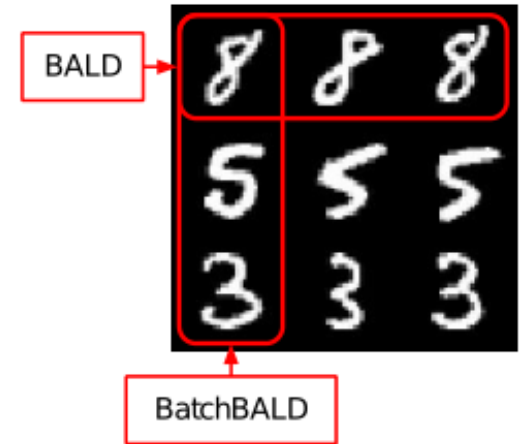


# BatchBALD: Efficient and Diverse Batch Acquisition for Deep Bayesian Active Learning

(NIPS 2019)

# Motivation of Batch Active Learning

- Reduce iteration times to improve data efficiency and reduce total run time
  - Re-training a deep model is computationally expensive
  - Mobilising an expert is time consuming
- Considering the joint and complementary relationship between sample points, diverse.
  - Naively acquiring the top k scoring points can lead to redundant acquisitions.



# BALD & Dropout variation inference

- Likelihood:  $p(y | \theta, x)$
- Bayesian: prior:  $p(\theta)$ , posterior:  $p(\theta | \mathcal{D})$

BALD: "Bayesian Active Learning by Disagreement"

$$I[\theta, y | \mathbf{x}, \mathcal{D}] = \arg \max_{\mathbf{x}} H[y | \mathbf{x}, \mathcal{D}] - \mathbb{E}_{\theta \sim p(\theta | \mathcal{D})} [H[y | \mathbf{x}, \theta]]$$

- First term captures general uncertainty of model
- Second term captures the uncertainty of a given draw of the model parameters

where,  $H(y | x, D) = \sum_c p(y = c | x, D) \log(p(y = c | x, D))$

$$p(y = c | x, D) = \int_{\theta \sim p(\theta | D)} p(y = c | x, D) \mathbf{p}(\theta | \mathbf{D}) d\theta \longrightarrow \left\{ \begin{array}{l} \text{Dropout as variation approximation} \\ \text{Kullback-Leibler (KL) divergence} \\ \text{MCMC} \end{array} \right.$$

# BatchBALD I - Introduction

→ Score batches of points **jointly**

Joint entropy: **Hard** to compute

$$a_{BatchBALD}(\{x_1, \dots, x_n\}, p(\omega | \mathcal{D}_{train})) = \mathbb{H}(y_1, \dots, y_n | x_1, \dots, x_b, \mathcal{D}_{train}) \\ - \mathbb{E}_{p(\omega | \mathcal{D}_{train})}[\mathbb{H}(y_1, \dots, y_n | x_1, \dots, x_b, \omega)].$$

Expectation over entropy: easy to compute

## BatchBALD II - MC estimator

$$\begin{aligned}\mathbb{H}(y_1, \dots, y_n) &= \mathbb{E}_{p(y_1, \dots, y_n)} [-\log p(y_1, \dots, y_n)] \\ &= \mathbb{E}_{p(\boldsymbol{\omega})} \mathbb{E}_{p(y_1, \dots, y_n | \boldsymbol{\omega})} [-\log \mathbb{E}_{p(\boldsymbol{\omega})} [p(y_1, \dots, y_n | \boldsymbol{\omega})]] \\ &\approx - \sum_{\hat{y}_{1:n}} \left( \frac{1}{k} \sum_{j=1}^k p(\hat{y}_{1:n} | \hat{\boldsymbol{\omega}}_j) \right) \log \left( \frac{1}{k} \sum_{j=1}^k p(\hat{y}_{1:n} | \hat{\boldsymbol{\omega}}_j) \right).\end{aligned}$$

- $k$  - the number of samples of the parameter distribution
- $\hat{y}$  - a configuration of a batch of labels
- Size of  $\hat{y}$  grows exponentially with  $|\mathcal{D}_{pool}|^b$

# BatchBALD III - Greedy Approximation

To avoid combinatorial explosion: we grow the acquisition batch one by one.

$$\frac{1}{k} \sum_{j=1}^k p(\hat{y}_{1:n} | \hat{\omega}_j) = \frac{1}{k} \sum_{j=1}^k p(\hat{y}_{1:n-1} | \hat{\omega}_j) p(\hat{y}_n | \hat{\omega}_j) = \left( \frac{1}{k} \hat{P}_{1:n-1} \hat{P}_n^T \right)_{\hat{y}_{1:n-1}, \hat{y}_n} .$$

---

**Algorithm 1:** Greedy BatchBALD 1 -  $1/e$ -approximate algorithm

---

**Input:** acquisition size  $b$ , unlabelled dataset  $\mathcal{D}_{\text{pool}}$ , model parameters  $p(\omega | \mathcal{D}_{\text{train}})$

1  $A_0 \leftarrow \emptyset$

2 **for**  $n \leftarrow 1$  **to**  $b$  **do**

3     **foreach**  $x \in \mathcal{D}_{\text{pool}} \setminus A_{n-1}$  **do**  $s_x \leftarrow a_{\text{BatchBALD}}(A_{n-1} \cup \{x\}, p(\omega | \mathcal{D}_{\text{train}}))$

4      $x_n \leftarrow \arg \max_{x \in \mathcal{D}_{\text{pool}} \setminus A_{n-1}} s_x$

5      $A_n \leftarrow A_{n-1} \cup \{x_n\}$

6 **end**

**Output:** acquisition batch  $A_n = \{x_1, \dots, x_b\}$

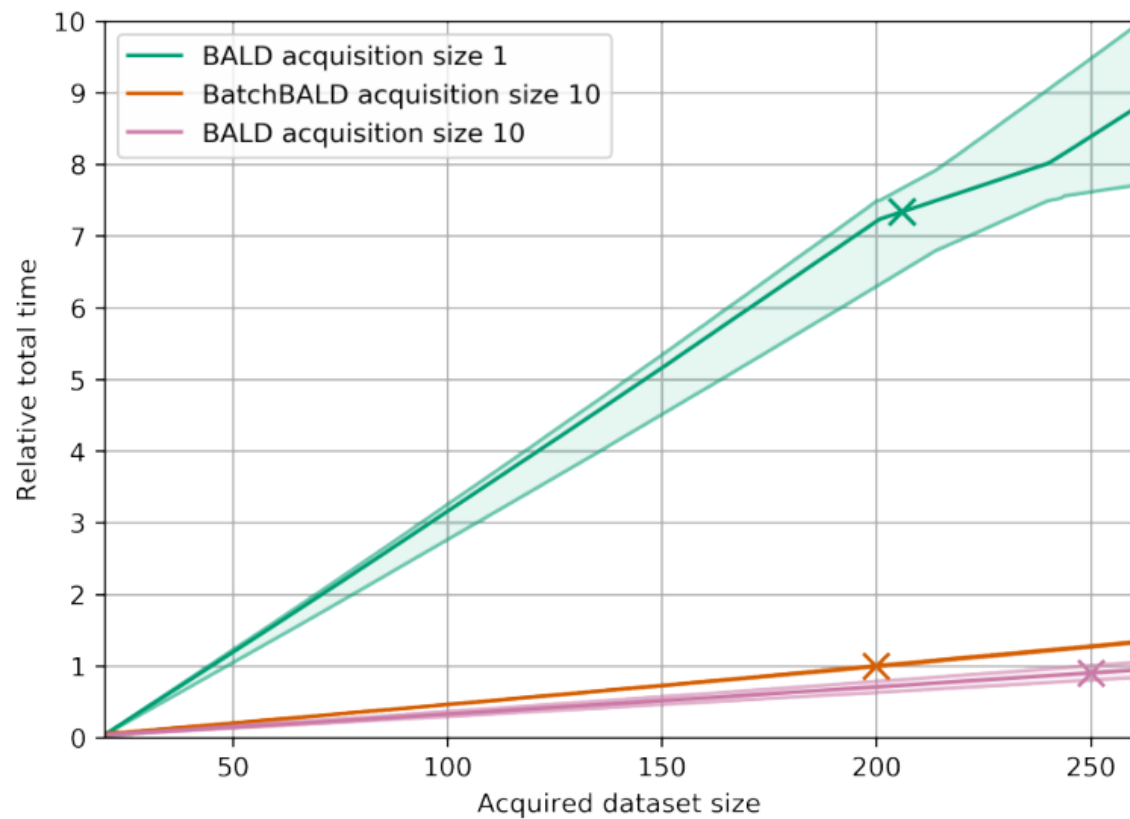

---

# Greedy Approximation and Submodularity



---

- We can compute the greedy joint entropy exact for the first 4 points, afterwards we use the Monte Carlo estimator
- We prove in the paper that the greedy approximation is a submodular function and therefore its error is bounded by  $1 - 1/e$

# Results MNIST




# Results EMNIST




# Diversity


BatchBALD



BALD



# Results CINIC-10 (CIFAR+ImageNet)



# Bayesian Batch Active Learning as Sparse Subset Approximation

(NIPS 2019)

# Related paper

---

- Active Learning For Convolutional Neural Networks: A Core-Set Approach, ICLR, 2018
- Automated Scalable Bayesian Inference via Hilbert Coresets, ICLR, 2019

# A Core-Set Approach

$$\begin{aligned}
 E_{\mathbf{x}, y \sim p_{\mathcal{Z}}} [l(\mathbf{x}, y; A_{\mathbf{s}})] &\leq \underbrace{E_{\mathbf{x}, y \sim p_{\mathcal{Z}}} [l(\mathbf{x}, y; A_{\mathbf{s}})] - \frac{1}{n} \sum_{i \in [n]} l(\mathbf{x}_i, y_i; A_{\mathbf{s}})}_{\text{Generalization Error}} + \underbrace{\frac{1}{|\mathbf{s}|} \sum_{j \in \mathbf{s}} l(\mathbf{x}_j, y_j; A_{\mathbf{s}})}_{\text{Training Error}} \\
 &+ \underbrace{\left[ \frac{1}{n} \sum_{i \in [n]} l(\mathbf{x}_i, y_i; A_{\mathbf{s}}) - \frac{1}{|\mathbf{s}|} \sum_{j \in \mathbf{s}} l(\mathbf{x}_j, y_j; A_{\mathbf{s}}) \right]}_{\text{Core-Set Loss}}
 \end{aligned}$$

$$\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$$

$$l(\cdot, \cdot; \mathbf{w}) : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{R}$$

$$\{\mathbf{x}_i, y_i\} \in [n] \sim p(\mathcal{Z})$$

$$\mathbf{s} = \{s(j) \in [n]\}_{j \in [m]}$$

$A_{\mathbf{s}}$ : which outputs a set of parameters  $\mathbf{w}$  given a labelled set  $\mathbf{s}$ .




Figure 1: **Visualization of the Theorem 1** Consider the set of selected points  $\mathbf{s}$  and the points in the remainder of the dataset  $[n] \setminus \mathbf{s}$ , our results shows that if  $\mathbf{s}$  is the  $\delta_{\mathbf{s}}$  cover of the dataset,  $\left| \frac{1}{n} \sum_{i \in [n]} l(\mathbf{x}_i, y_i, A_{\mathbf{s}}) - \frac{1}{|\mathbf{s}|} \sum_{j \in \mathbf{s}} l(\mathbf{x}_j, y_j, A_{\mathbf{s}}) \right| \leq \mathcal{O}(\delta_{\mathbf{s}}) + \mathcal{O}\left(\sqrt{\frac{1}{n}}\right)$

# Bayesian Coreset

$$p(\boldsymbol{\theta} | \mathcal{D}_0 \cup (\mathcal{X}_p, \mathcal{Y}_p)) = \frac{p(\boldsymbol{\theta} | \mathcal{D}_0) p(\mathcal{Y}_p | \mathcal{X}_p, \boldsymbol{\theta})}{p(\mathcal{Y}_p | \mathcal{X}_p, \mathcal{D}_0)} \quad (1)$$

$p(y|x, \theta)$  parameterized by  $\theta \in \Theta$

$\mathcal{D}_0 = \{x_n, y_n\}_{n=1}^N$ : labeled dataset

$\mathcal{X}_p = \{x_m\}_{m=1}^M$ : unlabeled pool set

$$\begin{aligned} \mathbb{E}_{\mathcal{Y}_p} [\log p(\boldsymbol{\theta} | \mathcal{D}_0 \cup (\mathcal{X}_p, \mathcal{Y}_p))] &= \mathbb{E}_{\mathcal{Y}_p} [\log p(\boldsymbol{\theta} | \mathcal{D}_0) + \log p(\mathcal{Y}_p | \mathcal{X}_p, \boldsymbol{\theta}) - \log p(\mathcal{Y}_p | \mathcal{X}_p, \mathcal{D}_0)] \\ &= \log p(\boldsymbol{\theta} | \mathcal{D}_0) + \mathbb{E}_{\mathcal{Y}_p} [\log p(\mathcal{Y}_p | \mathcal{X}_p, \boldsymbol{\theta})] + \mathbb{H}[\mathcal{Y}_p | \mathcal{X}_p, \mathcal{D}_0] \\ &= \log p(\boldsymbol{\theta} | \mathcal{D}_0) + \sum_{m=1}^M \underbrace{\left( \mathbb{E}_{\mathbf{y}_m} [\log p(\mathbf{y}_m | \mathbf{x}_m, \boldsymbol{\theta})] + \mathbb{H}[\mathbf{y}_m | \mathbf{x}_m, \mathcal{D}_0] \right)}_{\mathcal{L}_m(\boldsymbol{\theta})} \end{aligned} \quad (2)$$

$$L(\boldsymbol{\theta}) = \sum_{m=1}^M L_m(\boldsymbol{\theta})$$

$$L(w, \boldsymbol{\theta}) := \sum_{m=1}^M w_m * L_m(\boldsymbol{\theta}) \text{ satisfies } |L(w, \boldsymbol{\theta}) - L(\boldsymbol{\theta})| < \varepsilon |L(\boldsymbol{\theta})|, \forall \boldsymbol{\theta} \in \Theta, w_m = 0 \text{ or } 1 \quad (3)$$

# Uniform Bayesian Coreset

$$L(w, \theta) := \sum_{m=1}^M w_m * L_m(\theta) \text{ satisfies } |L(w, \theta) - L(\theta)| < \epsilon |L(\theta)|, \forall \theta \in \Theta, w_m = 0 \text{ or } 1 \quad (3)$$

The algorithm proposed by Huggins et al.(2016) to construct a Bayesian coreset is as follows. First, compute the **sensitivity**  $\sigma_m$  of each data point,

$$\sigma_m := \sup_{\theta \in \Theta} \left| \frac{L_m(\theta)}{L(\theta)} \right|$$

and then subsample the dataset by taking  $b$  independent draws with probability proportional to  $\sigma_m$  (resulting in a coreset of size  $\leq b$ ) via

$$\sigma := \sum_{m=1}^M \sigma_m \quad (b_1, \dots, b_M) \sim \text{Multi} \left( b, \left( \frac{\sigma_m}{\sigma} \right)_{m=1}^M \right) \quad W_m = \frac{\sigma}{\sigma_m} \frac{b_m}{b}$$

Since  $E[W_m] = 1$ , we have that  $E[L(W, \theta)] = L(\theta)$ , and we expect that  $L(W, \theta) \rightarrow L(\theta)$  in some sense as  $b$  increases,  $\epsilon^2 = O(\frac{1}{b})$

# Coreset as sparse vectors

Given a function space  $g : \Theta \rightarrow \mathbb{R}$  with bounded uniform norm weighted by  $\mathcal{L}(\theta)$ :

$$\|g\| := \sup_{\theta \in \Theta} \left| \frac{g(\theta)}{\mathcal{L}(\theta)} \right|.$$

we view  $\mathcal{L}_m : \Theta \mapsto \mathbb{R}$  and  $\mathcal{L} = \sum_m \mathcal{L}_m$  as vectors in this function space.

we convert the problem of constructing a batch to a sparse subset approximation problem, i.e.

$$\mathbf{w}^* = \underset{\mathbf{w}}{\text{minimize}} \|\mathcal{L} - \mathcal{L}(\mathbf{w})\|^2 \quad \text{subject to} \quad w_m \in \{0, 1\} \quad \forall m, \quad \sum_m \mathbb{1}_m \leq b. \quad (4)$$

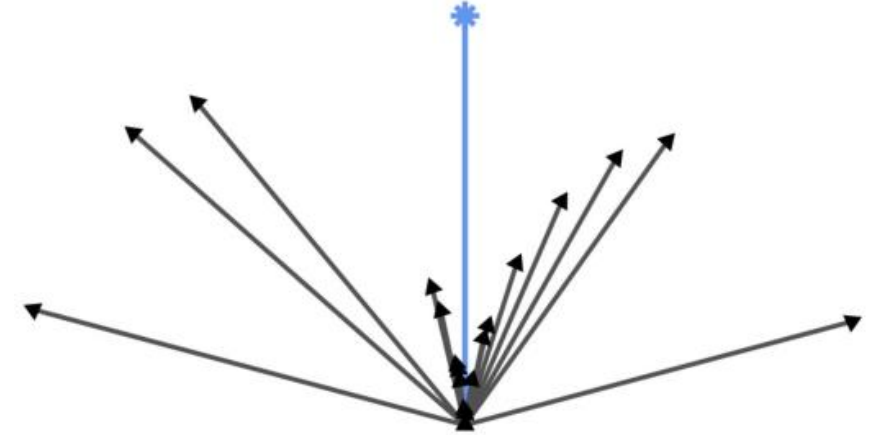
# Inner product and Hilbert Coreset

- The uniform norm lacks a sense of “directionality” as it does not correspond to an inner-product.
- Constructing Bayesian coresets in a Hilbert space by inner product  $\langle \mathcal{L}_m, \mathcal{L}_n \rangle$

we relax the binary weight constraint to be non-negative and replace the cardinality constraint with a polytope constraint.

Let  $\sigma_m = \|\mathcal{L}_m\|$ ,  $\sigma = \sum_{m=1}^M \sigma_m$ , and  $K \in \mathbb{R}^{M \times M}$  be a kernel matrix with  $K_{m,n} = \langle \mathcal{L}_m, \mathcal{L}_n \rangle$ . The related optimization problem is:

$$\underset{\mathbf{w}}{\text{minimize}} \quad (\mathbf{1} - \mathbf{w})^T \mathbf{K} (\mathbf{1} - \mathbf{w}) \quad \text{subject to} \quad w_m \geq 0 \quad \forall m, \quad \sum_m w_m \sigma_m = \sigma, \quad (5)$$



# Inner product and Hilbert Coreset

where we used  $\|\mathcal{L} - \mathcal{L}(\mathbf{w})\|^2 = (\mathbf{1} - \mathbf{w})^T \mathbf{K} (\mathbf{1} - \mathbf{w})$ .

eg. Let  $M = 2$ ,

$$L = L_1 + L_2, \mathbf{w} = (w_1, w_2)^T, L(\mathbf{w}) = w_1 * L_1 + w_2 * L_2, \quad K = \begin{bmatrix} K_{11} & K_{12} \\ K_{21} & K_{22} \end{bmatrix}$$

$$\begin{aligned} (\mathbf{1} - \mathbf{w})^T K (\mathbf{1} - \mathbf{w}) &= (1 - w_1, 1 - w_2) \begin{bmatrix} K_{11} & K_{12} \\ K_{21} & K_{22} \end{bmatrix} (1 - w_1, 1 - w_2)^T \\ &= (1 - w_1)^2 K_{11} + (1 - w_1)(1 - w_2)(K_{21} + K_{12}) + (1 - w_2)^2 K_{22} \end{aligned}$$

$$\begin{aligned} \|L - L(\mathbf{w})\|^2 &= \langle L - L(\mathbf{w}), L - L(\mathbf{w}) \rangle \\ &= \langle (\mathbf{1} - \mathbf{w}_1)L_1 + (\mathbf{1} - \mathbf{w}_2)L_2, (\mathbf{1} - \mathbf{w}_1)L_1 + (\mathbf{1} - \mathbf{w}_2)L_2 \rangle \\ &= (1 - w_1)^2 \langle L_1, L_1 \rangle + (1 - w_1)(1 - w_2)(\langle L_2, L_1 \rangle + \langle L_1, L_2 \rangle) + (1 - w_2)^2 \langle L_2, L_2 \rangle \end{aligned}$$


# Frank-Wolfe optimization

## Algorithm 1 Active Bayesian Coresets with Frank-Wolfe Optimization

```

1: procedure ACS-FW( $b, \{\mathcal{L}_n\}_{n=1}^N, \langle \cdot, \cdot \rangle$ )
2:    $\sigma_n \leftarrow \sqrt{\langle \mathcal{L}_n, \mathcal{L}_n \rangle} \quad \forall n$  ▷ Compute norms
3:    $\sigma \leftarrow \sum_n \sigma_n$ 
4:    $\mathbf{w} \leftarrow \mathbf{0}$  ▷ Initialize weights to 0
5:   for  $t \in 1, \dots, b$  do
6:      $f \leftarrow \arg \max_{n \in N} \left[ \left\langle \mathcal{L} - \mathcal{L}(\mathbf{w}), \frac{1}{\sigma_n} \mathcal{L}_n \right\rangle \right]$  ▷ Greedily select point  $f$ 
7:      $\gamma \leftarrow \frac{\left[ \left\langle \frac{\sigma}{\sigma_f} \mathcal{L}_f - \mathcal{L}(\mathbf{w}), \mathcal{L} - \mathcal{L}(\mathbf{w}) \right\rangle \right]}{\left[ \left\langle \frac{\sigma}{\sigma_f} \mathcal{L}_f - \mathcal{L}(\mathbf{w}), \frac{\sigma}{\sigma_f} \mathcal{L}_f - \mathcal{L}(\mathbf{w}) \right\rangle \right]}$  ▷ Perform line search for step-size  $\gamma$ 
8:      $\mathbf{w} \leftarrow (1 - \gamma)\mathbf{w} + \gamma \frac{\sigma}{\sigma_f} \mathbf{1}_f$  ▷ Update weight for newly selected point
9:   end for
10:  return  $\mathbf{w}$ 
11: end procedure

```



$$6: \text{ Let } \mathbf{g} = L - L(\mathbf{w}), \quad \mathbf{d} = \frac{\sigma}{\sigma_f} L_f - L(\mathbf{w})$$

$$7: \gamma = \frac{\langle \mathbf{g}, \mathbf{d} \rangle}{\langle \mathbf{d}, \mathbf{d} \rangle} = \frac{\|\mathbf{g}\| \cos \theta}{\|\mathbf{d}\|}$$

$$8: (1 - \gamma)\mathbf{w} + \gamma \frac{\sigma}{\sigma_f} \mathbf{1}_f = \mathbf{w} + \gamma \left( \mathbf{w} - \frac{\sigma}{\sigma_f} \mathbf{1}_f \right)$$

$$8: L[\mathbf{w}] = L(\mathbf{w}) + \gamma \left( L(\mathbf{w}) - \frac{\sigma}{\sigma_f} L_f \right) = L(\mathbf{w}) + \gamma \mathbf{d}$$



$$\bar{\eta}^2 := \max_{n, m \in [N]} \left\| \frac{\mathcal{L}_n}{\sigma_n} - \frac{\mathcal{L}_m}{\sigma_m} \right\|^2,$$

$$\|L(\mathbf{w}) - L\| \leq \frac{2\sigma\bar{\eta}}{\sqrt{3M+1}}.$$

# Choice of inner products

- weighted Fisher inner product:

$$\langle \mathcal{L}_n, \mathcal{L}_m \rangle_{\hat{\pi}, \mathcal{F}} = \mathbb{E}_{\hat{\pi}} \left[ \nabla_{\theta} \mathcal{L}_n(\theta)^T \nabla_{\theta} \mathcal{L}_m(\theta) \right],$$

$$\|\mathcal{L}(w) - \mathcal{L}\|_{\hat{\pi}, \mathcal{F}}^2 := \mathbb{E}_{\hat{\pi}} \left[ \|\nabla \mathcal{L}(\theta) - \nabla \mathcal{L}(w, \theta)\|_2^2 \right],$$

where  $\hat{\pi}$  to be the current posterior  $p(\theta | \mathcal{D}_0)$ .

- weighted Euclidean inner product:

$$\langle \mathcal{L}_n, \mathcal{L}_m \rangle_{\hat{\pi}, 2} = \mathbb{E}_{\hat{\pi}} \left[ \mathcal{L}_n(\theta) \mathcal{L}_m(\theta) \right].$$

$$\|\mathcal{L}(w) - \mathcal{L}\|_{\hat{\pi}, 2} := \mathbb{E}_{\hat{\pi}} \left[ (\mathcal{L}(\theta) - \mathcal{L}(w, \theta))^2 \right]$$

only requires tractable likelihood computations.

# Random projections for non-linear models

- Is limited to models for which the inner product can be evaluated in closed form.
- Requires  $\mathcal{O}|\mathcal{P}^2|$  computations
- Given random projections:

$$\hat{\mathcal{L}}_n = \frac{1}{\sqrt{J}} [\mathcal{L}_n(\boldsymbol{\theta}_1), \dots, \mathcal{L}_n(\boldsymbol{\theta}_J)]^T, \quad \boldsymbol{\theta}_j \sim \hat{\pi},$$

$$\langle \mathcal{L}_n, \mathcal{L}_m \rangle_{\hat{\pi}, 2} \approx \hat{\mathcal{L}}_n^T \hat{\mathcal{L}}_m,$$

---

## Algorithm 2 ACS-FW with Random Projections (for Weighted Euclidean Inner Product)

---

- 1: **procedure** ACS-FW( $b, J$ )
  - 2:      $\boldsymbol{\theta}_j \sim \hat{\pi} \quad j = 1, \dots, J$  ▷ Sample parameters
  - 3:      $\hat{\mathcal{L}}_n = \frac{1}{\sqrt{J}} [\mathcal{L}_n(\boldsymbol{\theta}_1), \dots, \mathcal{L}_n(\boldsymbol{\theta}_J)]^T \quad \forall n$  ▷ Compute random feature projections
  - 4:     **return** ACS-FW( $b, \{\hat{\mathcal{L}}_n\}_{n=1}^N, (\cdot)^T (\cdot)$ ) ▷ Call Algorithm 1 using projections
  - 5: **end procedure**
-

# Experiments

Avoid correlated queries:





Figure 1: Batch construction of different AL methods on *cifar10*, shown as a t-SNE projection [12]. Given 5000 labeled points (colored by class), a batch of 200 points (black crosses) is queried.




# Experiments


Competitive with greedy methods in the small-data:

Table 1: Final test RMSE on UCI regression datasets averaged over 40 (*year*: 5) seeds. MAXENT-I and MAXENT-SG require order(s) of magnitudes more model updates and are thus not directly comparable.

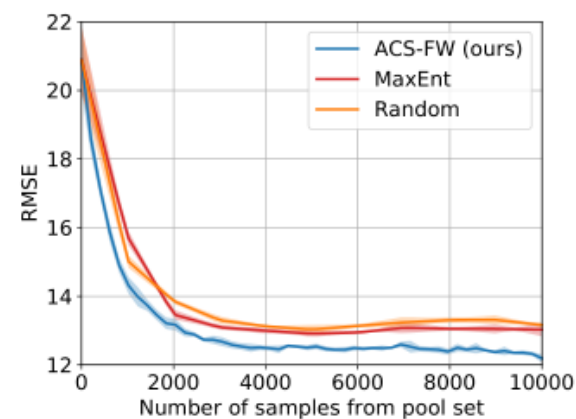
	N	d	RANDOM	MAXENT	ACS-FW	MAXENT-I	MAXENT-SG
yacht	308	6	1.272±0.0593	<b>0.923±0.0319</b>	1.031±0.0438	0.865±0.0276	0.971±0.0350
boston	506	13	4.068±0.0852	<b>3.640±0.0652</b>	3.799±0.0858	3.467±0.0676	3.458±0.0682
energy	768	8	0.959±0.0337	1.443±0.0857	<b>0.855±0.0259</b>	0.927±0.0461	1.055±0.0740
power	9568	4	5.108±0.0468	<b>5.022±0.0428</b>	<b>4.984±0.0366</b>	4.834±0.0313	4.855±0.0339
year	515 345	90	13.165±0.0307	13.030±0.0975	<b>12.194±0.0596</b>	N/A	N/A



(a) yacht



(b) energy




(c) year


Figure 3: Test RMSE on UCI regression datasets averaged over 40 (a-b) and 5 (c) seeds during AL. Error bars denote two standard errors.

# Experiments

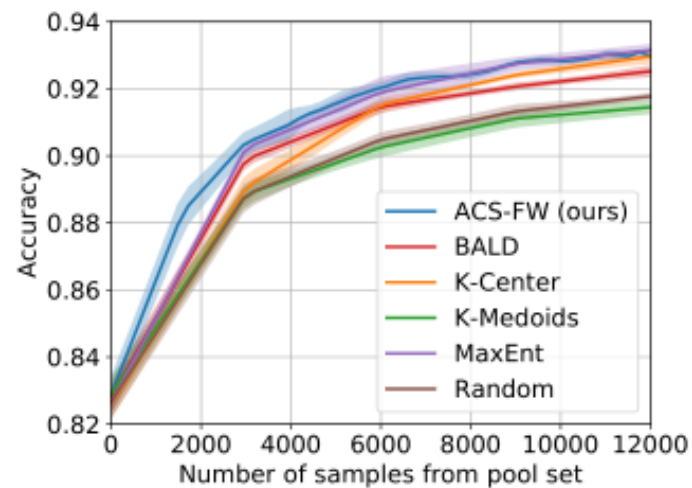
Scale to large datasets and models:



(a) cifar10



(b) SVHN



(c) Fashion MNIST

Figure 4: Test accuracy on classification tasks over 5 seeds. Error bars denote two standard errors.

# Experiments

## Runtime Evaluation:

Table 2: Runtime in seconds on UCI regression datasets averaged over 40 (*year*: 5) seeds. We report mean batch construction time (BT/it.) and total time (TT/it.) per AL iteration, as well as total cumulative time (total). MAXENT-I requires order(s) of magnitudes more model updates and is thus not directly comparable.

	RANDOM			MAXENT			ACS-FW			MAXENT-I		
	BT/it.	TT/it.	total	BT/it.	TT/it.	total	BT/it.	TT/it.	total	BT/it.	TT/it.	total
yacht	0.0	8.9	88.6	1.3	10.2	101.7	0.0	9.1	107.2	12.3	105.7	1057.4
boston	0.0	12.4	123.6	2.4	14.5	144.8	0.1	12.4	132.7	23.5	157.9	1578.6
energy	0.0	12.1	121.4	3.9	16.0	159.6	0.1	12.6	137.8	37.5	170.5	1704.9
power	0.4	9.4	94.0	53.0	61.7	617.0	0.8	10.2	179.8	517.3	609.1	6090.7
year	30.2	381.2	3811.6	3391.5	3746.5	37464.6	53.0	463.8	28475.2	N/A	N/A	N/A

**Thanks**

---