Efficient nonmyopic batch active search

ICML 2017

Introduction

T ~ the number of queries we can provide the oracle

$\Sigma_{i=1}^{T} y_i$ ~ the number of targets identified

Previous work developed policies for active search by appealing to Bayesian decision theory. This policy in the general case requires exponential computation.

To overcome this intractability, the authors of that work proposed using myopic lookahead policies in practice, which compute the optimal policy only up to a limited number of steps into the future.

The Optimal Policy

$$\mathcal{X} \triangleq \{x_i\} \qquad y \triangleq \mathbb{1}\{x \in \mathcal{R}\}$$

$$\mathcal{D} \triangleq \{(x_i, y_i)\} \qquad u(\mathcal{D}) \triangleq \sum_{y_i \in \mathcal{D}} y_i$$

$$\mathcal{D}_i \triangleq \{(x_j, y_j)\}_{j=1}^{i} \qquad \Pr(y = 1 \mid x, \mathcal{D})$$

Bayesian decision theory:

$$x_i^* = \operatorname*{arg\,max}_{x_i \in \mathcal{X} \setminus \mathcal{D}_{i-1}} \mathbb{E}\big[u(\mathcal{D}_t) \mid x_i, \mathcal{D}_{i-1}\big]$$

The Optimal Policy

i = t:

$$\mathbb{E}\left[u(\mathcal{D}_t) \mid x_t, \mathcal{D}_{t-1}\right] = \sum\nolimits_{y_t} u(\mathcal{D}_t)\,\Pr(y_t \mid x_t, \mathcal{D}_{t-1})$$
$$= u(\mathcal{D}_{t-1}) + \Pr(y_t = 1 \mid x_t, \mathcal{D}_{t-1}).$$

i = t - 1:

$$\mathbb{E}\left[u(\mathcal{D}_t) \mid x_{t-1}, \mathcal{D}_{t-2}\right] = u(\mathcal{D}_{t-2}) +$$
$$\Pr(y_{t-1} = 1 \mid x_{t-1}, \mathcal{D}_{t-2}) +$$
$$\mathbb{E}_{y_{t-1}}\left[\max_{x_t} \Pr(y_t = 1 \mid x_t, \mathcal{D}_{t-1})\right]$$

The Optimal Policy

i < t:

$$\mathbb{E}\left[u(\mathcal{D}_t) \mid x_i, \mathcal{D}_{i-1}\right] = u(\mathcal{D}_{i-1}) + \underbrace{\Pr(y_i = 1 \mid x_i, \mathcal{D}_{i-1}) +}_{\text{exploitation, } < 1} \underbrace{\mathbb{E}_{y_i}\left[\max_{x'} \mathbb{E}\left[u(\mathcal{D}_t \setminus \mathcal{D}_i) \mid x', \mathcal{D}_i\right]\right]}_{\text{exploration, } < t-i}$$

$$\mathcal{O}\big((2n)^\ell\big)$$

$$\ell = t - i + 1$$

Efficient Nonmyopic Active Search

If we assume that after observing D_i, the labels of all remaining unlabeled points are conditionally independent, then this approximation recovers the Bayesian optimal policy exactly.

$$\max_{x'} \mathbb{E}\big[u(\mathcal{D}_t \setminus \mathcal{D}_i) \mid x', \mathcal{D}_i\big] \approx {\textstyle\sum'_{t-i}} \Pr(y = 1 \mid x, \mathcal{D}_i)$$

$$\begin{aligned}\mathbb{E}\big[u(\mathcal{D}_t) \mid x_i, \mathcal{D}_{i-1}\big] \approx u(\mathcal{D}_{i-1}) &+ \\ \Pr(y_i = 1 \mid x_i, \mathcal{D}_{i-1}) &+ \\ \underbrace{\mathbb{E}_{y_i}\Big[{\textstyle\sum'_{t-i}} \Pr\big(y = 1 \mid x, \mathcal{D}_i\big)\Big]}_{\text{exploration, } < t-i}\end{aligned}$$

$$\mathcal{O}(n^2 \log n)$$

Experiments

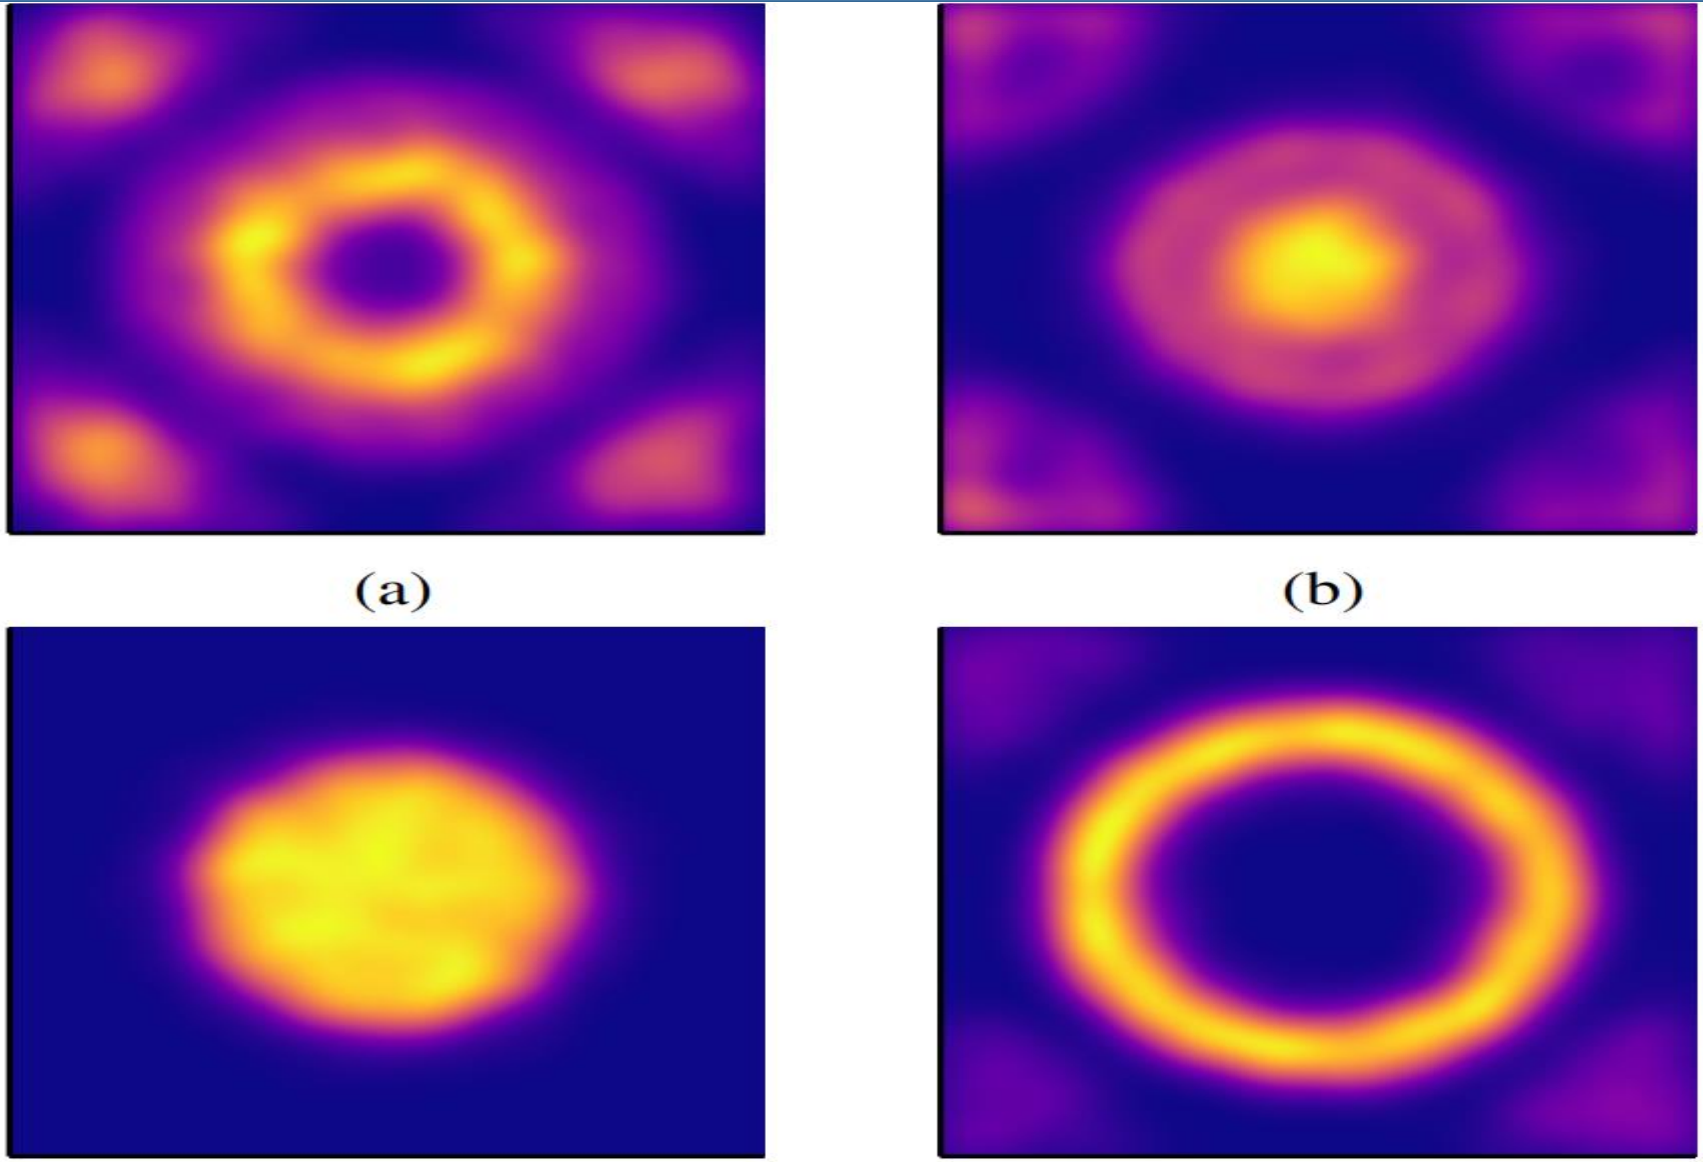

(a) (b) (c) (d)

Experiments

39788 computer science papers published in the top-50 most popular computer science venues

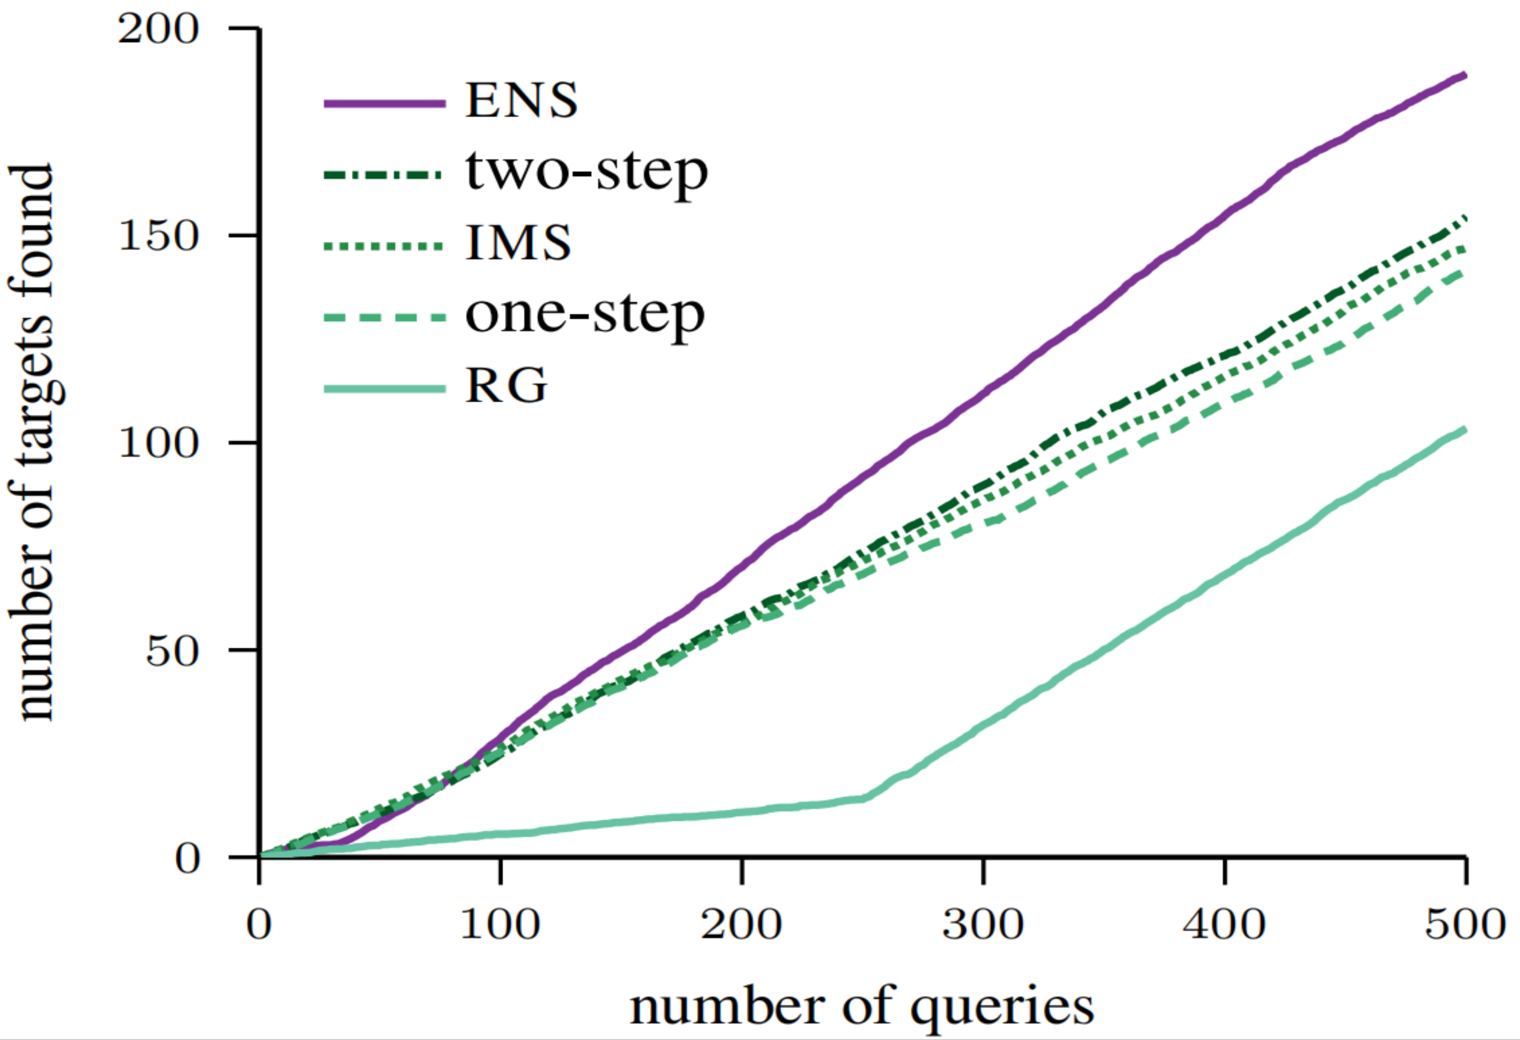

Experiments

CiteSeer^x data

policy	query number 100	300	500	700	900
RG	19.7	60.0	104	140	176
IMS	26.3	86.3	147	214	281
one-step	25.5	80.5	141	209	273
two-step	24.9	89.8	155	220	287
ENS–900	25.9	94.3	163	239	**308**
ENS–700	28.0	105	188	**259**	
ENS–500	28.7	**112**	**189**		
ENS–300	26.4	105			
ENS–100	**30.7**				

BMG data

policy	query number 100	300	500	700	900
RG	48.6	144	243	336	427
IMS	93.6	276	451	629	799
one-step	90.8	273	450	633	798
two-step	91.0	273	452	632	802
ENS–900	89.0	270	453	635	**815**
ENS–700	91.3	276	460	**645**	
ENS–500	92.4	279	**466**		
ENS–300	92.8	**279**			
ENS–100	**94.5**				

Efficient nonmyopic batch active search

NIPS 2018

$$\mathbb{E}[u(\mathcal{D}_t \setminus \mathcal{D}_i) \mid X, \mathcal{D}_i] = \sum_{x \in X} \Pr(y = 1 \mid x, \mathcal{D}_i) + \\ \mathbb{E}_{Y \mid X, \mathcal{D}_i} \left[\max_{X' : |X'| = T - b - |\mathcal{D}_i|} \mathbb{E} \left[u(Y') \mid X', \mathcal{D}_i, X, Y \right] \right]$$

	1	5	10	15	20	25	50	75	100
UGB	-	257.6	257.9	258.3	250.1	246.0	218.8	206.2	172.1
greedy	269.8	268.1	264.1	261.6	258.2	257.0	240.1	227.2	208.2
ss-one-1	269.8	260.7	254.6	245.2	233.6	223.4	200.8	182.9	178.9
ss-one-m	269.8	264.5	257.7	250.0	244.4	236.5	211.7	195.4	179.4
ss-one-s	269.8	266.8	261.3	256.7	248.7	244.1	214.9	202.4	181.3
ss-one-0	269.8	268.1	264.1	261.6	258.2	257.0	240.1	227.2	208.2
ss-two-1	281.1	237.1	219.8	210.8	212.1	196.2	172.1	158.8	152.9
ss-two-m	281.1	252.6	246.4	237.2	232.9	225.1	200.2	181.6	167.2
ss-two-s	281.1	248.9	242.5	235.3	226.6	219.2	196.7	175.3	158.3
ss-two-0	281.1	252.5	247.6	247.9	244.4	240.4	225.6	213.8	199.1
ss-ENS-1	**295.1**	269.4	247.9	227.2	223.1	210.3	185.3	152.6	148.7
ss-ENS-m	*295.1*	293.8	290.2	285.3	281.6	274.4	249.4	217.2	203.1
ss-ENS-s	*295.1*	289.9	278.3	269.8	262.6	255.0	220.8	185.5	161.2
ss-ENS-0	*295.1*	293.6	289.1	288.1	*287.5*	280.7	269.2	257.2	241.0
batch-ENS-16	*295.1*	**300.8**	**296.2**	293.9	**292.1**	*288.0*	275.8	*272.3*	252.9
batch-ENS-32	*295.1*	*300.8*	*295.5*	**297.9**	*290.6*	**288.8**	**281.4**	**275.5**	**263.5**

Multi-Class Active Learning by Uncertainty Sampling with Diversity Maximization

IJCV 2015

$$\mathrm{H}(\ell_i) = -\sum_{j}^{c} P(\mathcal{C}_j|x_i) log \left(P(\mathcal{C}_j|x_i)\right)$$

$$\max_{\sum_i f_i=1,\, f_i\geq 0} \sum_{x_i\in\mathcal{P}} -f_i \times \sum_{j}^{c} P(C_j|x_i) log \left(P(C_j|x_i)\right) - \Omega(f_i)$$

The term $\Omega(f_i)$ is a function on f encoding the data distribution information, in other words, the diversity criterion in decision making.

$$\min_{f_i} \Omega(f_i) = \min_{f_i} \sum_{i=1}^{n} \sum_{j=1}^{n} f_i f_j K_{ij}$$

$$K_{ij} = \frac{-\|x_i - x_j\|^2}{\sigma^2}$$

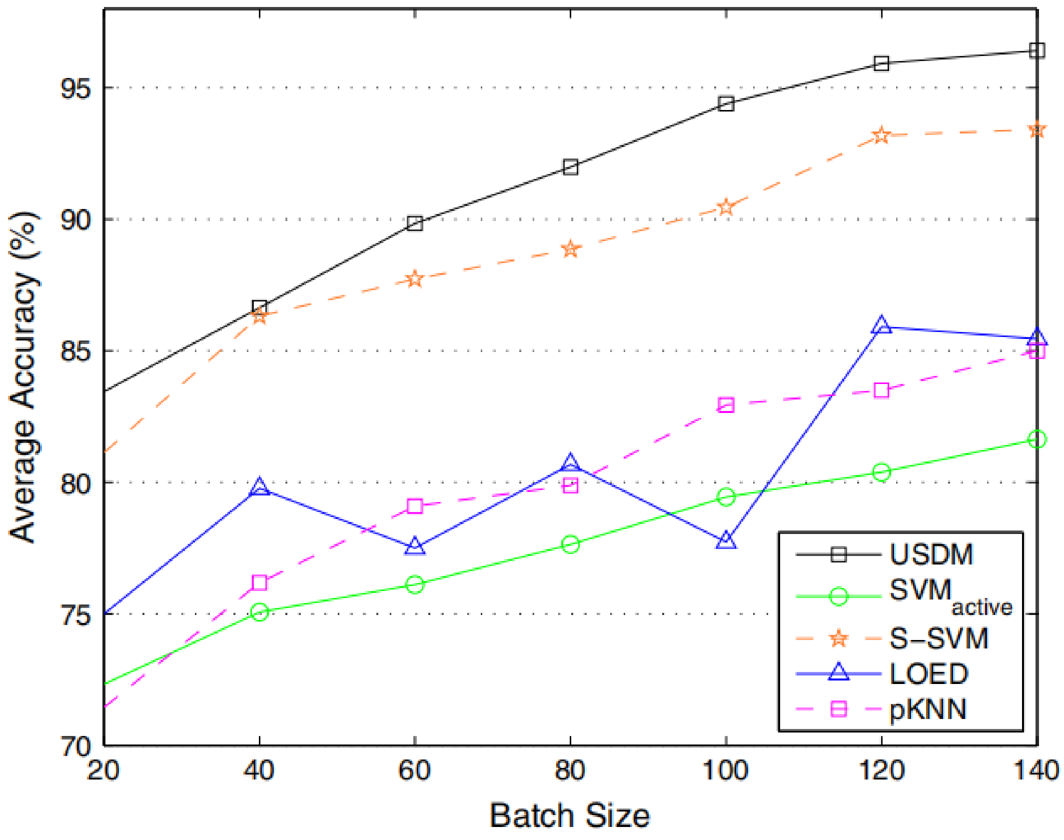

Average Accuracy (%)
Batch Size
USDM
SVM active
S−SVM
LOED
pKNN

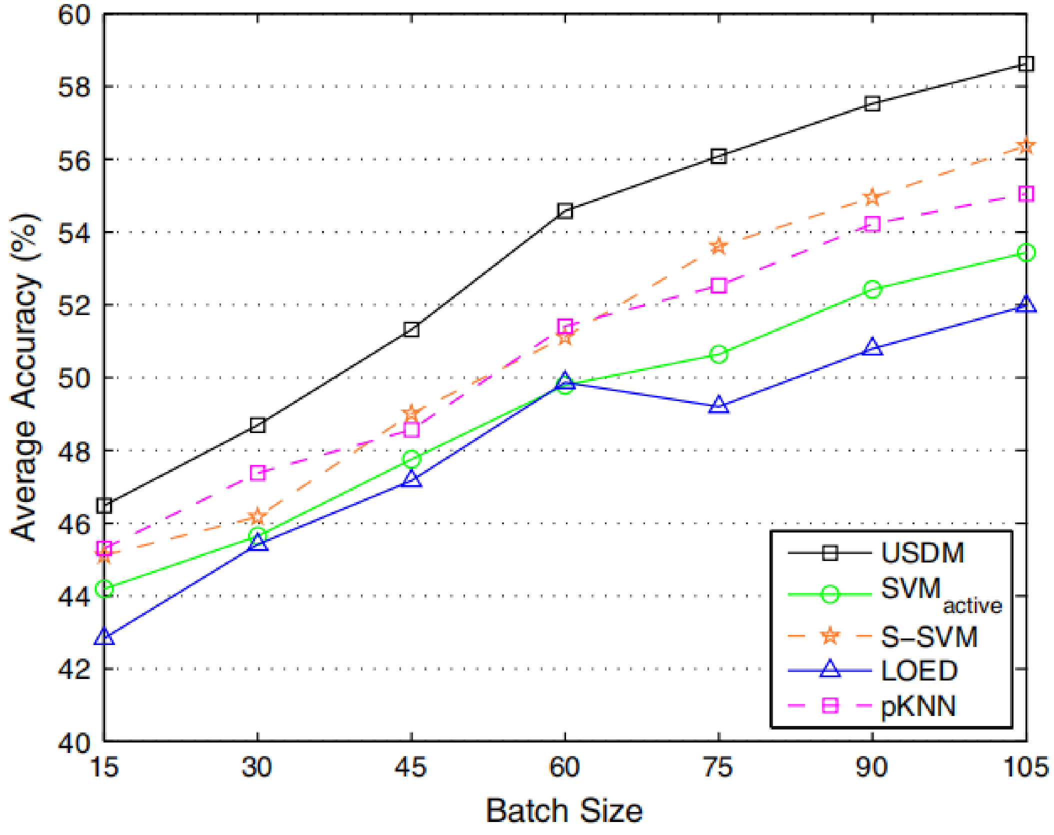

Average Accuracy (%)
Batch Size
USDM
SVM active
S−SVM
LOED
pKNN