

The 20th ACM SIGKDD International Conference
on Knowledge Discovery and Data Mining

Large Margin Distribution Machine

Teng Zhang and Zhi-Hua Zhou

LAMDA Group

National Key Laboratory for Novel Software Technology

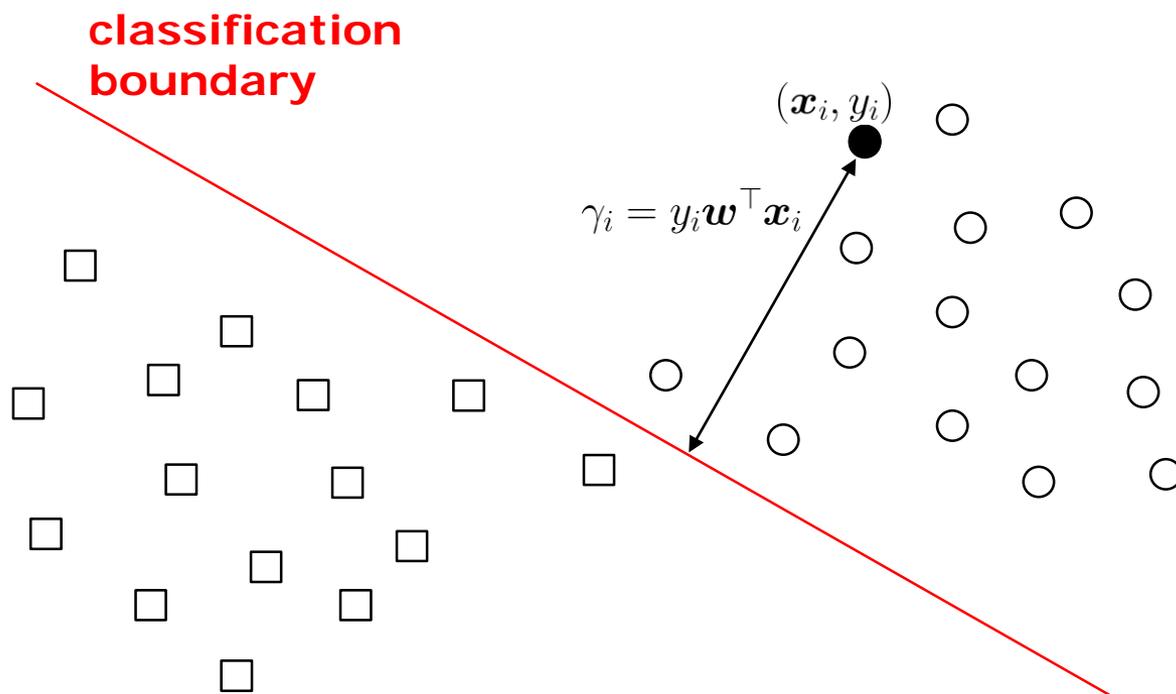
Nanjing University, China

{zhangt, Zhouzh}@lamda.nju.edu.cn



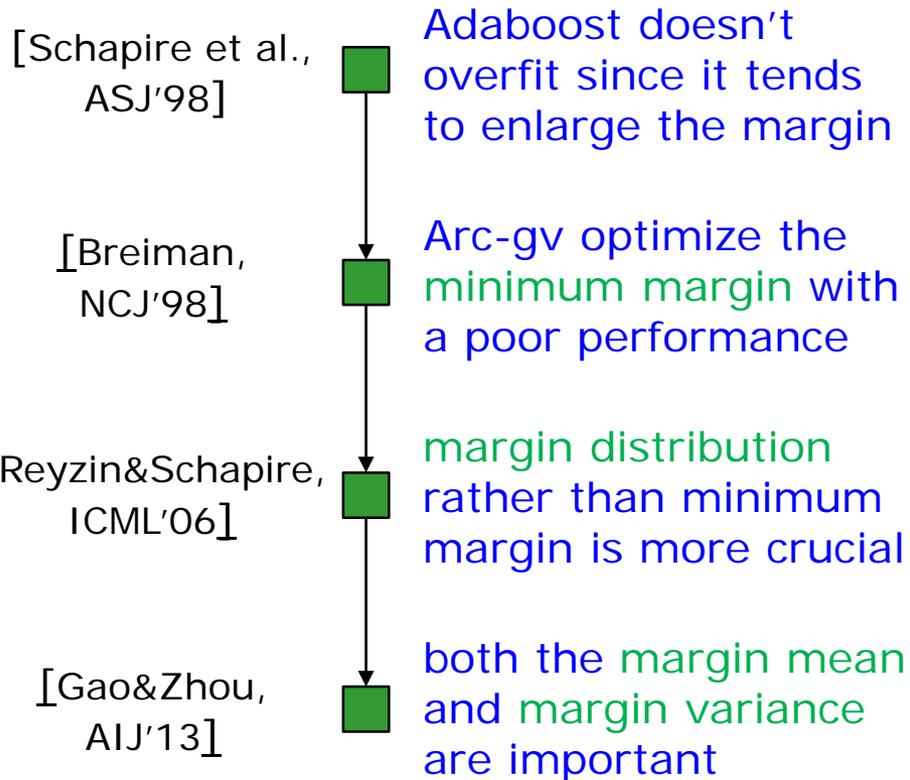
Margin

Given m training examples $\{(\mathbf{x}_i \in \mathbb{R}^d, y_i \in \{+1, -1\})\}_{i=1}^m$, the margin of instance (\mathbf{x}_i, y_i) is formulated as $\gamma_i = y_i \mathbf{w}^\top \mathbf{x}_i$.

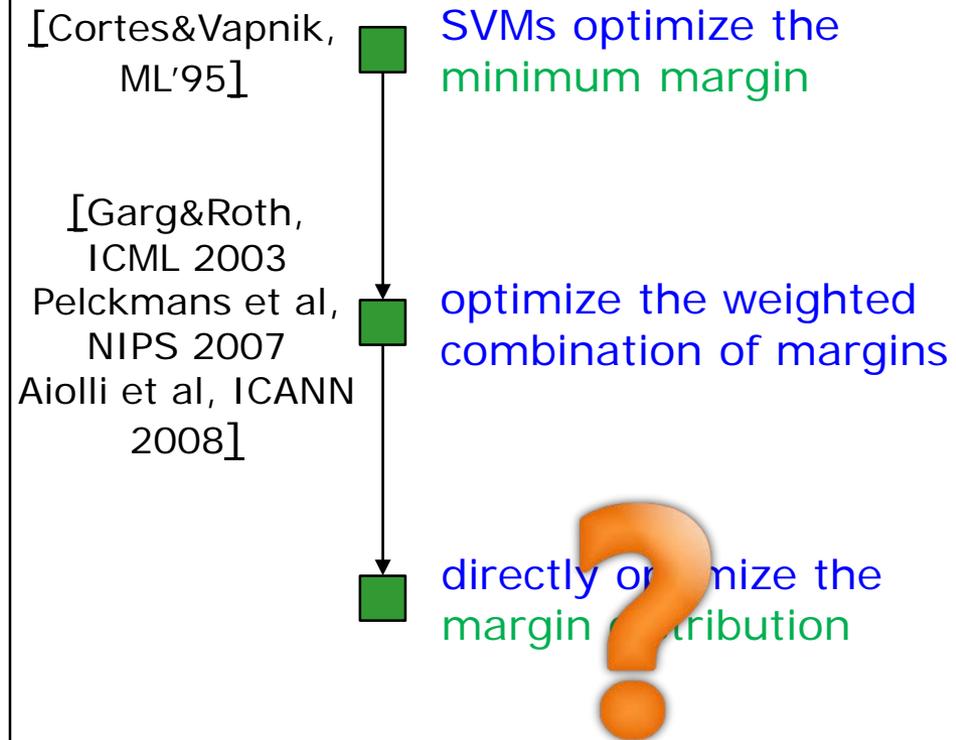


Margin Theory

Boosting



SVM



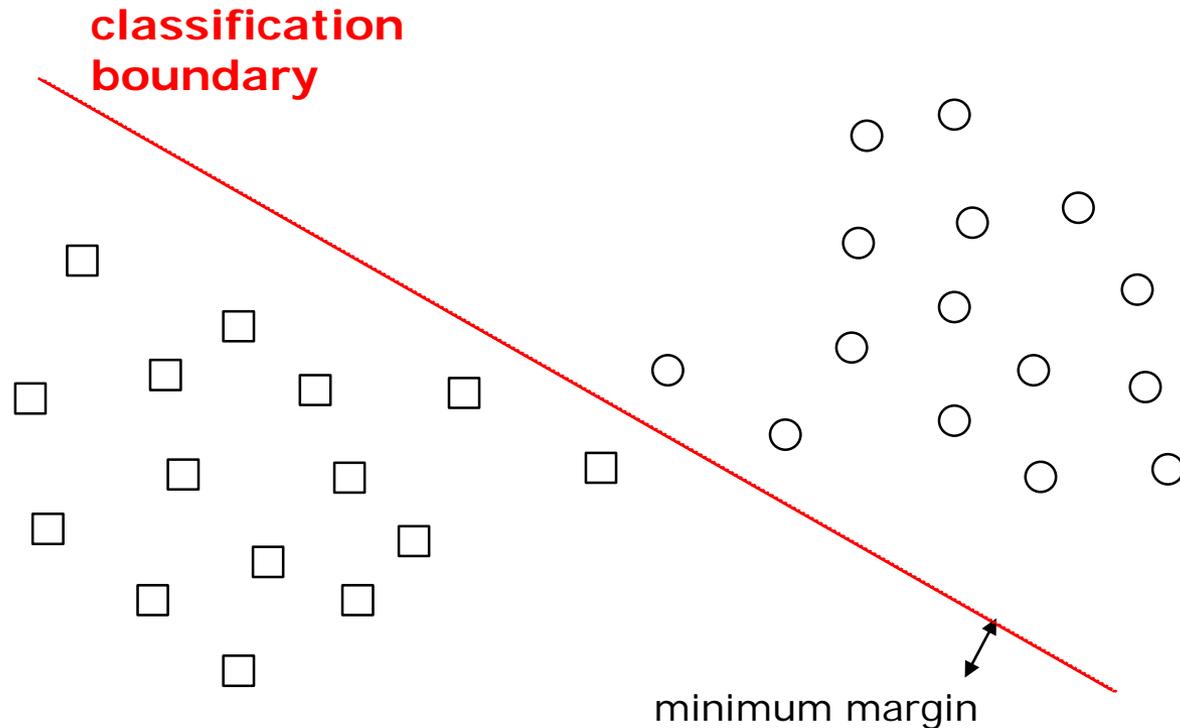
Outline

- SVM and its extensions
 - Our Method
 - Experiments
 - Conclusion
-

SVM

Try to maximize the minimum margin

$$\max \left\{ \min_{i=1}^m \gamma_i \right\} \quad \gamma_i = y_i \mathbf{w}^\top \phi(\mathbf{x}_i), \forall i = 1, \dots, m.$$



SVM ignores the whole margin distribution!

SVMs beyond minimum margin maximization

- MDO [Garg&Roth, ICML 2003]

Optimize the weighted margin combination, however, the setting of weights is heuristic, and the objective function is non-convex.

All the methods consider the margin mean but ignore the influence of margin variance. However, margin mean can not characterize the margin distribution well!
size.

- KM-OMD [Aioli et al, ICANN 2008]

Optimize the weighted margin combination. It considers hard-margin only.

Outline

- SVM and its extensions
 - Our Method
 - Experiments
 - Conclusion
-

Basic Idea

Inspired by [Gao&Zhou, AIJ'13], the **mean** and the **variance** of the margin are important.

Margin mean:

$$\bar{\gamma} = \frac{1}{m} \sum_{i=1}^m y_i \mathbf{w}^\top \phi(\mathbf{x}_i)$$

Margin variance:

$$\hat{\gamma} = \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m (y_i \mathbf{w}^\top \phi(\mathbf{x}_i) - y_j \mathbf{w}^\top \phi(\mathbf{x}_j))^2$$

Intuitive Interpretation

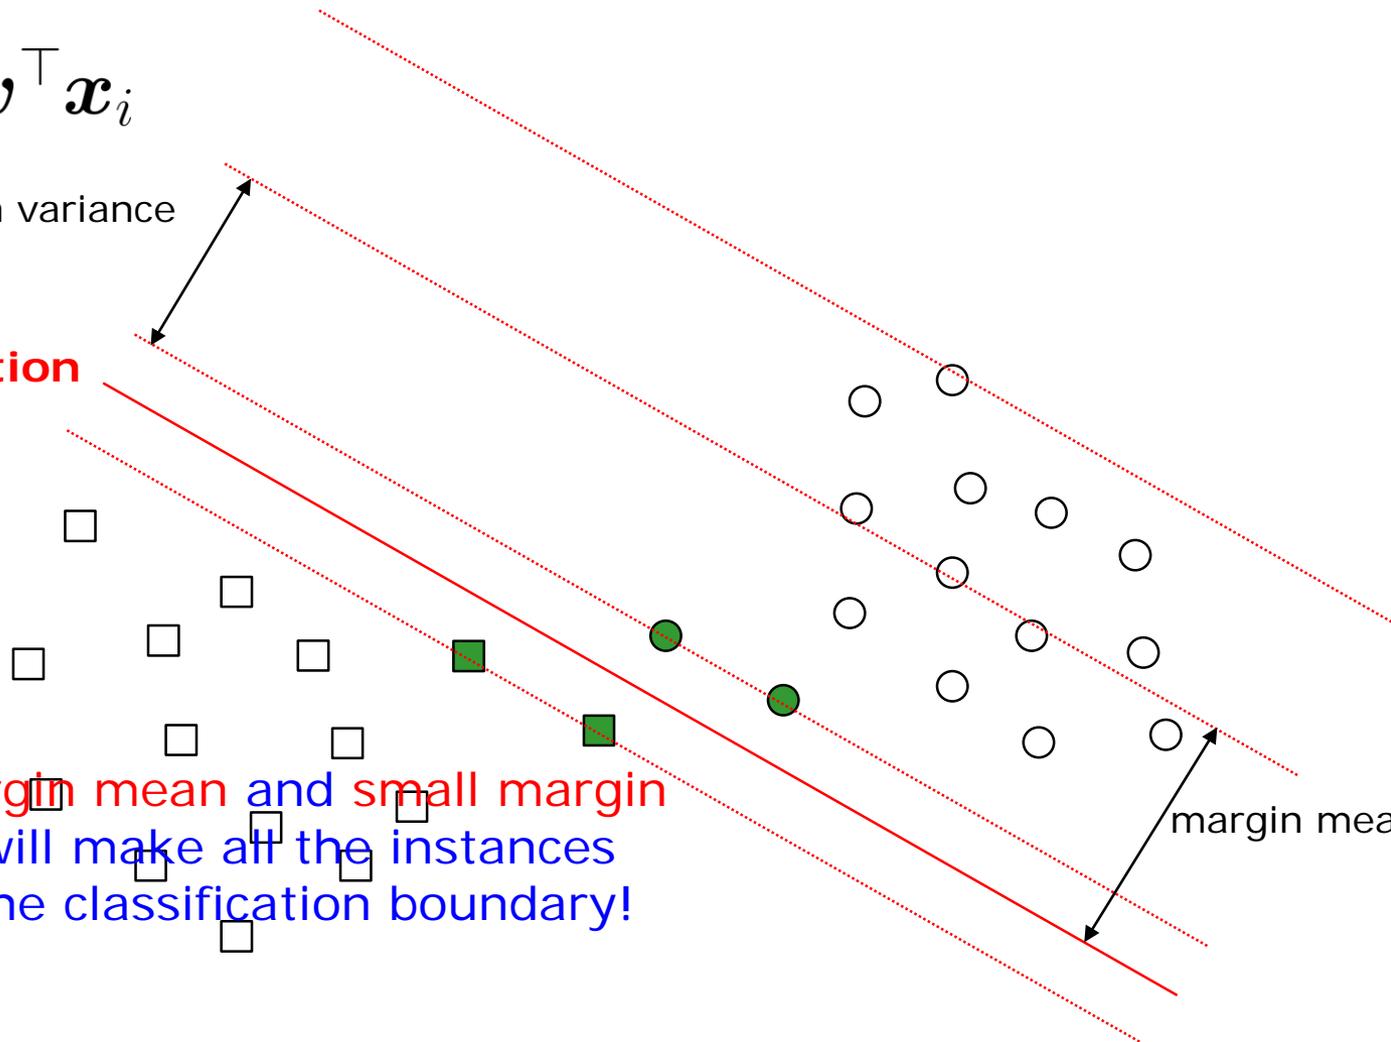
$$\gamma_i = y_i \mathbf{w}^\top \mathbf{x}_i$$

margin variance

classification boundary

Large margin mean and small margin variance will make all the instances far from the classification boundary!

margin mean



The LDM Method

LDM: Large margin Distribution Machine

Idea: maximize the margin mean and minimize the margin variance simultaneously.

Formulation:

margin variance

$$\min_{\mathbf{w}, \xi} \frac{1}{2} \mathbf{w}^\top \mathbf{w} + \frac{2\lambda_1}{m^2} (m \mathbf{w}^\top \mathbf{X} \mathbf{X}^\top \mathbf{w} - \mathbf{w}^\top \mathbf{X} \mathbf{y} \mathbf{y}^\top \mathbf{X}^\top \mathbf{w})$$

margin mean

$$- \frac{\lambda_2}{m} (\mathbf{X} \mathbf{y})^\top \mathbf{w} + C \sum_{i=1}^m \xi_i$$

$$\text{s.t. } y_i \mathbf{w}^\top \phi(\mathbf{x}_i) \geq 1 - \xi_i,$$

$$\xi_i \geq 0, \quad i = 1, \dots, m.$$

The LDM Method (cont.)

With **representer theorem**, the dual of LDM is

$$\begin{aligned} \min_{\beta} f(\beta) &= \frac{1}{2} \beta^\top \mathbf{H} \beta + \left(\frac{\lambda_2}{m} \mathbf{H} \mathbf{e} - \mathbf{e} \right)^\top \beta, \\ \text{s.t. } & 0 \leq \beta_i \leq C, \quad i = 1, \dots, m. \end{aligned}$$

- convex quadratic programming
- decoupled box constraint

It can be solved by the **coordinate descent method** efficiently since a **closed-form solution** can be achieved in each iteration.

$$\beta_i^{new} = \min \left(\max \left(\beta_i - \frac{[\nabla f(\beta)]_i}{h_{ii}}, 0 \right), C \right)$$

Large Scale Kernel LDM

For large scale problems, we can solve the prime LDM directly by **average stochastic gradient descent (ASGD)**.

Key idea of ASGD: utilize **unbiased estimation of the gradient**

THEOREM 2. If two examples (\mathbf{x}_i, y_i) and (\mathbf{x}_j, y_j) are sampled from training set randomly, then

$$\begin{aligned} \nabla g(\mathbf{w}, \mathbf{x}_i, \mathbf{x}_j) = & 4\lambda_1 \mathbf{x}_i \mathbf{x}_i^\top \mathbf{w} - 4\lambda_1 y_i \mathbf{x}_i y_j \mathbf{x}_j^\top \mathbf{w} + \mathbf{w} \\ & - \lambda_2 y_i \mathbf{x}_i - mC \begin{cases} y_i \mathbf{x}_i & i \in \mathbf{I}, \\ 0 & \text{otherwise,} \end{cases} \end{aligned} \quad (15)$$

where $\mathbf{I} \equiv \{i \mid y_i \mathbf{w}^\top \mathbf{x}_i < 1\}$ is an unbiased estimation of $\nabla g(\mathbf{w})$.

the unbiased estimation of the gradient of the objective function can be obtained by sampling **two examples** randomly in each iteration.

Theoretical Analysis

Based on **leave-one-out cross-validation estimate**, we can derive a bound on the expectation of error for LDM.

THEOREM 3. *Let α denote the optimal solution of (18), and $E[R(\alpha)]$ be the expectation of the probability of error, then we have*

$$E[R(\alpha)] \leq \frac{E[h \sum_{i \in I_1} \alpha_i + |I_2|]}{m}, \quad (19)$$

where $I_1 \equiv \{i \mid 0 < \alpha_i < C\}$, $I_2 \equiv \{i \mid \alpha_i = C\}$, $h = \max\{\text{diag}\{\mathbf{H}\}\}$ and $\mathbf{H} = \mathbf{Y}\mathbf{X}^\top \mathbf{Q}^{-1} \mathbf{X}\mathbf{Y}$.

Observation

- A similar bound also holds for SVM [Vapnik, 1995], and the only difference is that $\mathbf{H} = \mathbf{Y}\mathbf{X}^\top \mathbf{X}\mathbf{Y}$.
 - \mathbf{Q} encodes the information of the **margin distribution** with the result that the value of h for LDM is **much smaller** than SVM.
-

Outline

- Introduction
 - Our Method
 - Experiments
 - Conclusion
-

Data

20 regular scale data sets and 12 large scale data sets

Scale	Dataset	#Instance	#Feature	Dataset	#Instance	#Feature
<i>regular</i>	<i>promoters</i>	106	57	<i>haberman</i>	306	14
	<i>planning</i>	182	12	<i>vehicle</i>	435	16
	<i>colic</i>	188	13	<i>clean1</i>	476	166
<i>large</i>	<i>farm-ads</i>	4,143	54,877	<i>ijcnn1</i>	141,691	22
	<i>news20</i>	19,996	1,355,191	<i>skin</i>	245,057	3
	<i>adult-a</i>	32,561	123	<i>covtype</i>	581,012	54
	<i>w8a</i>	49,749	300	<i>rcv1</i>	697,641	47,236
	<i>cod-rna</i>	59,535	8	<i>url</i>	2,396,130	3,231,961
	<i>real-sim</i>	72,309	20,958	<i>kdd2010</i>	8,407,752	20,216,830

The size of data set ranges from 106 to more than 8,000,000, and the dimensionality ranges from 2 to more than 20,000,000, covering a broad range of properties.

Settings

Compared methods:

LDM, MDO, MAMC, KM-OMD, SVM

Parameters:

selected by 5-fold cross validation

Evaluation:

30 times hold-out tests, 1/2 training, 1/2 testing

Results

Accuracy on twenty regular scale data sets with linear kernel

Dataset	SVM	MDO	MAMC	KM-OMD	LDM
<i>promoters</i>	0.723±0.071	0.713±0.067	0.520±0.096○	0.736±0.061	0.721±0.069
<i>planning-relax</i>	0.683±0.031	0.605±0.185○	0.706±0.034●	0.479±0.050○	0.706±0.034●
<i>colic</i>	0.814±0.035	0.781±0.154	0.661±0.062○	0.813±0.028	0.832±0.026●
<i>parkinsons</i>	0.846±0.038	0.732±0.270○	0.764±0.035○	0.814±0.024○	0.865±0.030●
<i>colic.ORIG</i>	0.618±0.027	0.624±0.040	0.623±0.027	0.635±0.045●	0.619±0.042
<i>sonar</i>	0.725±0.039	0.734±0.035	0.533±0.045○	0.766±0.033●	0.736±0.036
<i>vote</i>	0.934±0.022	0.587±0.435○	0.884±0.022○	0.957±0.013●	0.970±0.014●
<i>house</i>	0.942±0.015	0.943±0.015	0.883±0.029○	0.957±0.020●	0.968±0.011●
<i>heart</i>	0.799±0.029	0.826±0.026●	0.537±0.057○	0.836±0.026●	0.791±0.030
<i>breast-cancer</i>	0.717±0.033	0.710±0.031	0.706±0.027	0.696±0.031○	0.725±0.027●
<i>haberman</i>	0.734±0.030	0.728±0.029	0.738±0.020	0.667±0.040○	0.738±0.020
<i>vehicle</i>	0.959±0.012	0.956±0.012	0.566±0.160○	0.960±0.010	0.959±0.013
<i>clean1</i>	0.803±0.035	0.798±0.031	0.561±0.025○	0.821±0.027●	0.814±0.019●
<i>wdbc</i>	0.963±0.012	0.966±0.010	0.623±0.020○	0.968±0.009●	0.968±0.011●
<i>isolet</i>	0.995±0.003	0.501±0.503○	0.621±0.207○	0.995±0.003	0.997±0.002●
<i>credit-a</i>	0.861±0.014	0.862±0.013	0.596±0.063○	0.863±0.013	0.864±0.013●
<i>austra</i>	0.857±0.013	0.842±0.055	0.567±0.044○	0.858±0.013	0.859±0.015
<i>australian</i>	0.844±0.019	0.842±0.020	0.576±0.049○	0.858±0.016●	0.866±0.014●
<i>fourclass</i>	0.724±0.014	0.377±0.238○	0.641±0.020○	0.736±0.014●	0.723±0.014
<i>german</i>	0.711±0.030	0.737±0.014●	0.697±0.017○	0.729±0.017●	0.738±0.016●
Ave. accuracy	0.813	0.743	0.650	0.807	0.823
LDM: w/t/l	12/8/0	9/10/1	17/3/0	10/5/5	

w/t/l counts: after t-tests (95% SI)

bold: best

●/○: significantly better/worse than SVM

LDM achieves the best accuracy on **13** data sets

LDM outperforms SVM, MDO, MAMC and KM-OMD for **12, 9, 17, 10** times respectively

Results (cont.)

Accuracy on twenty regular scale data sets with RBF kernel

Dataset	SVM	MDO	MAMC	KM-OMD	LDM
<i>promoters</i>	0.684±0.100	N/A	0.638±0.121○	0.701±0.085	0.715±0.074●
<i>planning-relax</i>	0.708±0.035	N/A	0.706±0.034	0.683±0.031○	0.707±0.034
<i>colic</i>	0.822±0.033	N/A	0.623±0.037○	0.825±0.024	0.841±0.018●
<i>parkinsons</i>	0.929±0.029	N/A	0.852±0.036○	0.906±0.033○	0.927±0.029
<i>colic.ORIG</i>	0.638±0.043	N/A	0.623±0.027	0.621±0.039	0.641±0.044
<i>sonar</i>	0.842±0.034	N/A	0.753±0.052○	0.821±0.051○	0.846±0.032
<i>vote</i>	0.946±0.016	N/A	0.913±0.019○	0.930±0.029○	0.968±0.013●
<i>house</i>	0.953±0.020	N/A	0.561±0.139○	0.938±0.022○	0.964±0.013●
<i>heart</i>	0.808±0.025	N/A	0.540±0.043○	0.805±0.048	0.822±0.029●
<i>breast-cancer</i>	0.729±0.030	N/A	0.706±0.027○	0.691±0.024○	0.753±0.027●
<i>haberman</i>	0.727±0.024	N/A	0.742±0.021●	0.676±0.042○	0.731±0.027
<i>vehicle</i>	0.992±0.007	N/A	0.924±0.025○	0.988±0.008○	0.993±0.006
<i>clean1</i>	0.890±0.020	N/A	0.561±0.025○	0.772±0.043○	0.891±0.024
<i>wdbc</i>	0.951±0.011	N/A	0.740±0.042○	0.941±0.040	0.961±0.010●
<i>isolet</i>	0.998±0.002	N/A	0.994±0.004○	0.995±0.003○	0.998±0.002
<i>credit-a</i>	0.858±0.014	N/A	0.542±0.032○	0.845±0.029○	0.861±0.013
<i>austra</i>	0.853±0.013	N/A	0.560±0.018○	0.854±0.017	0.857±0.014●
<i>australian</i>	0.815±0.014	N/A	0.554±0.015○	0.860±0.014●	0.854±0.016●
<i>fourclass</i>	0.998±0.003	N/A	0.791±0.014○	0.838±0.014○	0.998±0.003
<i>german</i>	0.731±0.019	N/A	0.697±0.017○	0.742±0.017●	0.743±0.016●
Ave. accuracy	0.844	N/A	0.701	0.822	0.854
LDM: w/t/l	10/10/0	N/A	18/1/1	15/5/0	

w/t/l counts: after t-tests (95% SI)

bold: best

●/○: significantly better/worse than SVM

LDM achieves the best accuracy on 15 data sets

LDM outperforms SVM, MAMC and KM-OMD for 10, 18, 15 times respectively

Results (cont.)

Accuracy on twelve large scale data sets with linear kernel

Results (cont.)

Accuracy on twelve large scale data sets with linear kernel

Dataset	SVM	MDO	MAMC	KM-OMD	LDM
<i>farm-ads</i>	0.880±0.007	0.880±0.007	0.759±0.038○	N/A	0.890±0.008●
<i>news20</i>	0.954±0.002	0.948±0.002○	0.772±0.017○	N/A	0.960±0.001●
<i>adult-a</i>	0.845±0.002	0.788±0.053○	0.759±0.002○	N/A	0.846±0.003●
<i>w8a</i>	0.983±0.001	0.985±0.001●	0.971±0.001○	N/A	0.983±0.001
<i>cod-rna</i>	0.899±0.001	0.774±0.203	0.667±0.001○	N/A	0.899±0.001
<i>real-sim</i>	0.961±0.001	0.955±0.002○	0.744±0.004○	N/A	0.971±0.001●
<i>ijcnn1</i>	0.921±0.003	0.921±0.002	0.904±0.001○	N/A	0.921±0.002
<i>skin</i>	0.934±0.001	0.929±0.003○	0.792±0.000○	N/A	0.934±0.001
<i>covtype</i>	0.762±0.001	0.760±0.003○	0.628±0.002○	N/A	0.763±0.001
<i>rcv1</i>	0.969±0.000	0.959±0.000○	0.913±0.000○	N/A	0.977±0.000●
<i>url</i>	0.993±0.006	0.993±0.006	0.670±0.000○	N/A	0.993±0.006
<i>kdd2010</i>	0.852±0.001	N/A	0.853±0.000●	N/A	0.881±0.001●
Ave. accuracy	0.913	0.899	0.786	N/A	0.919
LDM: w/t/l	6/6/0	7/3/1	12/0/0	N/A	

Results (cont.)

Accuracy on twelve large scale data sets with linear kernel

Dataset	SVM	MDO	MAMC	KM-OMD	LDM
<i>farm-ads</i>	0.880±0.007	0.880±0.007	0.759±0.038○	N/A	0.890±0.008●
<i>news20</i>	0.954±0.002	0.948±0.002○	0.772±0.017○	N/A	0.960±0.001●
<i>adult-a</i>	0.845±0.002	0.788±0.053○	0.759±0.002○	N/A	0.846±0.003●
<i>w8a</i>	0.983±0.001	0.985±0.001●	0.971±0.001○	N/A	0.983±0.001
<i>cod-rna</i>	0.899±0.001	0.774±0.203	0.667±0.001○	N/A	0.899±0.001
<i>real-sim</i>	0.961±0.001	0.955±0.002○	0.744±0.004○	N/A	0.971±0.001●
<i>ijcnn1</i>	0.921±0.003	0.921±0.002	0.904±0.001○	N/A	0.921±0.002
<i>skin</i>	0.934±0.001	0.929±0.003○	0.792±0.000○	N/A	0.934±0.001
<i>covtype</i>	0.762±0.001	0.760±0.003○	0.628±0.002○	N/A	0.763±0.001
<i>rcv1</i>	0.969±0.000	0.959±0.000○	0.913±0.000○	N/A	0.977±0.000●
<i>url</i>	0.993±0.006	0.993±0.006	0.670±0.000○	N/A	0.993±0.006
<i>kdd2010</i>	0.852±0.001	N/A	0.853±0.000●	N/A	0.881±0.001●
Ave. accuracy	0.913	0.899	0.786	N/A	0.919
LDM: w/t/l	6/6/0	7/3/1	12/0/0	N/A	

w/t/l counts: after t-tests (95% SI)

bold: best

●/○: significantly better/worse than SVM

Results (cont.)

Accuracy on twelve large scale data sets with linear kernel

Dataset	SVM	MDO	MAMC	KM-OMD	LDM
<i>farm-ads</i>	0.880±0.007	0.880±0.007	0.759±0.038○	N/A	0.890±0.008●
<i>news20</i>	0.954±0.002	0.948±0.002○	0.772±0.017○	N/A	0.960±0.001●
<i>adult-a</i>	0.845±0.002	0.788±0.053○	0.759±0.002○	N/A	0.846±0.003●
<i>w8a</i>	0.983±0.001	0.985±0.001●	0.971±0.001○	N/A	0.983±0.001
<i>cod-rna</i>	0.899±0.001	0.774±0.203	0.667±0.001○	N/A	0.899±0.001
<i>real-sim</i>	0.961±0.001	0.955±0.002○	0.744±0.004○	N/A	0.971±0.001●
<i>ijcnn1</i>	0.921±0.003	0.921±0.002	0.904±0.001○	N/A	0.921±0.002
<i>skin</i>	0.934±0.001	0.929±0.003○	0.792±0.000○	N/A	0.934±0.001
<i>covtype</i>	0.762±0.001	0.760±0.003○	0.628±0.002○	N/A	0.763±0.001
<i>rcv1</i>	0.969±0.000	0.959±0.000○	0.913±0.000○	N/A	0.977±0.000●
<i>url</i>	0.993±0.006	0.993±0.006	0.670±0.000○	N/A	0.993±0.006
<i>kdd2010</i>	0.852±0.001	N/A	0.853±0.000●	N/A	0.881±0.001●
Ave. accuracy	0.913	0.899	0.786	N/A	0.919
LDM: w/t/l	6/6/0	7/3/1	12/0/0	N/A	

w/t/l counts: after t-tests (95% SI)

bold: best

●/○: significantly better/worse than SVM

LDM achieves the best accuracy on 8 data sets

Results (cont.)

Accuracy on twelve large scale data sets with linear kernel

Dataset	SVM	MDO	MAMC	KM-OMD	LDM
<i>farm-ads</i>	0.880±0.007	0.880±0.007	0.759±0.038○	N/A	0.890±0.008●
<i>news20</i>	0.954±0.002	0.948±0.002○	0.772±0.017○	N/A	0.960±0.001●
<i>adult-a</i>	0.845±0.002	0.788±0.053○	0.759±0.002○	N/A	0.846±0.003●
<i>w8a</i>	0.983±0.001	0.985±0.001●	0.971±0.001○	N/A	0.983±0.001
<i>cod-rna</i>	0.899±0.001	0.774±0.203	0.667±0.001○	N/A	0.899±0.001
<i>real-sim</i>	0.961±0.001	0.955±0.002○	0.744±0.004○	N/A	0.971±0.001●
<i>ijcnn1</i>	0.921±0.003	0.921±0.002	0.904±0.001○	N/A	0.921±0.002
<i>skin</i>	0.934±0.001	0.929±0.003○	0.792±0.000○	N/A	0.934±0.001
<i>covtype</i>	0.762±0.001	0.760±0.003○	0.628±0.002○	N/A	0.763±0.001
<i>rcv1</i>	0.969±0.000	0.959±0.000○	0.913±0.000○	N/A	0.977±0.000●
<i>url</i>	0.993±0.006	0.993±0.006	0.670±0.000○	N/A	0.993±0.006
<i>kdd2010</i>	0.852±0.001	N/A	0.853±0.000●	N/A	0.881±0.001●
Ave. accuracy	0.913	0.899	0.786	N/A	0.919
LDM: w/t/l	6/6/0	7/3/1	12/0/0	N/A	

w/t/l counts: after t-tests (95% SI)

bold: best

●/○: significantly better/worse than SVM

LDM outperforms SVM, MDO and MAMC for 6, 7, 12 times respectively

LDM achieves the best accuracy on 8 data sets

Results (cont.)

Accuracy on twelve large scale data sets with linear kernel

Dataset	SVM	MDO	MAMC	KM-OMD	LDM
<i>farm-ads</i>	0.880±0.007	0.880±0.007	0.759±0.038○	N/A	0.890±0.008●
<i>news20</i>	0.954±0.002	0.948±0.002○	0.772±0.017○	N/A	0.960±0.001●
<i>adult-a</i>	0.845±0.002	0.788±0.053○	0.759±0.002○	N/A	0.846±0.003●
<i>w8a</i>	0.983±0.001	0.985±0.001●	0.971±0.001○	N/A	0.983±0.001
<i>cod-rna</i>	0.899±0.001	0.774±0.203	0.667±0.001○	N/A	0.899±0.001
<i>real-sim</i>	0.961±0.001	0.955±0.002○	0.744±0.004○	N/A	0.971±0.001●
<i>ijcnn1</i>	0.921±0.003	0.921±0.002	0.904±0.001○	N/A	0.921±0.002
<i>skin</i>	0.934±0.001	0.929±0.003○	0.792±0.000○	N/A	0.934±0.001
<i>covtype</i>	0.762±0.001	0.760±0.003○	0.628±0.002○	N/A	0.763±0.001
<i>rcv1</i>	0.969±0.000	0.959±0.000○	0.913±0.000○	N/A	0.977±0.000●
<i>url</i>	0.993±0.006	0.993±0.006	0.670±0.000○	N/A	0.993±0.006
<i>kdd2010</i>	0.852±0.001	N/A	0.853±0.000●	N/A	0.881±0.001●
Ave. accuracy	0.913	0.899	0.786	N/A	0.919
LDM: w/t/l	6/6/0	7/3/1	12/0/0	N/A	

w/t/l counts: after t-tests (95% SI)

bold: best

●/○: significantly better/worse than SVM

LDM outperforms SVM, MDO and MAMC for 6, 7, 12 times respectively

LDM achieves the best accuracy on 8 data sets

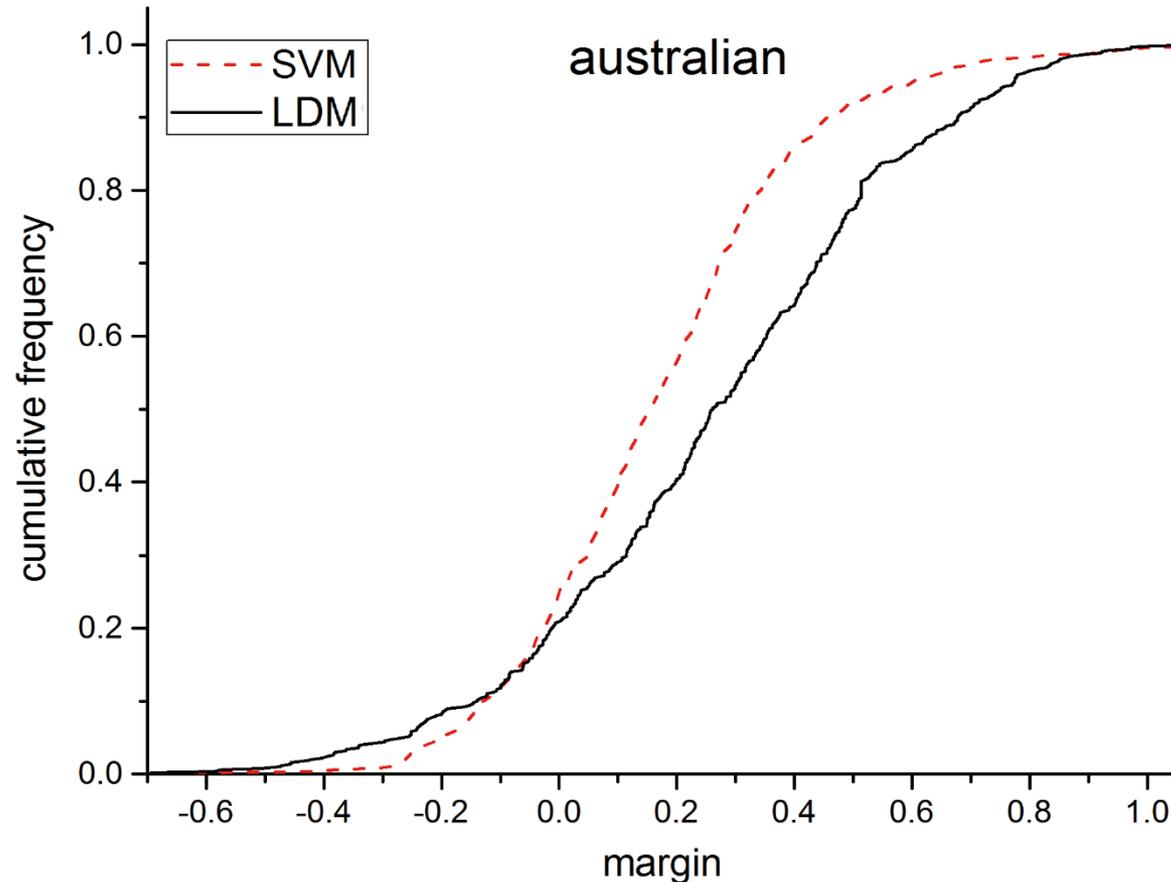
LDM is significantly better than the other methods

Margin Distribution

Cumulative frequency (y-axis) with respect to margin (x-axis) of SVM and LDM

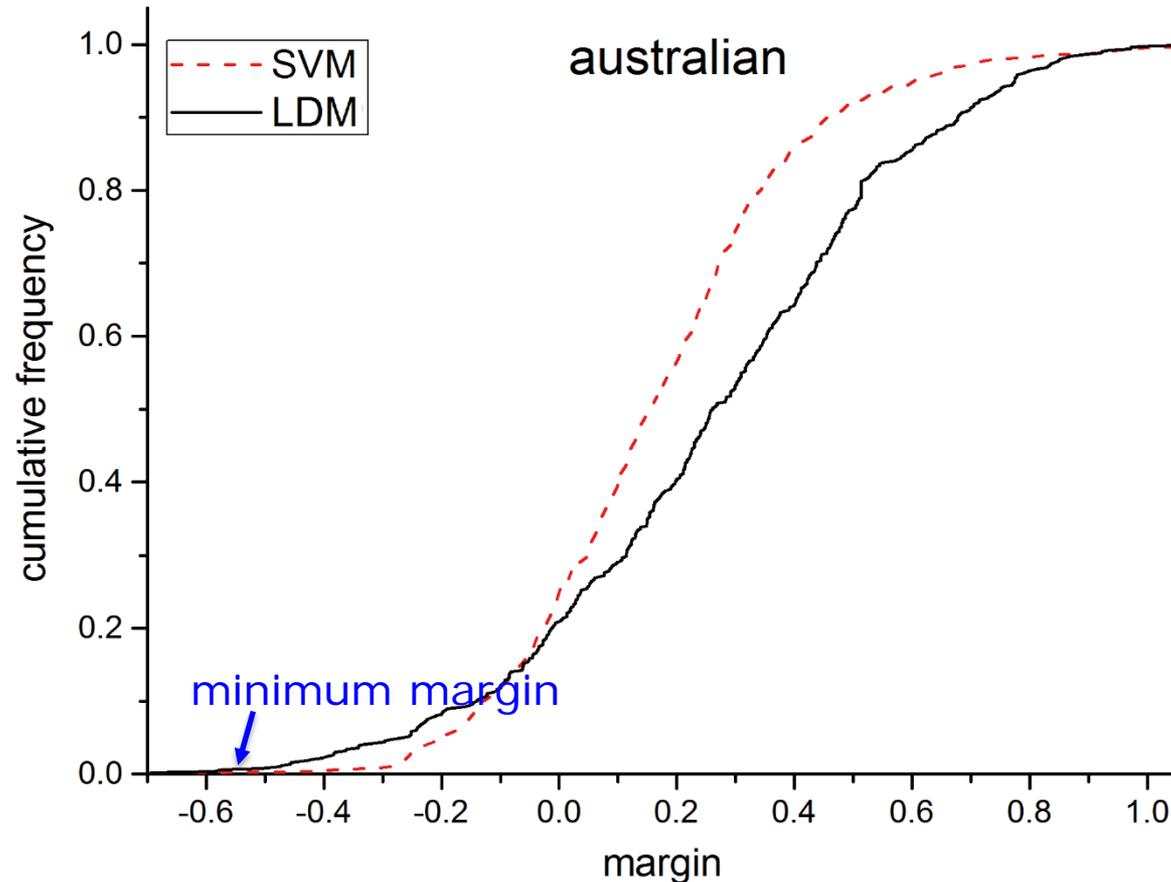
Margin Distribution

Cumulative frequency (y-axis) with respect to margin (x-axis) of SVM and LDM



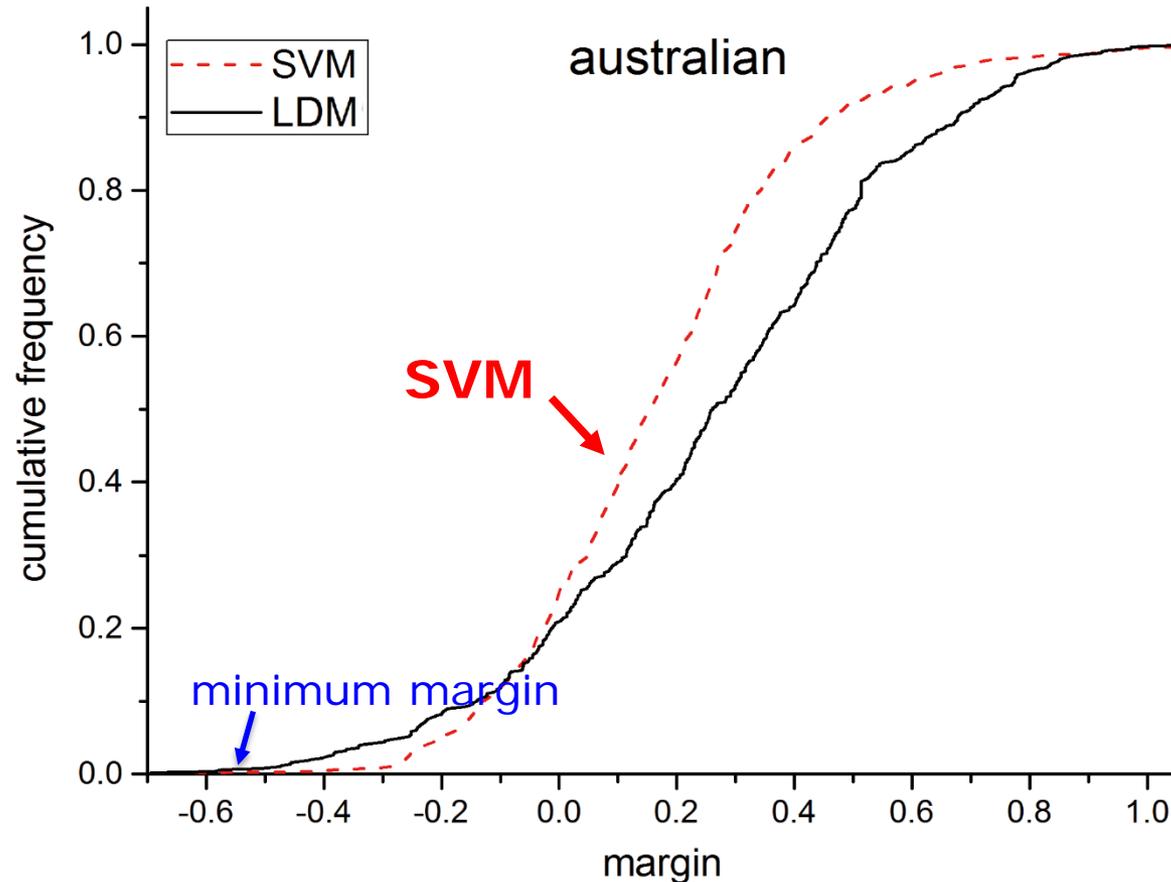
Margin Distribution

Cumulative frequency (y-axis) with respect to margin (x-axis) of SVM and LDM



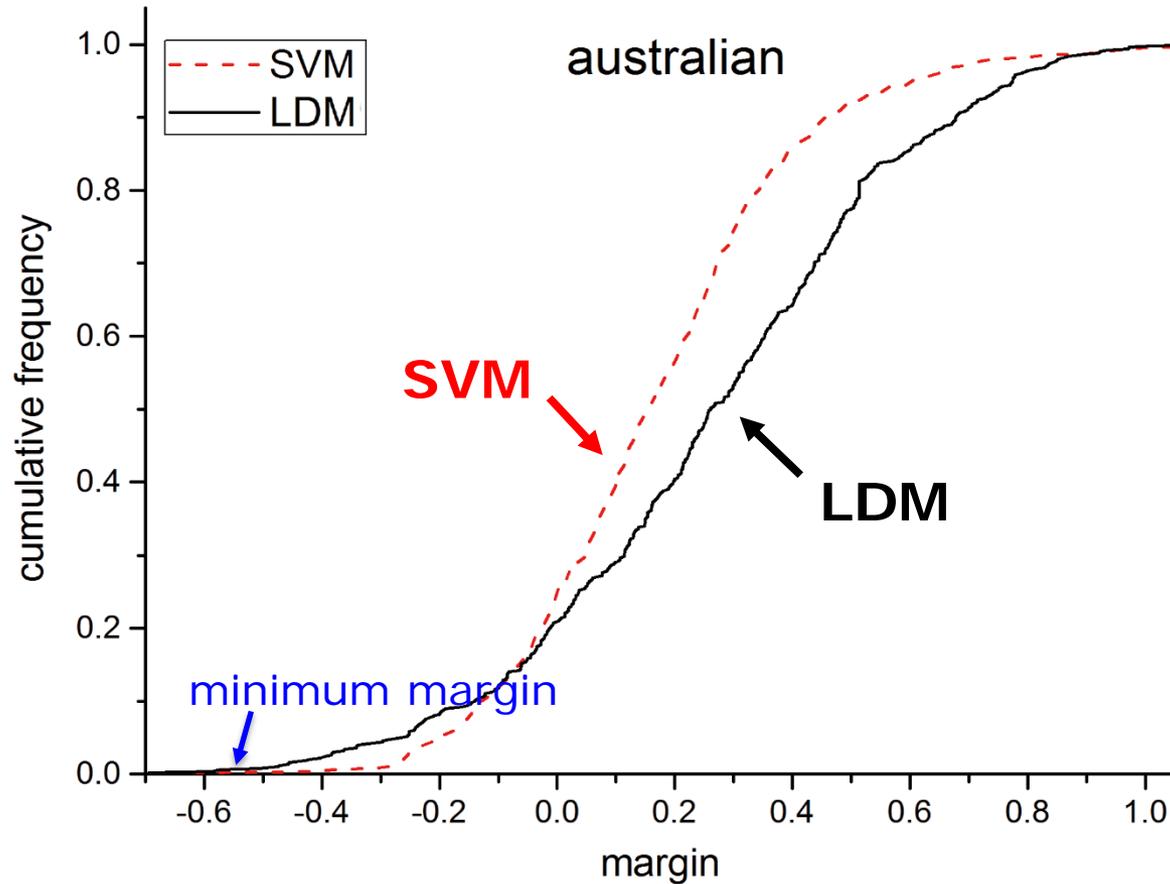
Margin Distribution

Cumulative frequency (y-axis) with respect to margin (x-axis) of SVM and LDM



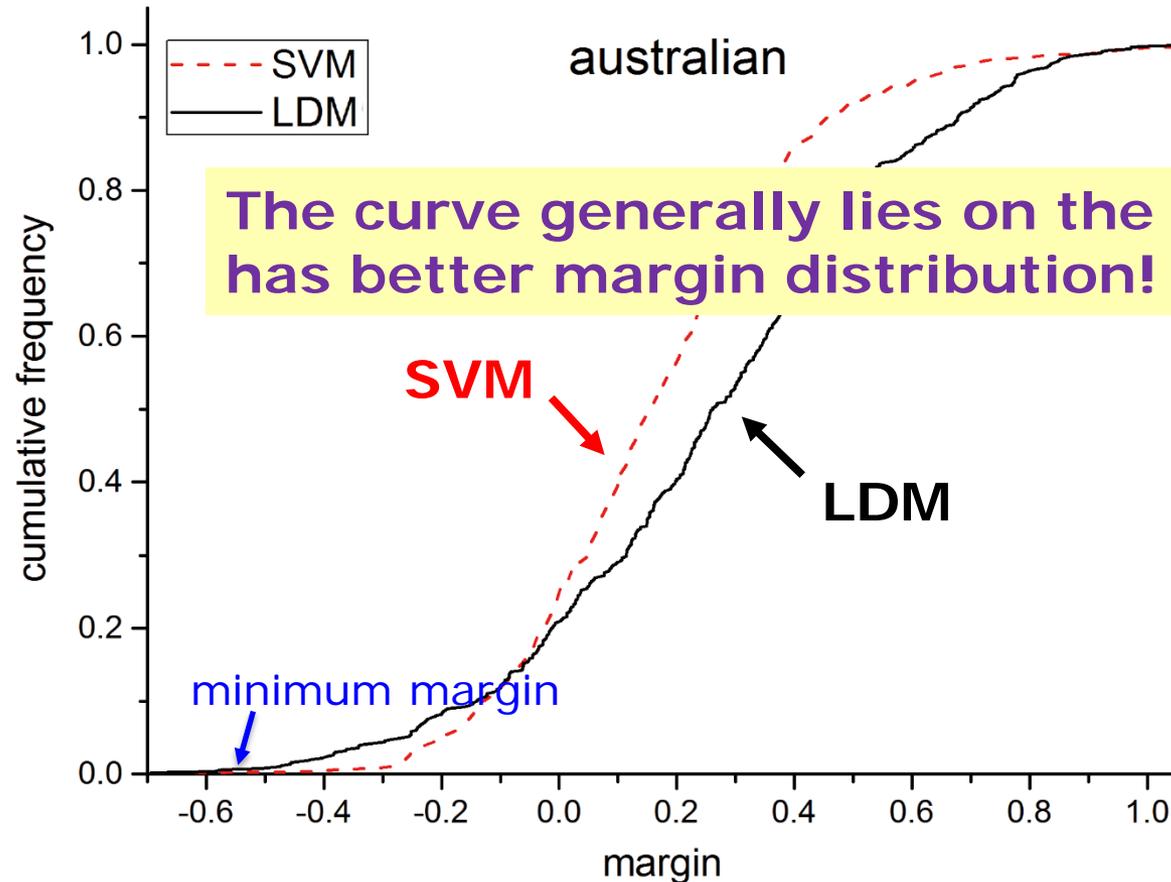
Margin Distribution

Cumulative frequency (y-axis) with respect to margin (x-axis) of SVM and LDM



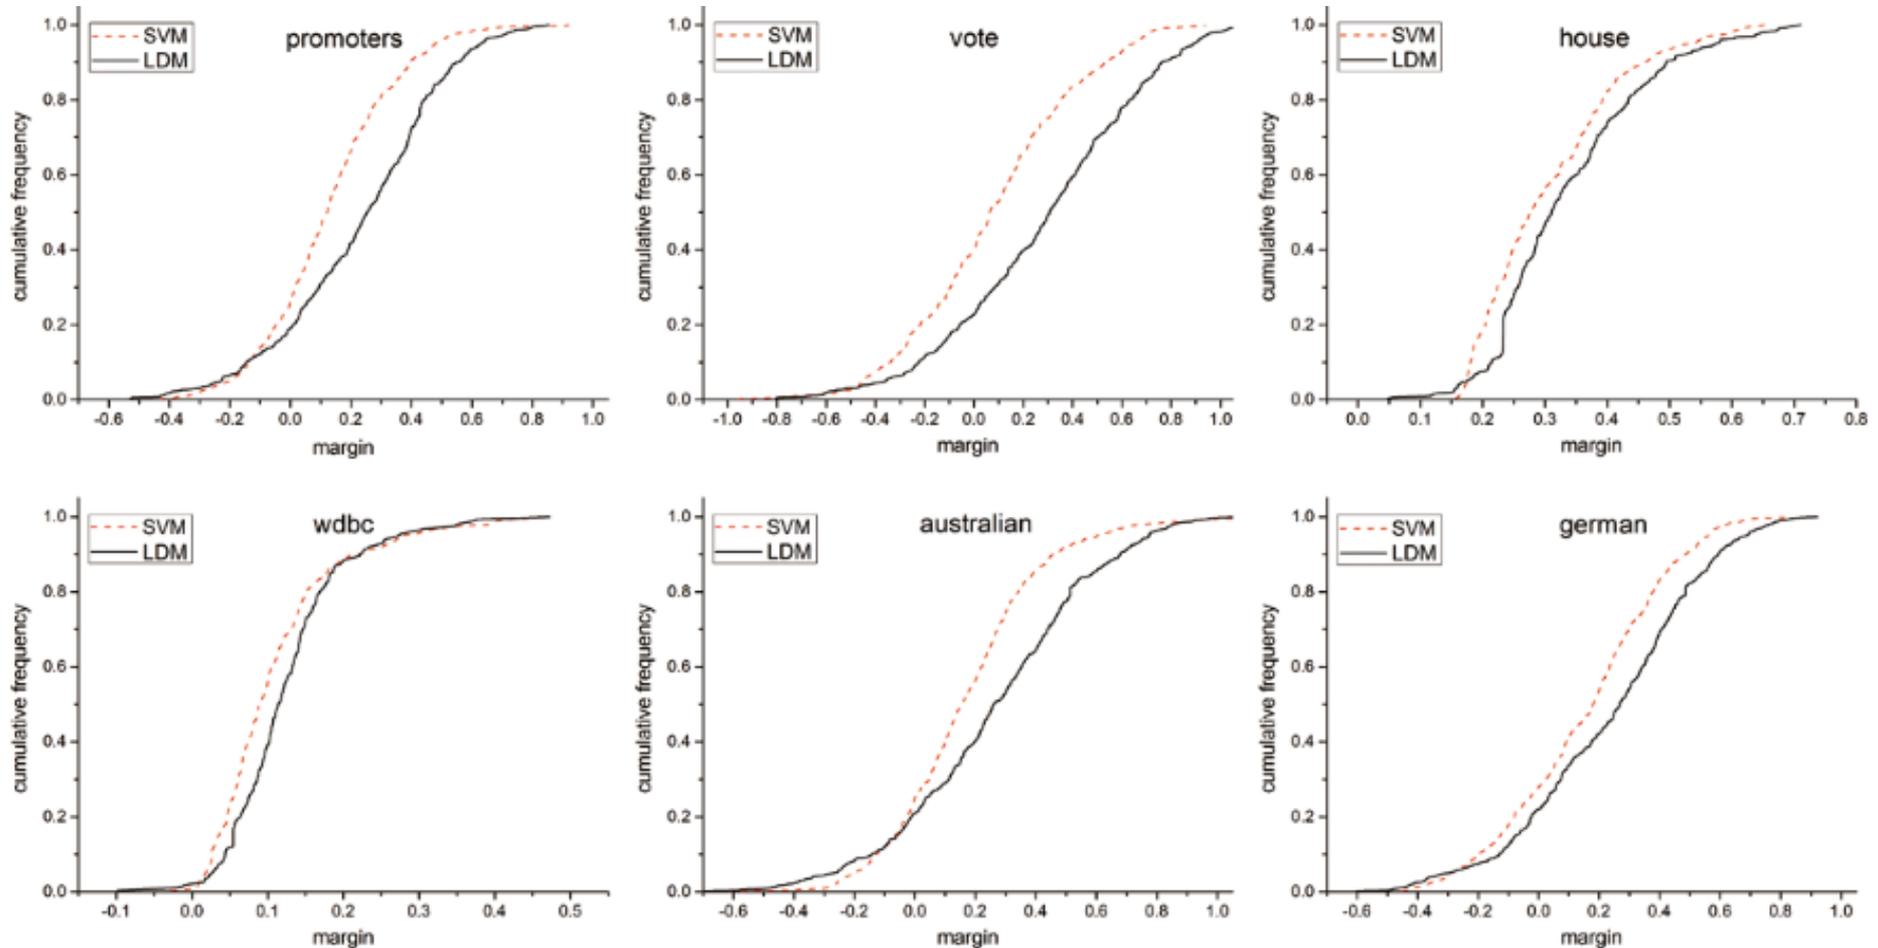
Margin Distribution

Cumulative frequency (y-axis) with respect to margin (x-axis) of SVM and LDM



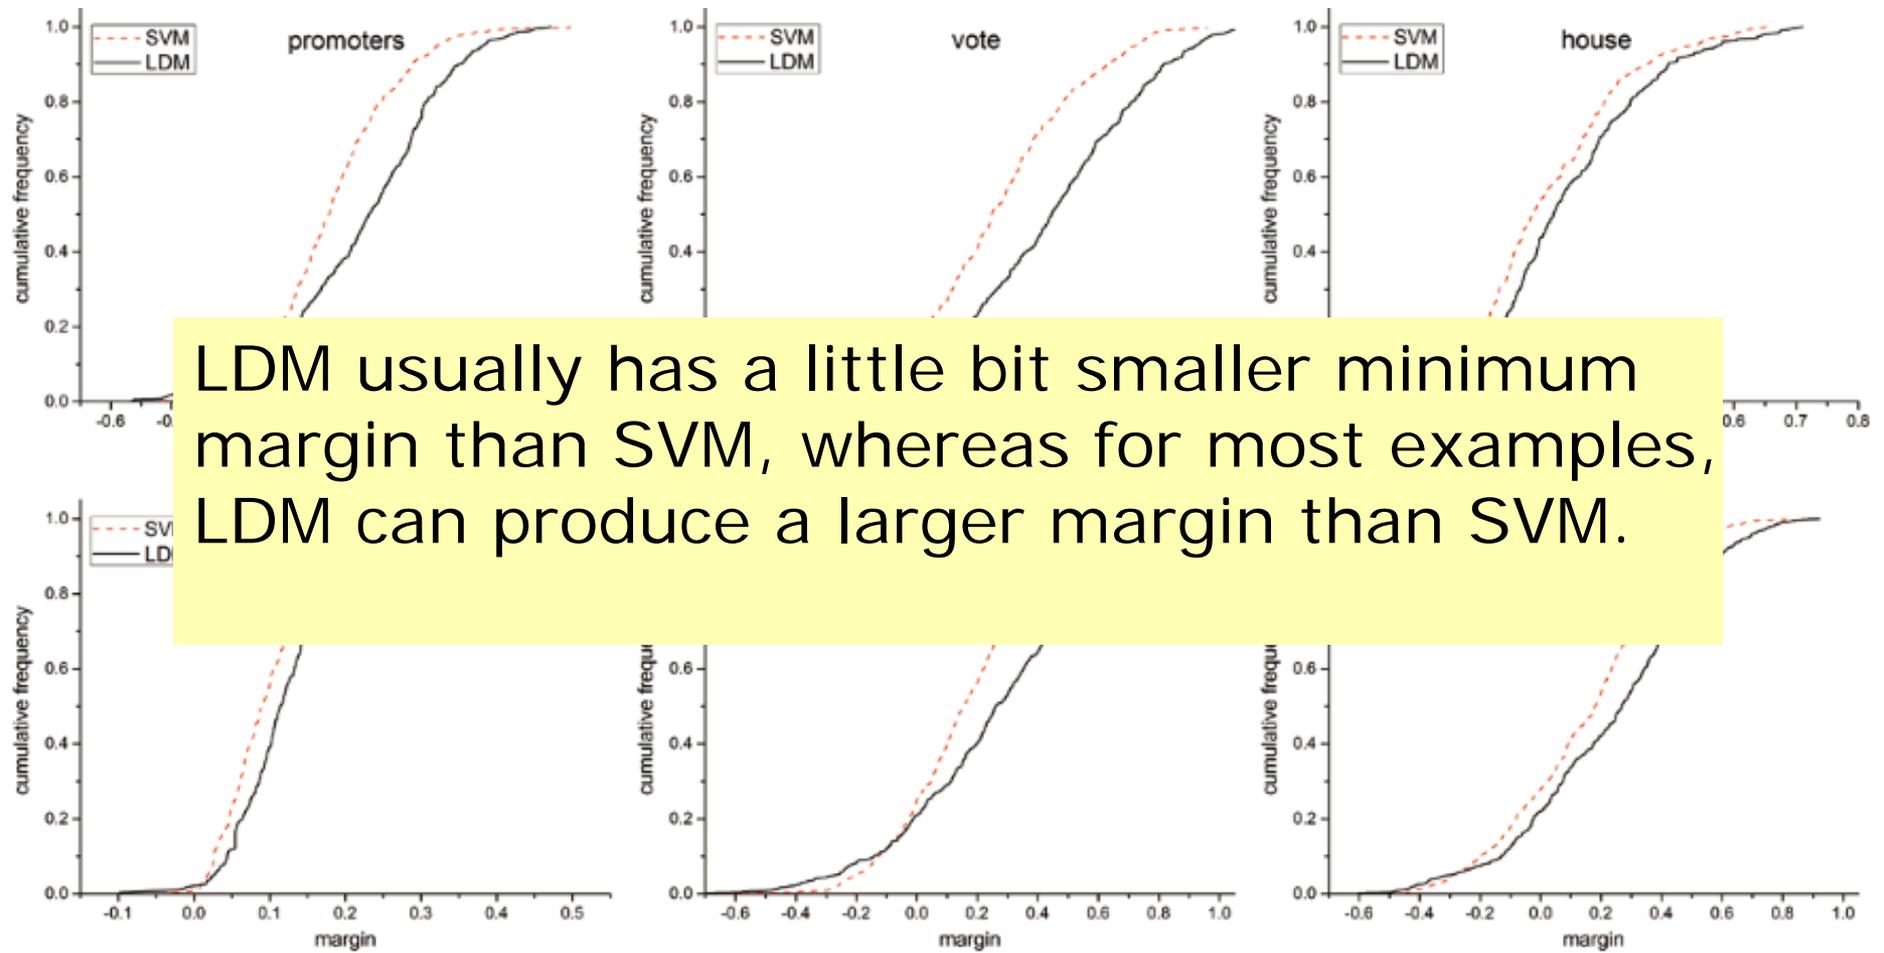
Margin Distribution

Cumulative frequency (y-axis) with respect to margin (x-axis) of SVM and LDM



Margin Distribution

Cumulative frequency (y-axis) with respect to margin (x-axis) of SVM and LDM



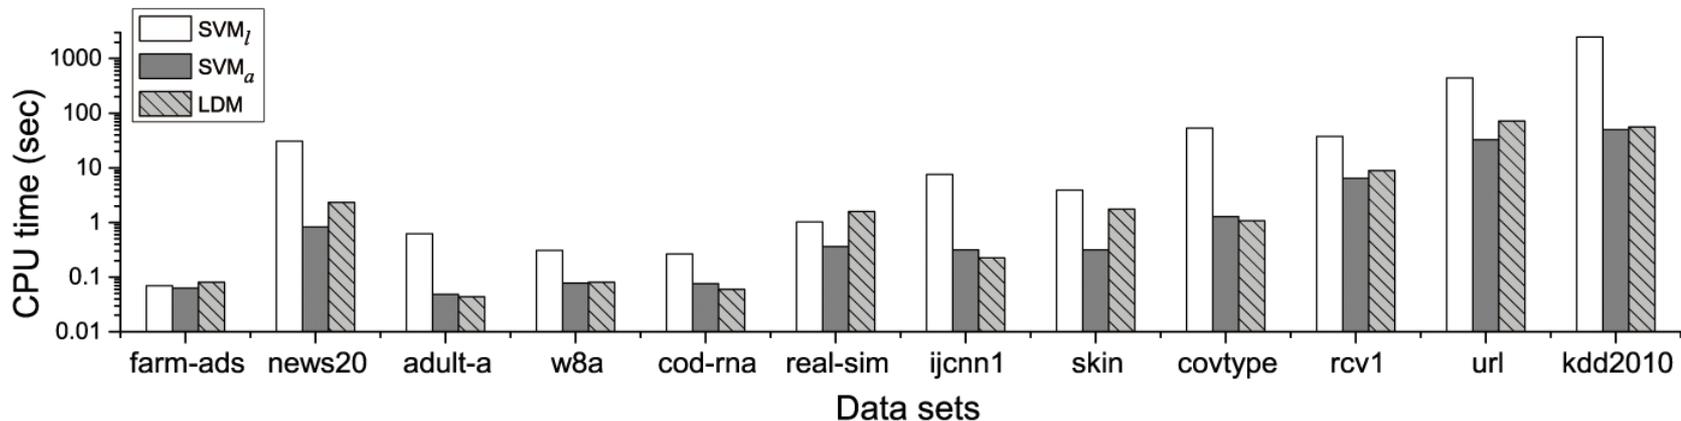
LDM usually has a little bit smaller minimum margin than SVM, whereas for most examples, LDM can produce a larger margin than SVM.

Scalability

The time cost of LDM and SVM on the twelve large scale data sets

Scalability

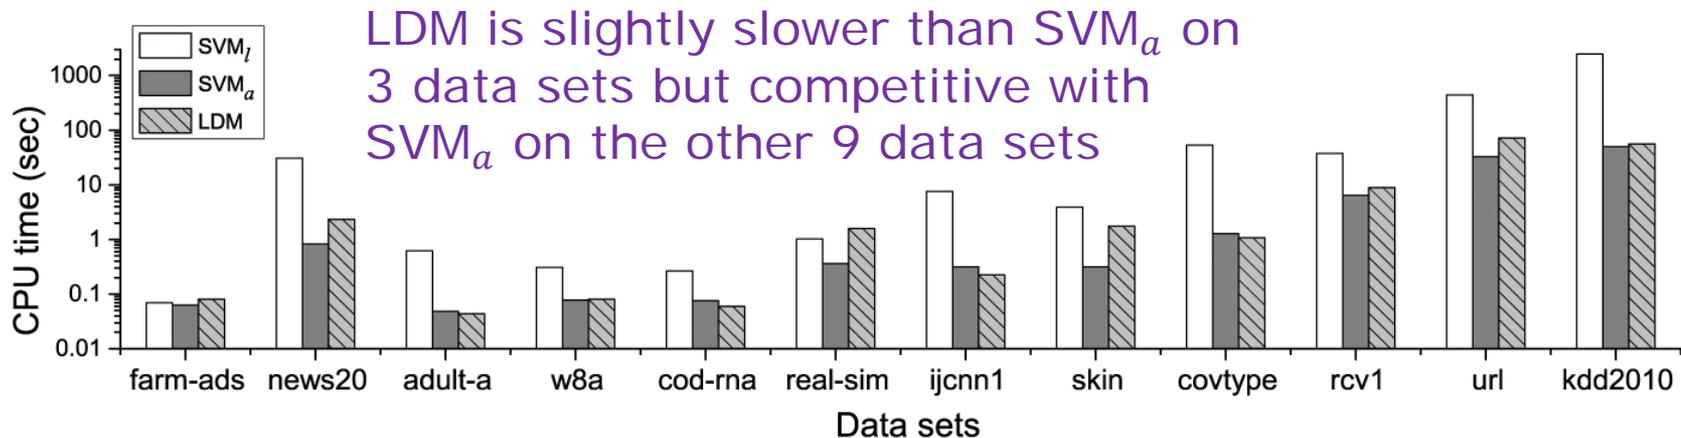
The time cost of LDM and SVM on the twelve large scale data sets



- SVM_l: SVM implemented by the LIBLINEAR package
- SVM_a: SVM implemented by ASGD
- LDM: LDM implemented by ASGD

Scalability

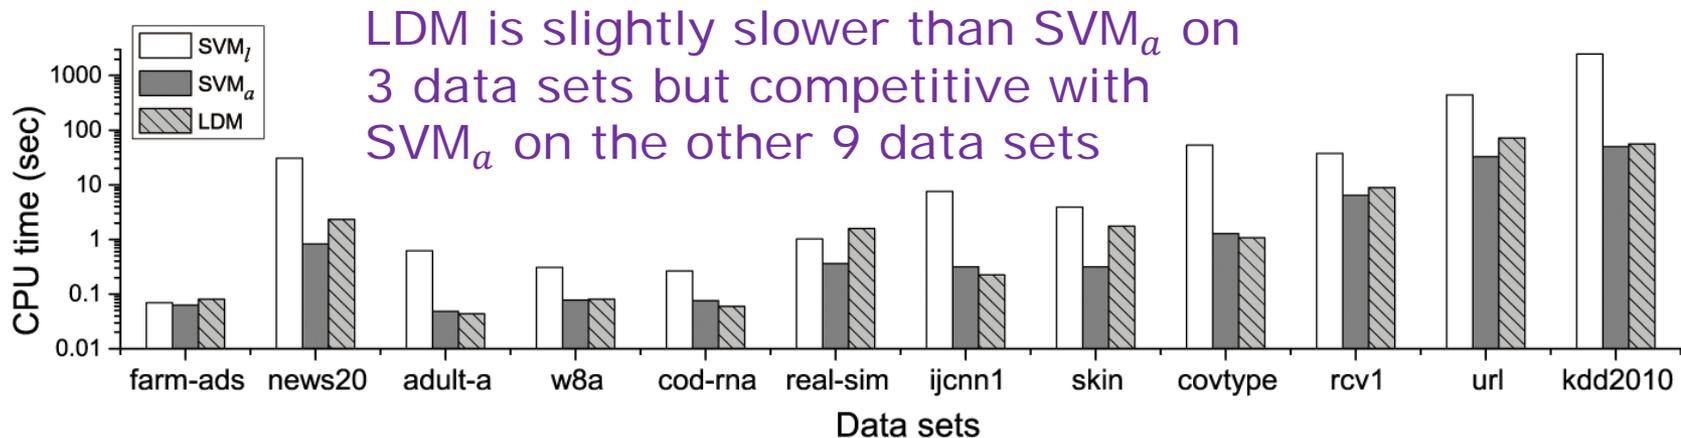
The time cost of LDM and SVM on the twelve large scale data sets



- SVM_l: SVM implemented by the LIBLINEAR package
- SVM_a: SVM implemented by ASGD
- LDM: LDM implemented by ASGD

Scalability

The time cost of LDM and SVM on the twelve large scale data sets



- SVM_l: SVM implemented by the LIBLINEAR package
- SVM_a: SVM implemented by ASGD
- LDM: LDM implemented by ASGD

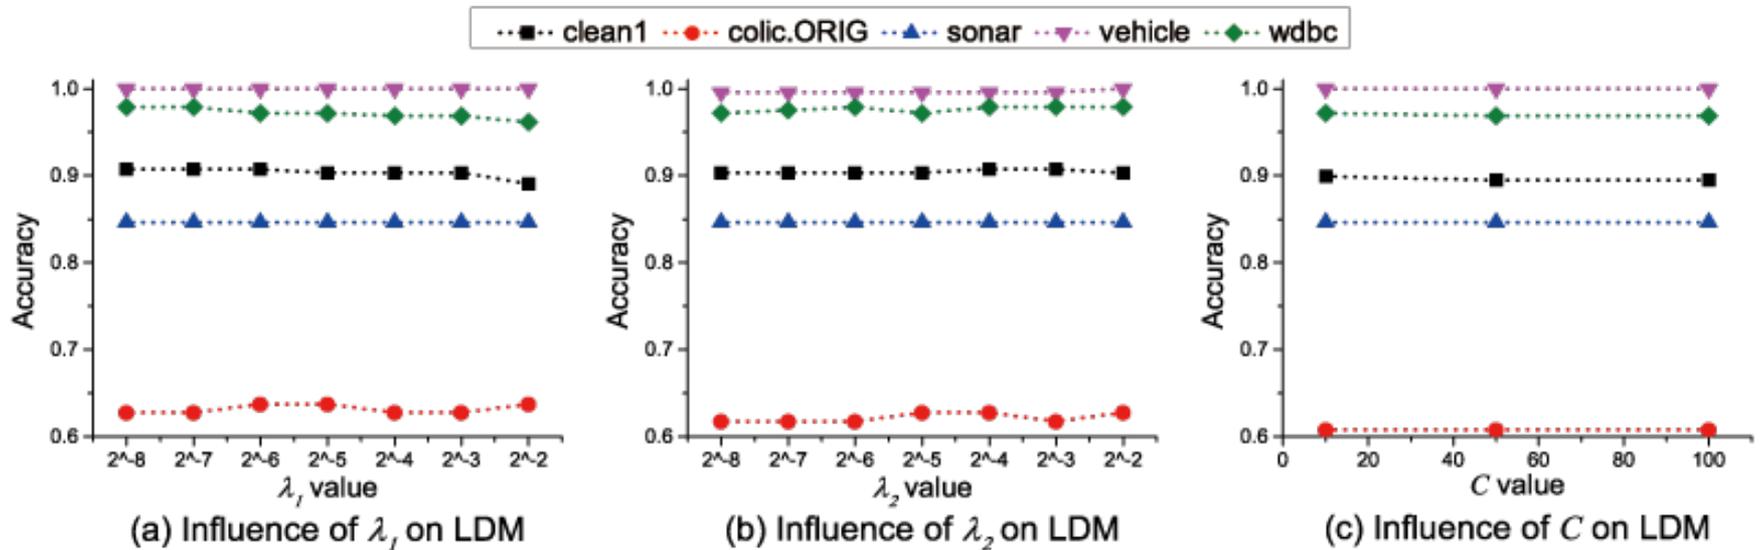
LDM is also computationally efficient.

Parameter Influence

LDM has three parameters: λ_1 , λ_2 and C

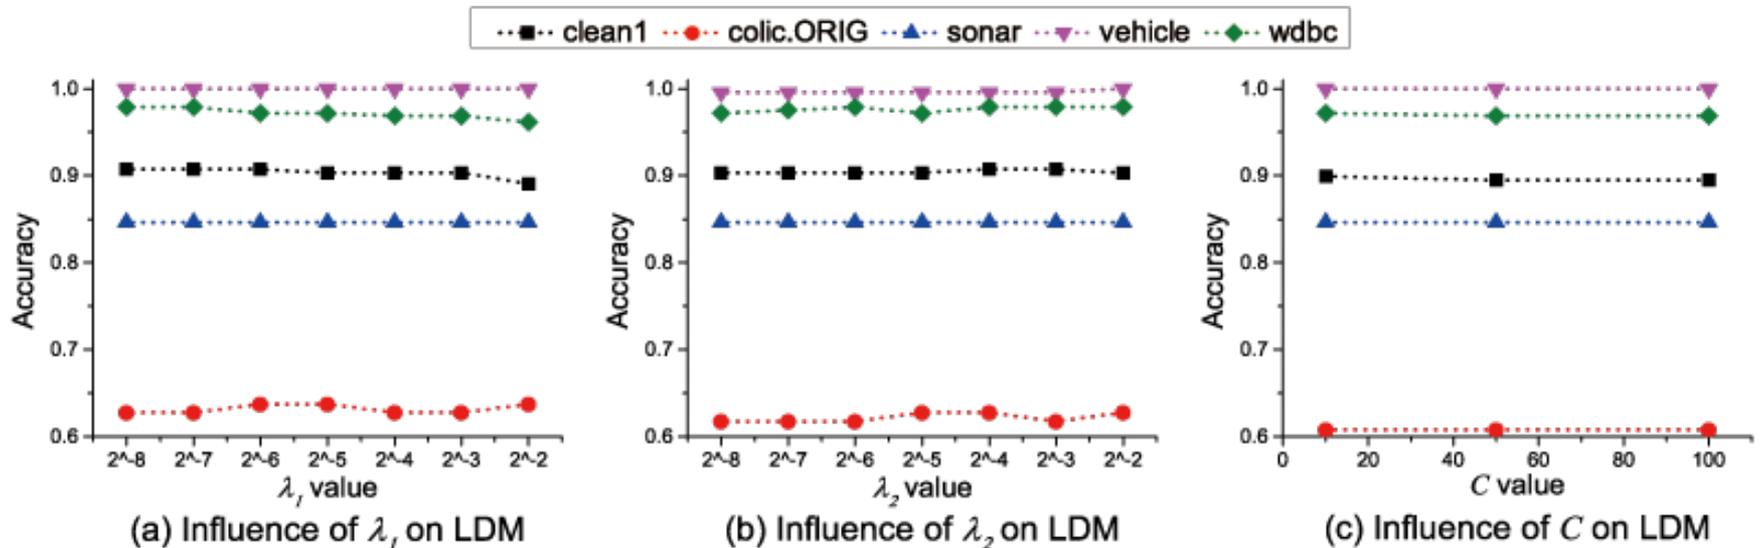
Parameter Influence

LDM has three parameters: λ_1 , λ_2 and C



Parameter Influence

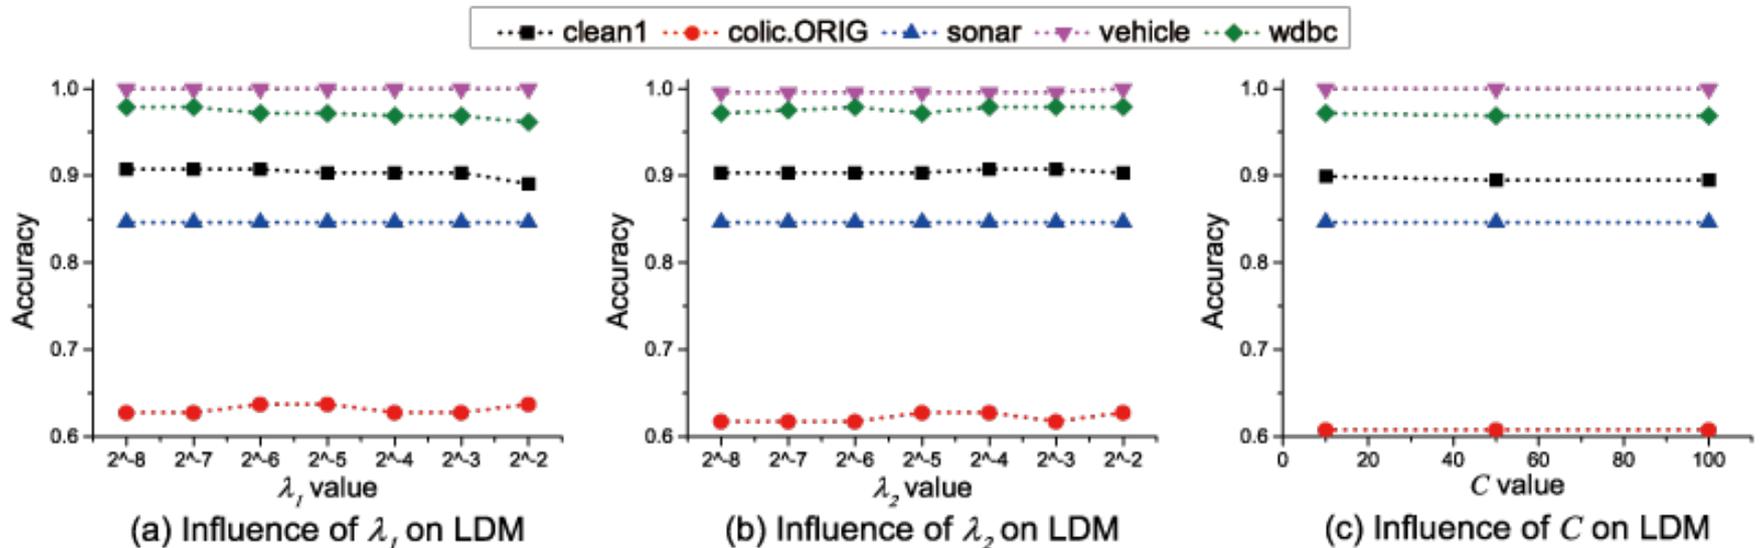
LDM has three parameters: λ_1 , λ_2 and C



- Vary one parameter and fix the other two as the value suggested by cross validation.
- λ_1 and λ_2 take value of $\{2^{-8}, \dots, 2^{-2}\}$, C takes value of $\{10, 50, 100\}$

Parameter Influence

LDM has three parameters: λ_1 , λ_2 and C



- Vary one parameter and fix the other two as the value suggested by cross validation.
- λ_1 and λ_2 take value of $\{2^{-8}, \dots, 2^{-2}\}$, C takes value of $\{10, 50, 100\}$

LDM is not very sensitive to the setting of the parameters

Outline

- Introduction
 - Our Method
 - Experiments
 - Conclusion
-

Conclusion

Main contribution:

- ✓ The first study on directly optimizing margin distribution for SVM
- ✓ We show that directly optimizing the margin distribution for SVM is beneficial.
- ✓ We propose an effective method LDM that optimizes margin distribution. (The code package for LDM method http://lamda.nju.edu.cn/code_LDM.ashx)

Future work:

- ✓ To generalize LDM to other learning settings
 - ✓
-

Conclusion

Main contribution:

- ✓ The first study on directly optimizing margin distribution for SVM
- ✓ We show that directly optimizing the margin distribution for SVM is beneficial.
- ✓ We propose an effective method LDM that optimizes margin distribution. (The code package for LDM method http://lamda.nju.edu.cn/code_LDM.ashx)

Future work:

- ✓ To generalize LDM to other learning settings
- ✓

Thanks!
