case of DP-means. In *ICML*, 2015.

[3] Olivier Bachem, Mario Lucic, and Andreas Krause. Practical coreset constructions for machine learning. *arXiv preprint arXiv:1703.06476*, 2017.

[4] Lukas Balles, Giovanni Zappella, and Cédric Archambeau. Gradient-matching coresets for continual learning, 2021.

[5] Zalán Borsos, Mojmir Mutny, and Andreas Krause. Coresets via bilevel optimization for continual learning and streaming. In *NeurIPS*, 2020.

[6] Zalán Borsos, Marco Tagliasacchi, and Andreas Krause. Semi-supervised batch active learning via bilevel optimization. In *ICASSP*, 2021.

[7] Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. In *ICLR*, 2016.

[8] Trevor Campbell and Boyan Beronov. Sparse variational inference: Bayesian coresets from scratch. In *NeurIPS*, 2019.

[9] Trevor Campbell and Tamara Broderick. Bayesian coreset construction via greedy iterative geodesic ascent. In *ICML*, 2018.

[10] Trevor Campbell and Tamara Broderick. Automated scalable Bayesian inference via hilbert coresets. *JMLR*, 2019.

[11] Yutian Chen, Max Welling, and Alexander J. Smola. Super-samples from kernel herding. In *UAI*, 2010.

[12] Chris Cremer, Quaid Morris, and David Duvenaud. Reinterpreting importance-weighted autoencoders. In *ICLR*, 2017.

[13] Justin Domke and Daniel R. Sheldon. Importance weighting and variational inference. In *NeurIPS*, 2018.

[14] Dheeru Dua and Casey Graff. Uci machine learning repository [http://archive. ics. uci. edu/ml]. irvine, ca: University of california. *School of Information and Computer Science*, 2019.

[15] Dan Feldman, Matthew Faulkner, and Andreas Krause. Scalable training of mixture models via coresets. In *NeurIPS*, 2011.

[16] Dan Feldman and Michael Langberg. A unified framework for approximating and clustering data. In *STOC*, 2011.

[17] Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning big data into tiny data: Constant-size coresets for k-means, PCA, and projective clustering. *SICOMP*, 2020.

[18] Riccardo Grazzi, Luca Franceschi, Massimiliano Pontil, and Saverio Salzo. On the iteration complexity of hypergradient computation. In *ICML*, 2020.

[19] Jonathan Huggins, Trevor Campbell, and Tamara Broderick. Coresets for scalable Bayesian logistic regression. *NeurIPS*, 2016.

[20] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010.

[21] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *AISTATS*, 2020.

[22] Mario Lucic, Olivier Bachem, and Andreas Krause. Strong coresets for hard and soft Bregman clustering with applications to exponential family mixtures. In *AISTATS*, 2016.

[23] Alaa Maalouf, Ibrahim Jubran, and Dan Feldman. Fast and accurate least-mean-squares solvers. *NeurIPS*, 2019.

[24] Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based hyperparameter optimization through reversible learning. In *ICML*, 2015.

[25] Dionysis Manousakas, Zuheng Xu, Cecilia Mascolo, and Trevor Campbell. Bayesian pseudo-coresets. *NeurIPS*, 2020.

[26] Baharan Mirzasoleiman, Jeff Bilmes, and Jure Leskovec. Coresets for data-efficient training of machine learning models. In *ICML*, 2020.

[27] Baharan Mirzasoleiman, Kaidi Cao, and Jure Leskovec. Coresets for robust training of deep neural networks against noisy labels. *NeurIPS*, 2020.

[28] Timothy Nguyen, Zhourong Chen, and Jaehoon Lee. Dataset meta-learning from kernel ridge-regression. In *ICLR*, 2021.

[29] Timothy Nguyen, Roman Novak, Lechao Xiao, and Jaehoon Lee. Dataset distillation with infinitely wide convolutional networks. In *NeurIPS*, 2021.

[30] Art B. Owen. *Monte Carlo theory, methods and examples*. 2013.

[31] Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training. In *NeurIPS*, 2021.

[32] Fabian Pedregosa. Hyperparameter optimization with approximate gradient. In *ICML*, 2016.

[33] Robert Pinsler, Jonathan Gordon, Eric Nalisnick, and José Miguel Hernández-Lobato. Bayesian batch active learning as sparse subset approximation. In *NeurIPS*, 2019.

[34] Rajesh Ranganath, Dustin Tran, Jaan Altosaar, and David Blei. Operator variational inference. *NeurIPS*, 2016.

[35] Rajesh Ranganath, Dustin Tran, and David Blei. Hierarchical variational models. In *ICML*, 2016.

[36] Edward Snelson and Zoubin Ghahramani. Sparse Gaussian processes using pseudo-inputs. *NeurIPS*, 2006.

[37] Ilia Sucholutsky and Matthias Schonlau. 'less than one'-shot learning: Learning N classes from M < N samples. *arXiv preprint arXiv:2009.08449*, 2020.

[38] Michalis Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *AISTATS*, 2009.

[39] Jakub Tomczak and Max Welling. VAE with a VampPrior. In *AISTATS*, 2018.

[40] Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J. Gordon. An empirical study of example forgetting during deep neural network learning. In *ICLR*, 2019.

[41] De Wang, Danesh Irani, and Calton Pu. Evolutionary study of web spam: Webb spam corpus 2011 versus webb spam corpus 2006. In *ICCCNAW*, 2012.

[42] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018.

[43] Jacky Zhang, Rajiv Khanna, Anastasios Kyrillidis, and Sanmi Koyejo. Bayesian coresets: Revisiting the nonconvex optimization perspective. In *AISTATS*, 2021.

[44] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *ICLR*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]