

Statistical Machine Learning (BE4M33SSU)

Lecture 4: Support Vector Machines

Czech Technical University in Prague
V.Franc

Primal SVM problem

- ◆ Find linear classifier $h(x; \mathbf{w}, b) = \text{sign}(\langle \phi(x), \mathbf{w} \rangle + b)$ by solving

$$(\mathbf{w}^*, b^*) = \underset{\mathbf{w} \in \mathbb{R}^n, b \in \mathbb{R}}{\text{argmin}} \left(\underbrace{\frac{1}{2} \|\mathbf{w}\|^2}_{\text{penalty term}} + C \underbrace{\sum_{i=1}^m \max\{0, 1 - y^i (\langle \mathbf{w}, \phi(x^i) \rangle + b)\}}_{\text{empirical error}} \right)$$

where $C > 0$ is the regularization constant.

Primal SVM problem

- ◆ Find linear classifier $h(x; \mathbf{w}, b) = \text{sign}(\langle \phi(x), \mathbf{w} \rangle + b)$ by solving

$$(\mathbf{w}^*, b^*) = \underset{\mathbf{w} \in \mathbb{R}^n, b \in \mathbb{R}}{\text{argmin}} \left(\underbrace{\frac{1}{2} \|\mathbf{w}\|^2}_{\text{penalty term}} + C \underbrace{\sum_{i=1}^m \max\{0, 1 - y^i(\langle \mathbf{w}, \phi(x^i) \rangle + b)\}}_{\text{empirical error}} \right)$$

where $C > 0$ is the regularization constant.

- ◆ It can be re-formulated as a convex *quadratic program*

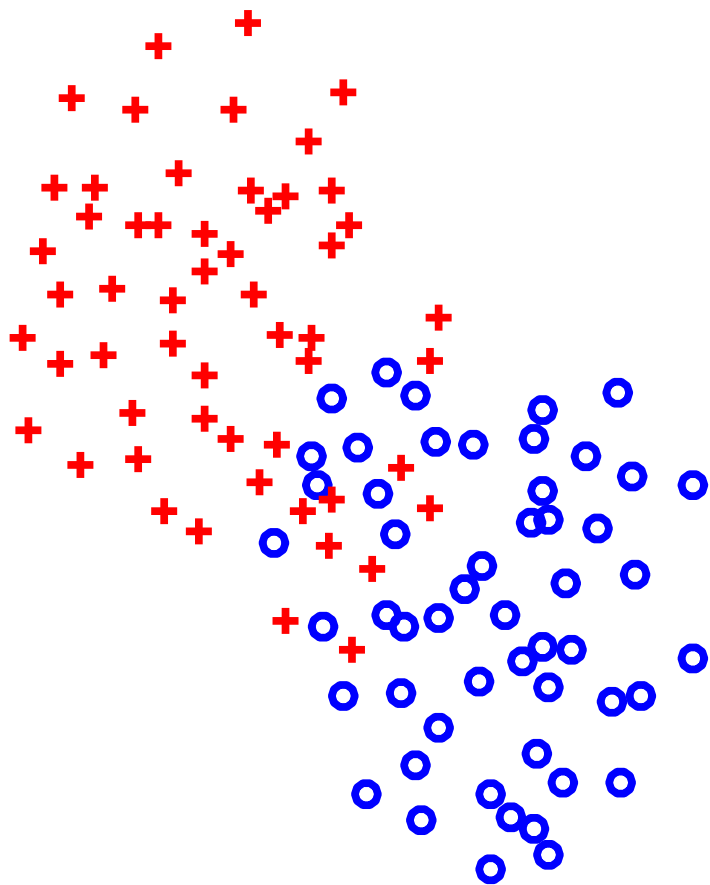
$$(\mathbf{w}^*, b^*, \boldsymbol{\xi}^*) = \underset{\substack{(\mathbf{w}, b) \in \mathbb{R}^{n+1} \\ \boldsymbol{\xi} \in \mathbb{R}^m}}{\text{argmin}} \left(\frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^m \xi_i \right)$$

subject to

$$\begin{aligned} y^i(\langle \mathbf{w}, \phi(x^i) \rangle + b) &\geq 1 - \xi_i, & i \in \{1, \dots, m\} \\ \xi_i &\geq 0, & i \in \{1, \dots, m\} \end{aligned}$$

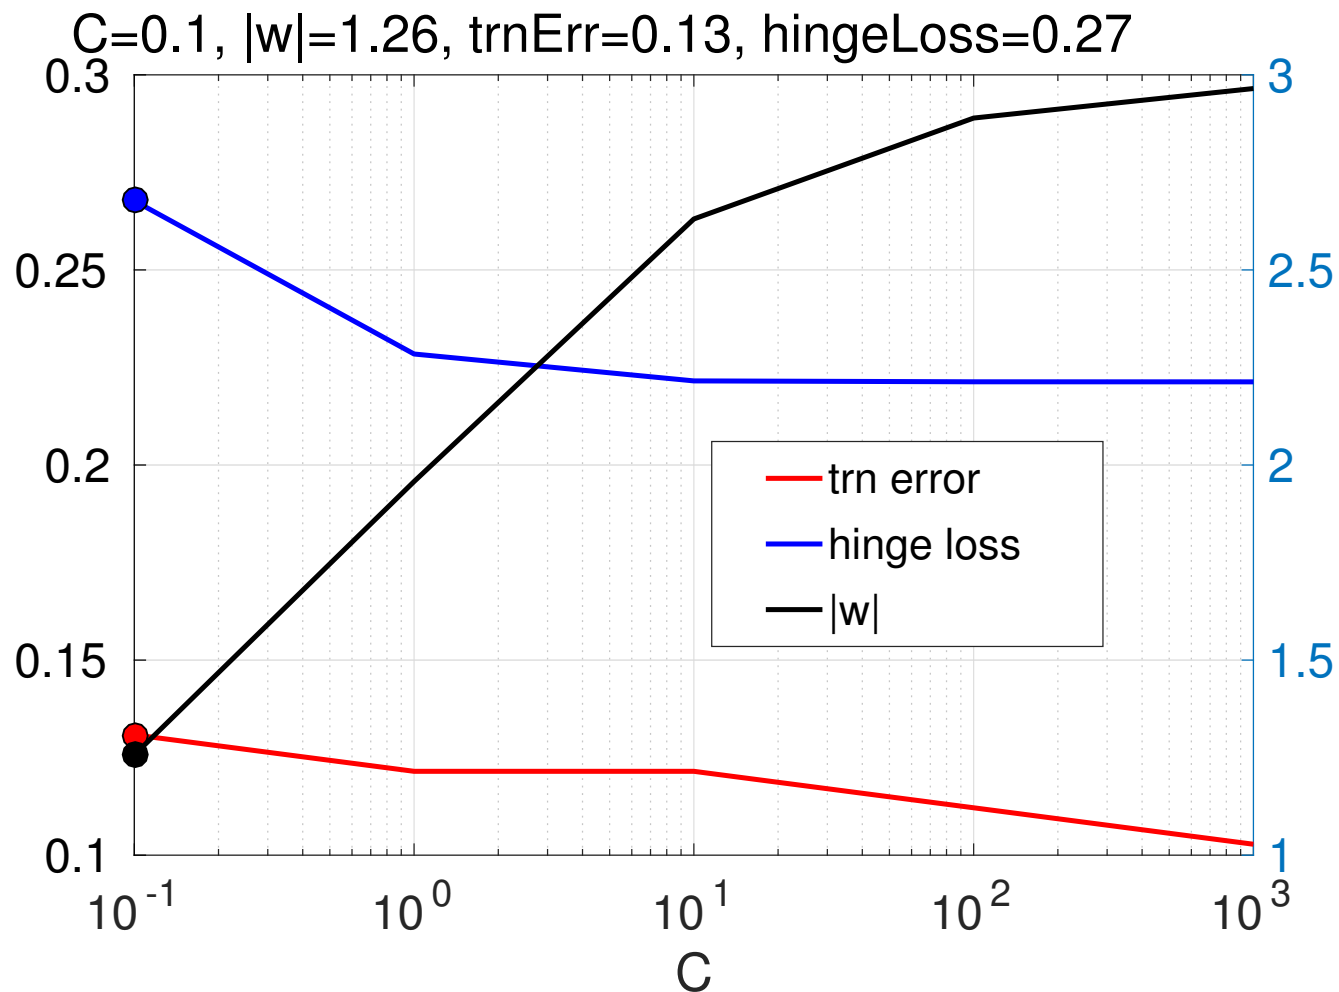
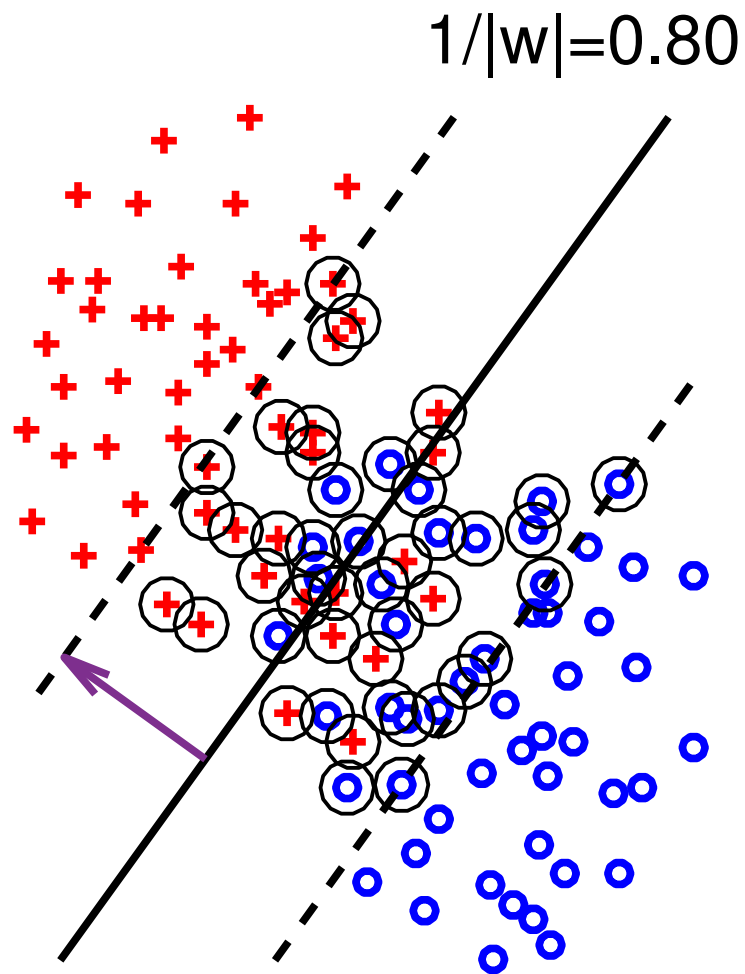
Example: Primal SVM problem

$$(\mathbf{w}^*, b^*) = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^n, b \in \mathbb{R}} \left(\underbrace{\frac{1}{2} \|\mathbf{w}\|^2}_{\text{penalty term}} + C \underbrace{\sum_{i=1}^m \max\{0, 1 - y^i (\langle \mathbf{w}, \phi(x^i) \rangle + b)\}}_{\text{empirical error}} \right)$$



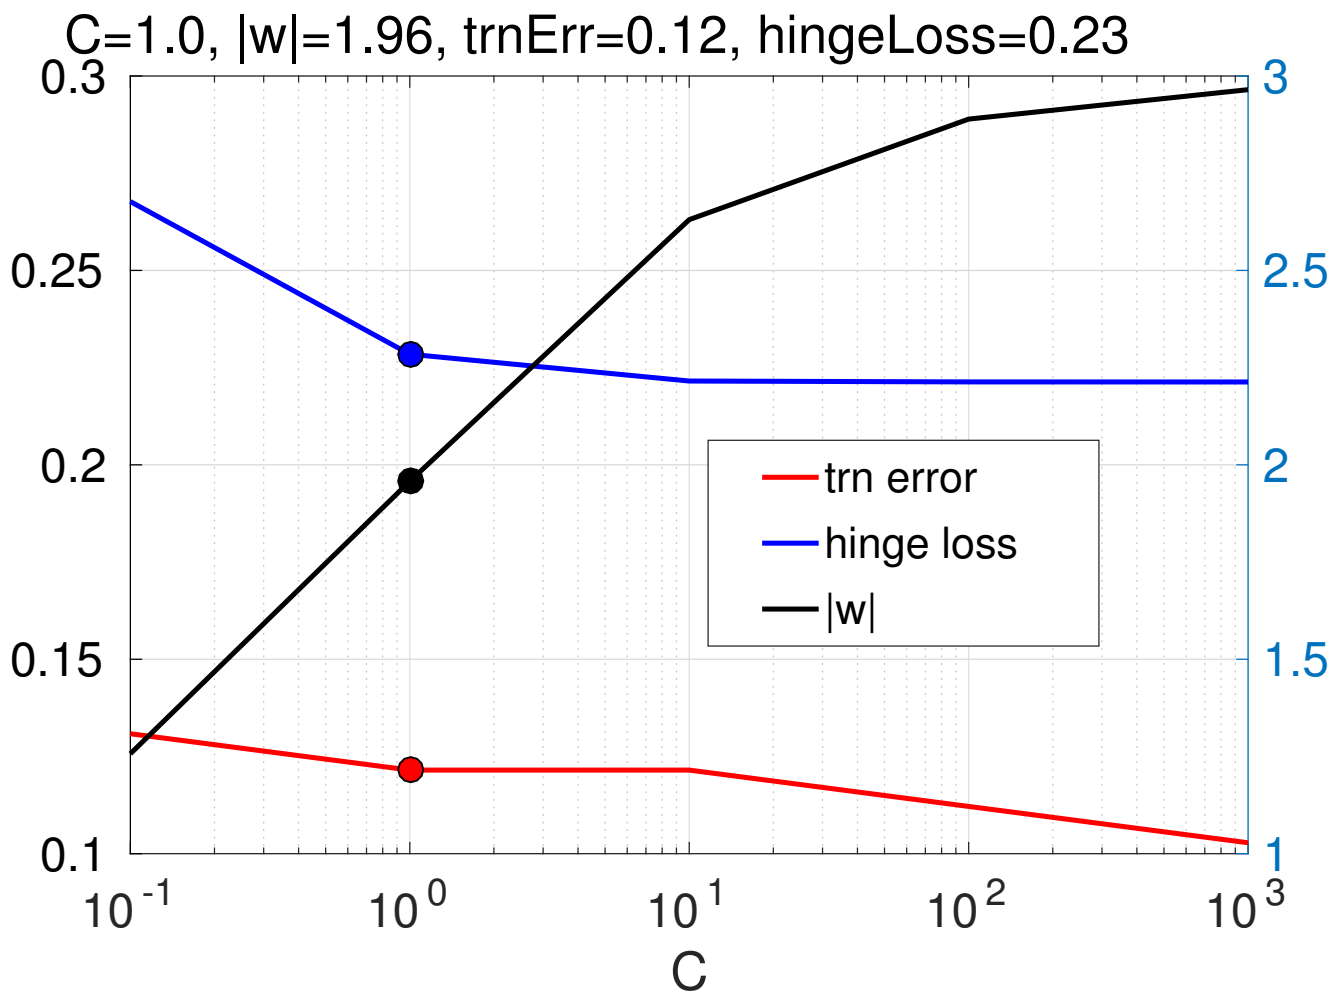
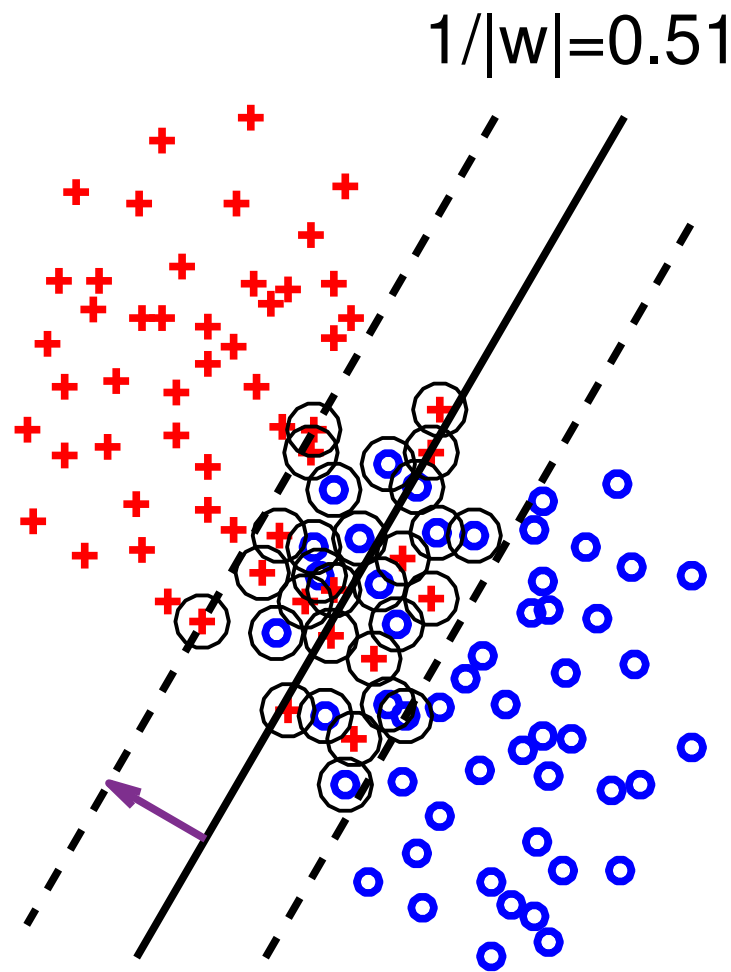
Example: Primal SVM problem

$$(\mathbf{w}^*, b^*) = \underset{\mathbf{w} \in \mathbb{R}^n, b \in \mathbb{R}}{\operatorname{argmin}} \left(\underbrace{\frac{1}{2} \|\mathbf{w}\|^2}_{\text{penalty term}} + C \underbrace{\sum_{i=1}^m \max\{0, 1 - y^i (\langle \mathbf{w}, \phi(x^i) \rangle + b)\}}_{\text{empirical error}} \right)$$



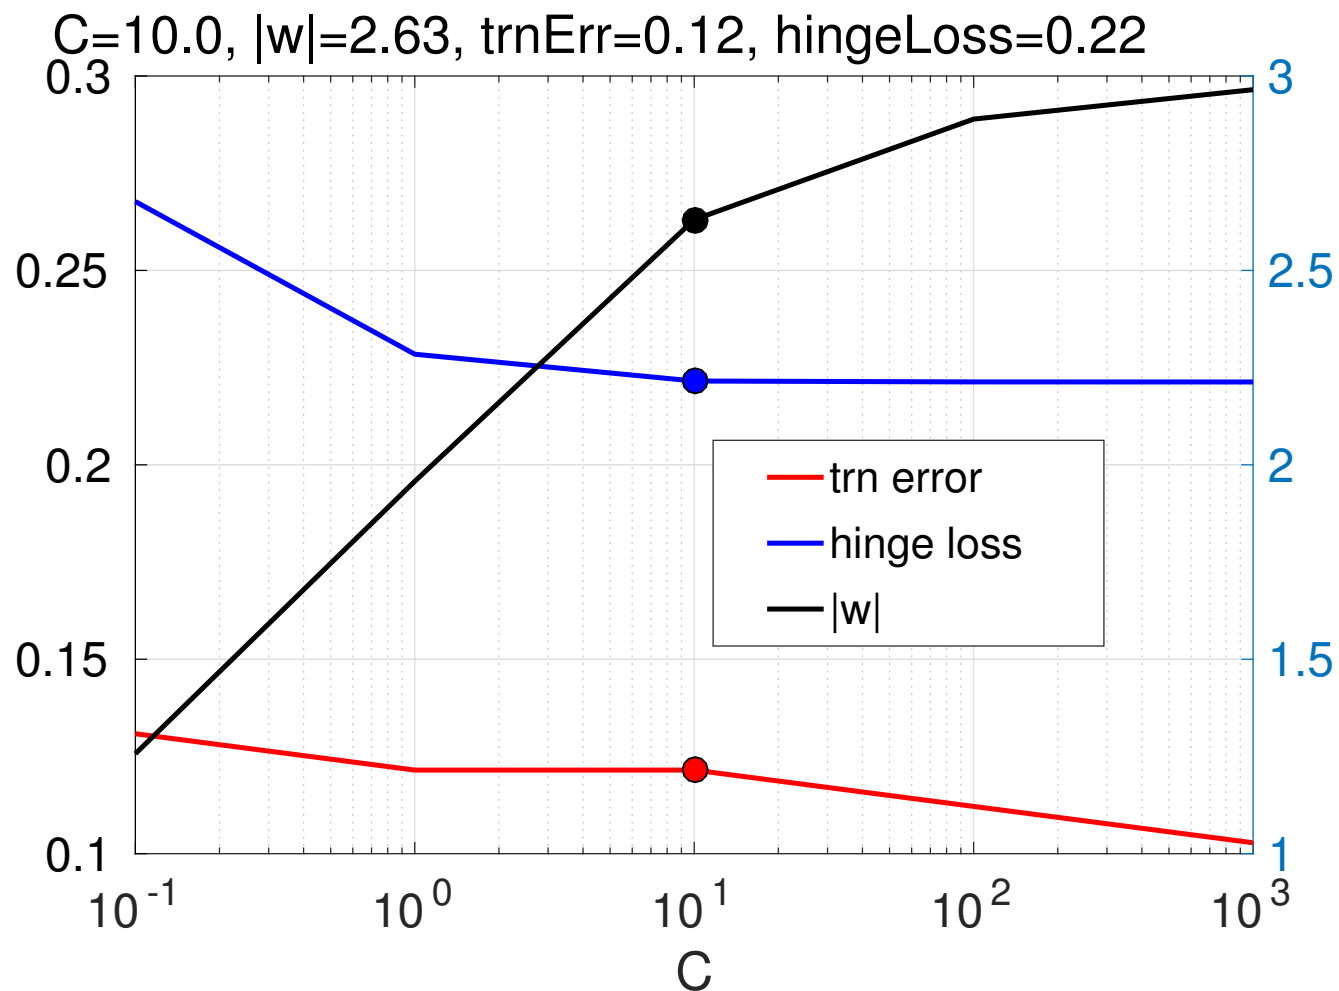
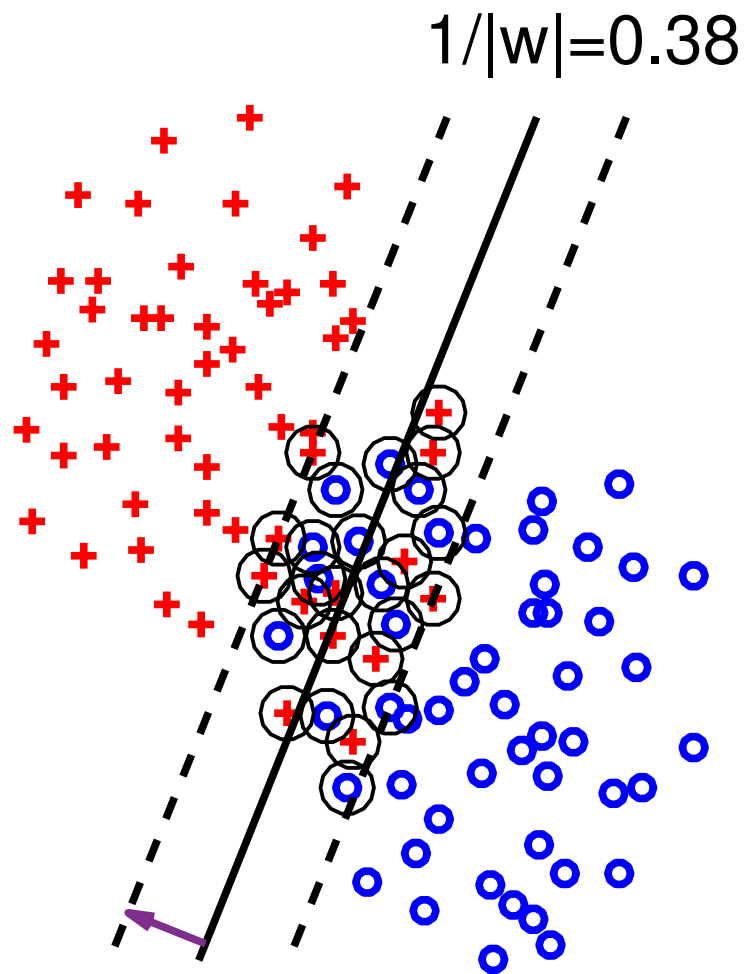
Example: Primal SVM problem

$$(\mathbf{w}^*, b^*) = \underset{\mathbf{w} \in \mathbb{R}^n, b \in \mathbb{R}}{\operatorname{argmin}} \left(\underbrace{\frac{1}{2} \|\mathbf{w}\|^2}_{\text{penalty term}} + C \underbrace{\sum_{i=1}^m \max\{0, 1 - y^i(\langle \mathbf{w}, \phi(x^i) \rangle + b)\}}_{\text{empirical error}} \right)$$



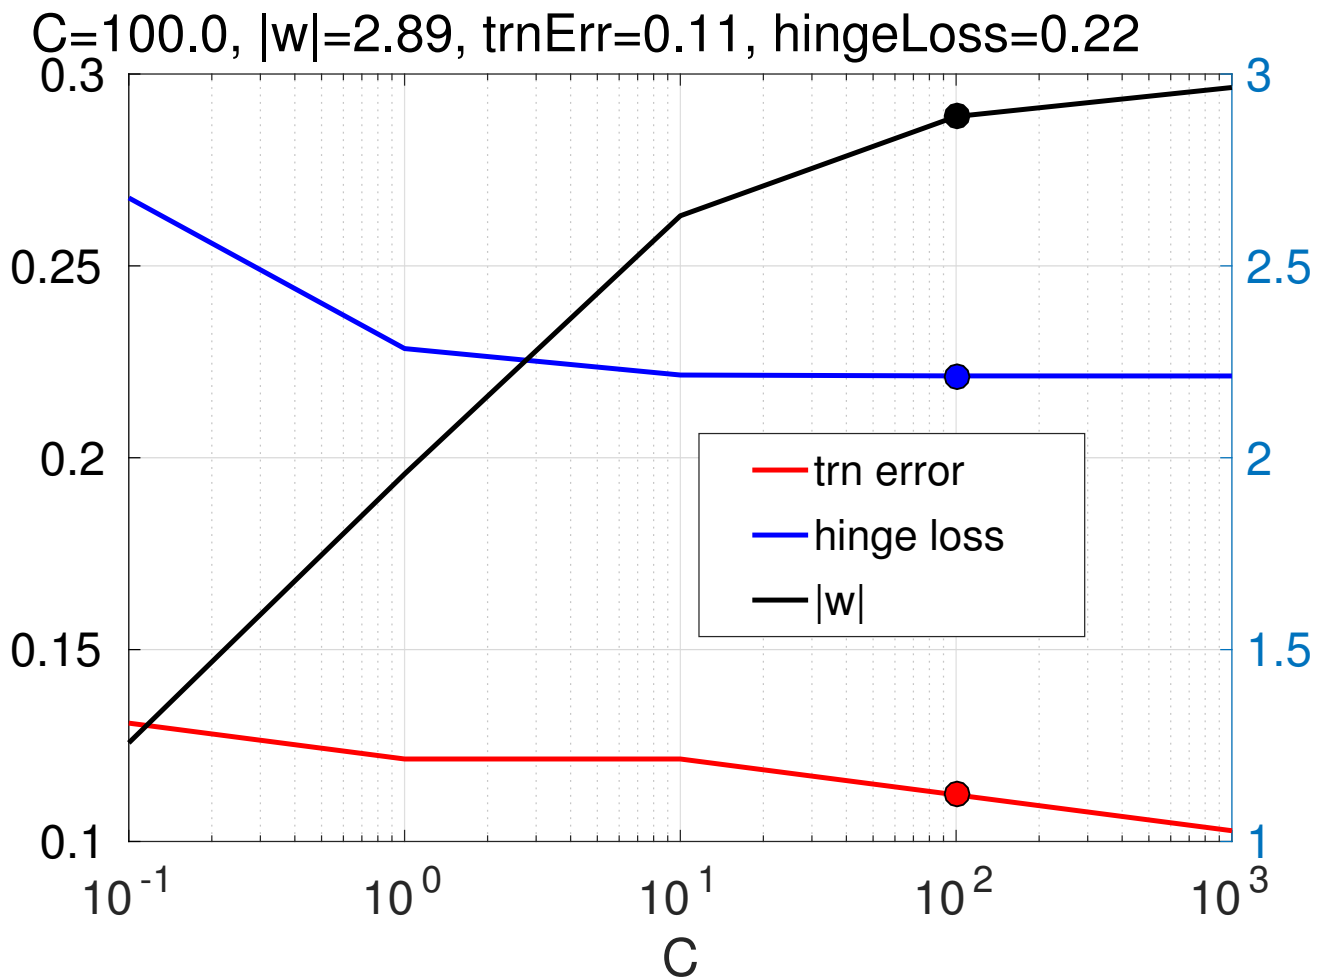
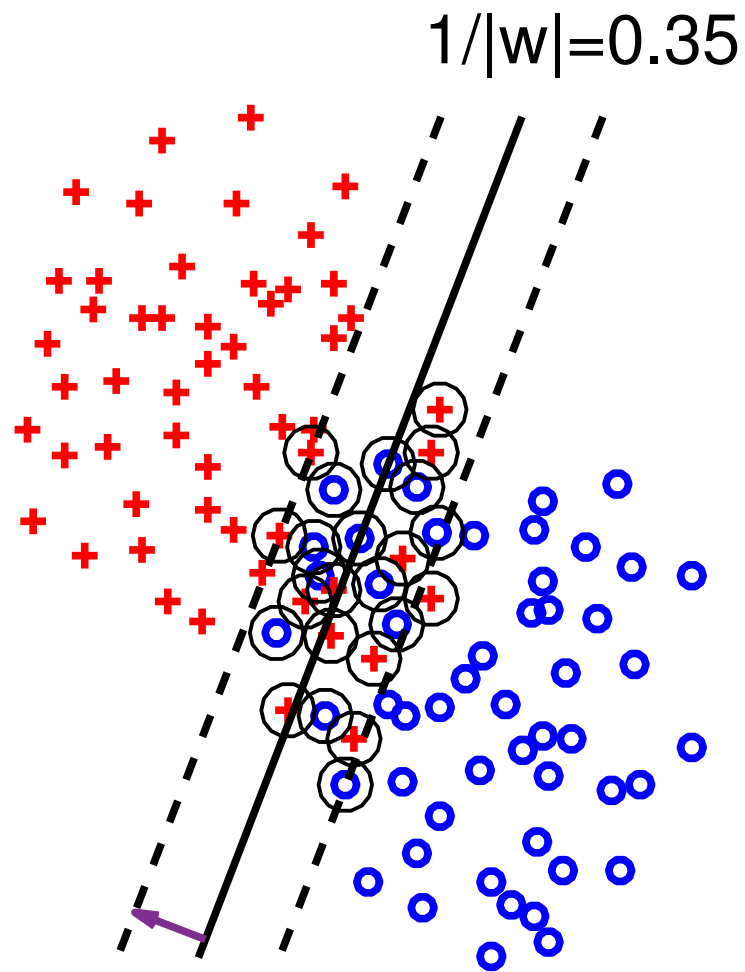
Example: Primal SVM problem

$$(\mathbf{w}^*, b^*) = \underset{\mathbf{w} \in \mathbb{R}^n, b \in \mathbb{R}}{\operatorname{argmin}} \left(\underbrace{\frac{1}{2} \|\mathbf{w}\|^2}_{\text{penalty term}} + C \underbrace{\sum_{i=1}^m \max\{0, 1 - y^i(\langle \mathbf{w}, \phi(x^i) \rangle + b)\}}_{\text{empirical error}} \right)$$



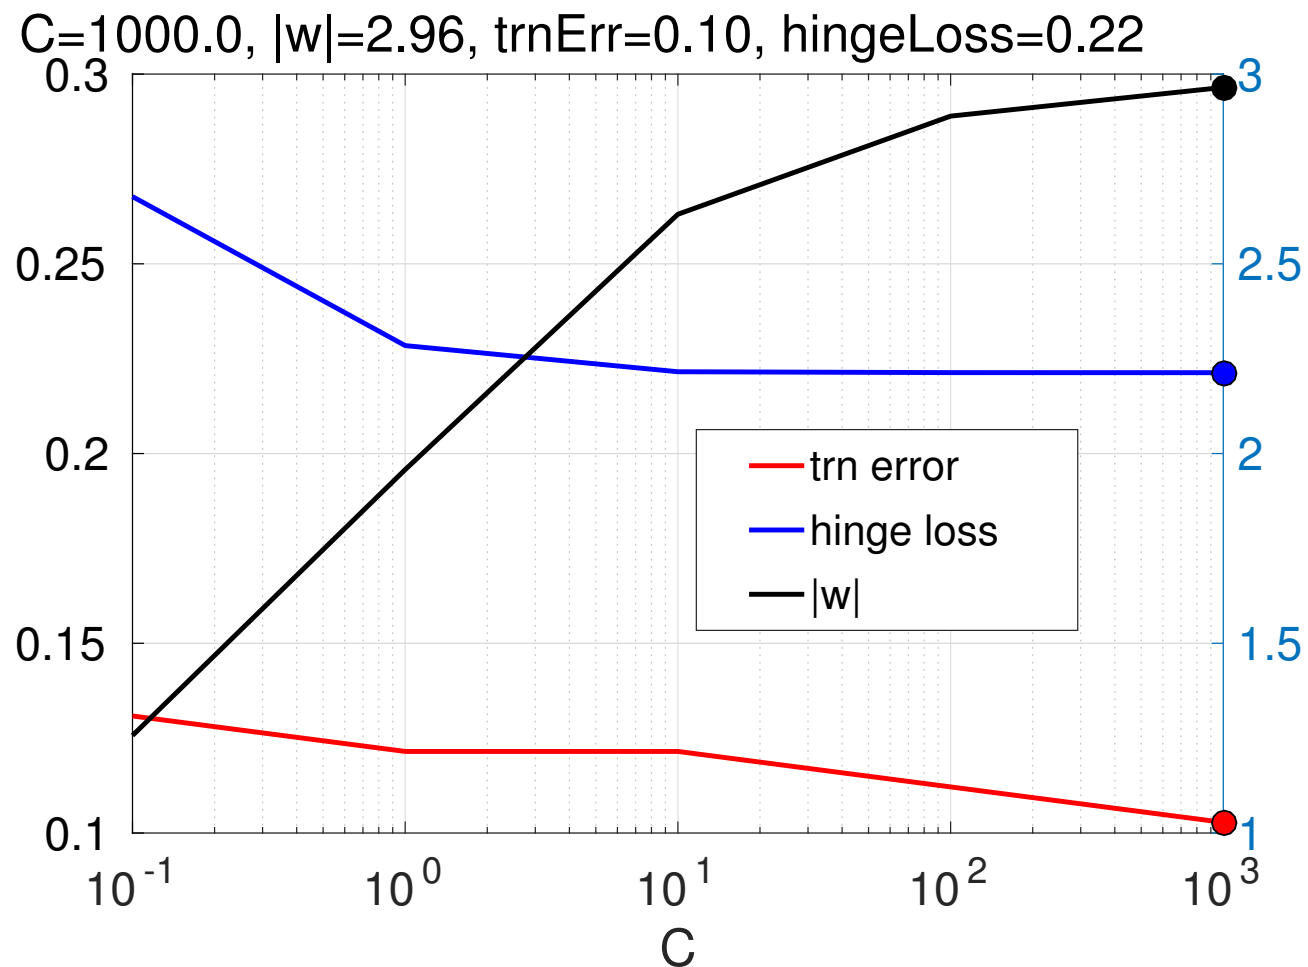
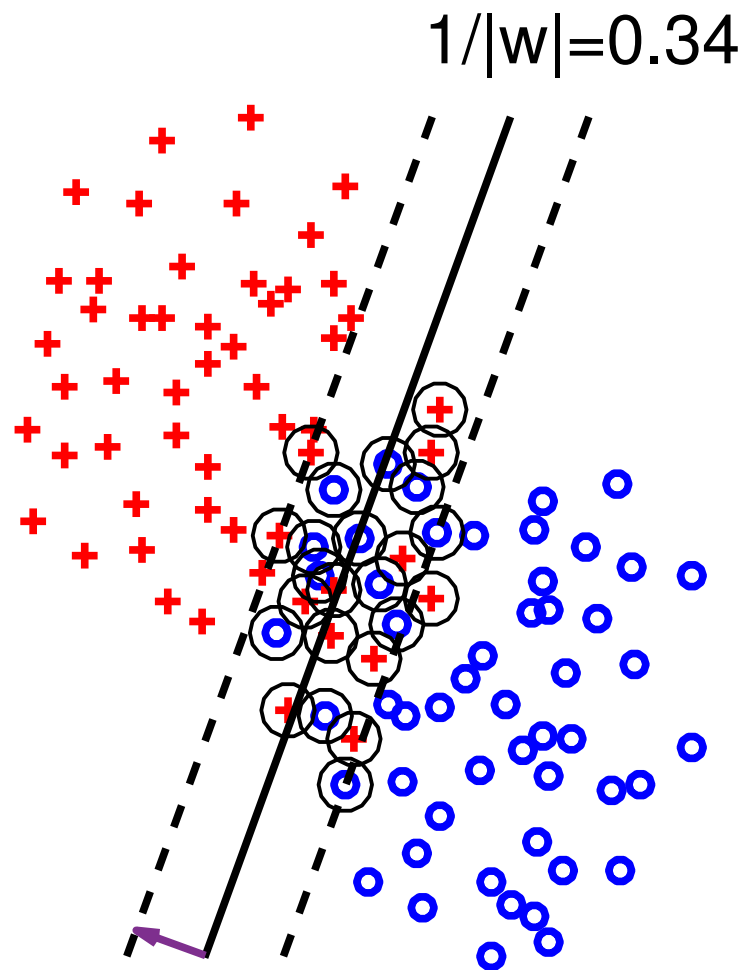
Example: Primal SVM problem

$$(\mathbf{w}^*, b^*) = \underset{\mathbf{w} \in \mathbb{R}^n, b \in \mathbb{R}}{\operatorname{argmin}} \left(\underbrace{\frac{1}{2} \|\mathbf{w}\|^2}_{\text{penalty term}} + C \underbrace{\sum_{i=1}^m \max\{0, 1 - y^i (\langle \mathbf{w}, \phi(x^i) \rangle + b)\}}_{\text{empirical error}} \right)$$



Example: Primal SVM problem

$$(\mathbf{w}^*, b^*) = \underset{\mathbf{w} \in \mathbb{R}^n, b \in \mathbb{R}}{\operatorname{argmin}} \left(\underbrace{\frac{1}{2} \|\mathbf{w}\|^2}_{\text{penalty term}} + C \underbrace{\sum_{i=1}^m \max\{0, 1 - y^i(\langle \mathbf{w}, \phi(x^i) \rangle + b)\}}_{\text{empirical error}} \right)$$



From Primal SVM to Dual SVM problem

- ◆ Lagrangian of the primal SVM problem:

$$\begin{aligned}
 L(\mathbf{w}, b, \boldsymbol{\xi}, \boldsymbol{\alpha}, \boldsymbol{\mu}) = & \underbrace{\frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^m \xi_i}_{\text{original objective}} \\
 & - \underbrace{\sum_{i=1}^m \alpha_i (y^i (\langle \mathbf{w}, \phi(x^i) \rangle + b) - 1 + \xi_i) - \sum_{i=1}^m \mu_i \xi_i}_{\text{constraint violation penalty}}
 \end{aligned}$$

- ◆ Strong duality:

$$\underbrace{\min_{\substack{\mathbf{w} \in \mathbb{R}^n \\ b \in \mathbb{R} \\ \boldsymbol{\xi} \in \mathbb{R}^m}} \max_{\substack{\boldsymbol{\alpha} \in \mathbb{R}_+^m \\ \boldsymbol{\mu} \in \mathbb{R}_+^m}} L(\mathbf{w}, b, \boldsymbol{\xi}, \boldsymbol{\alpha}, \boldsymbol{\mu})}_{\text{primal problem}} = \underbrace{\max_{\substack{\boldsymbol{\alpha} \in \mathbb{R}_+^m \\ \boldsymbol{\mu} \in \mathbb{R}_+^m}} \min_{\substack{\mathbf{w} \in \mathbb{R}^n \\ b \in \mathbb{R} \\ \boldsymbol{\xi} \in \mathbb{R}^m}} L(\mathbf{w}, b, \boldsymbol{\xi}, \boldsymbol{\alpha}, \boldsymbol{\mu})}_{\text{dual problem}}$$

Dual SVM problem

- ◆ The dual SVM formulation is a convex quadratic program

$$\begin{aligned} \boldsymbol{\alpha}^* = \operatorname{argmax}_{\boldsymbol{\alpha} \in \mathbb{R}^m} & \left(\sum_{i=1}^m \alpha_i - \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m \alpha_i \alpha_j y^i y^j \langle \boldsymbol{\phi}(x^i), \boldsymbol{\phi}(x^j) \rangle \right) \\ \text{s.t.} & \sum_{i=1}^m \alpha_i y^i = 0, \quad 0 \leq \alpha_i \leq C, \quad i \in \{1, \dots, m\} \end{aligned}$$

Dual SVM problem

- ◆ The dual SVM formulation is a convex quadratic program

$$\begin{aligned}
 \boldsymbol{\alpha}^* &= \operatorname{argmax}_{\boldsymbol{\alpha} \in \mathbb{R}^m} \left(\sum_{i=1}^m \alpha_i - \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m \alpha_i \alpha_j y^i y^j \langle \boldsymbol{\phi}(x^i), \boldsymbol{\phi}(x^j) \rangle \right) \\
 \text{s.t.} \quad &\sum_{i=1}^m \alpha_i y^i = 0, \quad 0 \leq \alpha_i \leq C, \quad i \in \{1, \dots, m\}
 \end{aligned}$$

- ◆ The primal variables (\boldsymbol{w}, b) are obtained from the dual variables $\boldsymbol{\alpha}$ by

$$\boldsymbol{w} = \sum_{i=1}^m y^i \boldsymbol{\phi}(x^i) \alpha_i = \sum_{i \in \mathcal{I}_{\text{SV}}} y^i \boldsymbol{\phi}(x^i) \alpha_i$$

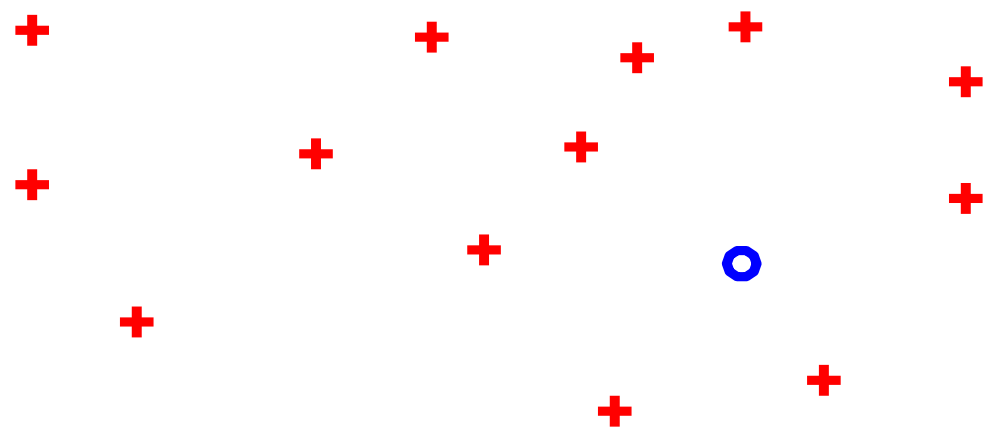
$$b = y^i - \langle \boldsymbol{w}, \boldsymbol{\phi}(x^i) \rangle, \quad \forall i \in \mathcal{I}_{\text{SV}}^b = \{j \mid 0 < \alpha_j < C\}$$

- ◆ $\boldsymbol{\alpha}$ is sparse; \boldsymbol{w} is lin. combination of Support Vectors $\mathcal{I}_{\text{SV}} = \{j \mid \alpha_j > 0\}$

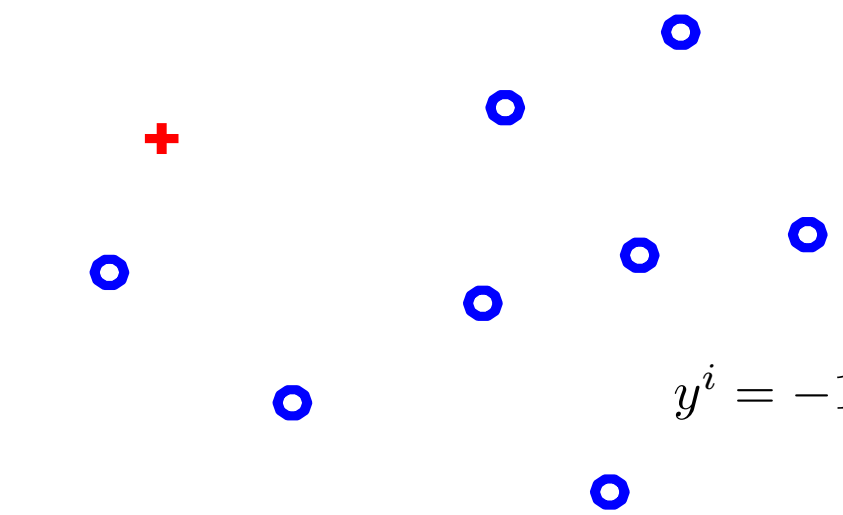
Example: SVM classifier

$$f(x) = \langle \mathbf{w}, \phi(x) \rangle + b = \langle \underbrace{\sum_{i=1}^m y^i \alpha_i \phi(x^i)}_{\mathbf{w}}, \phi(x) \rangle + b$$

$y^i = +1$

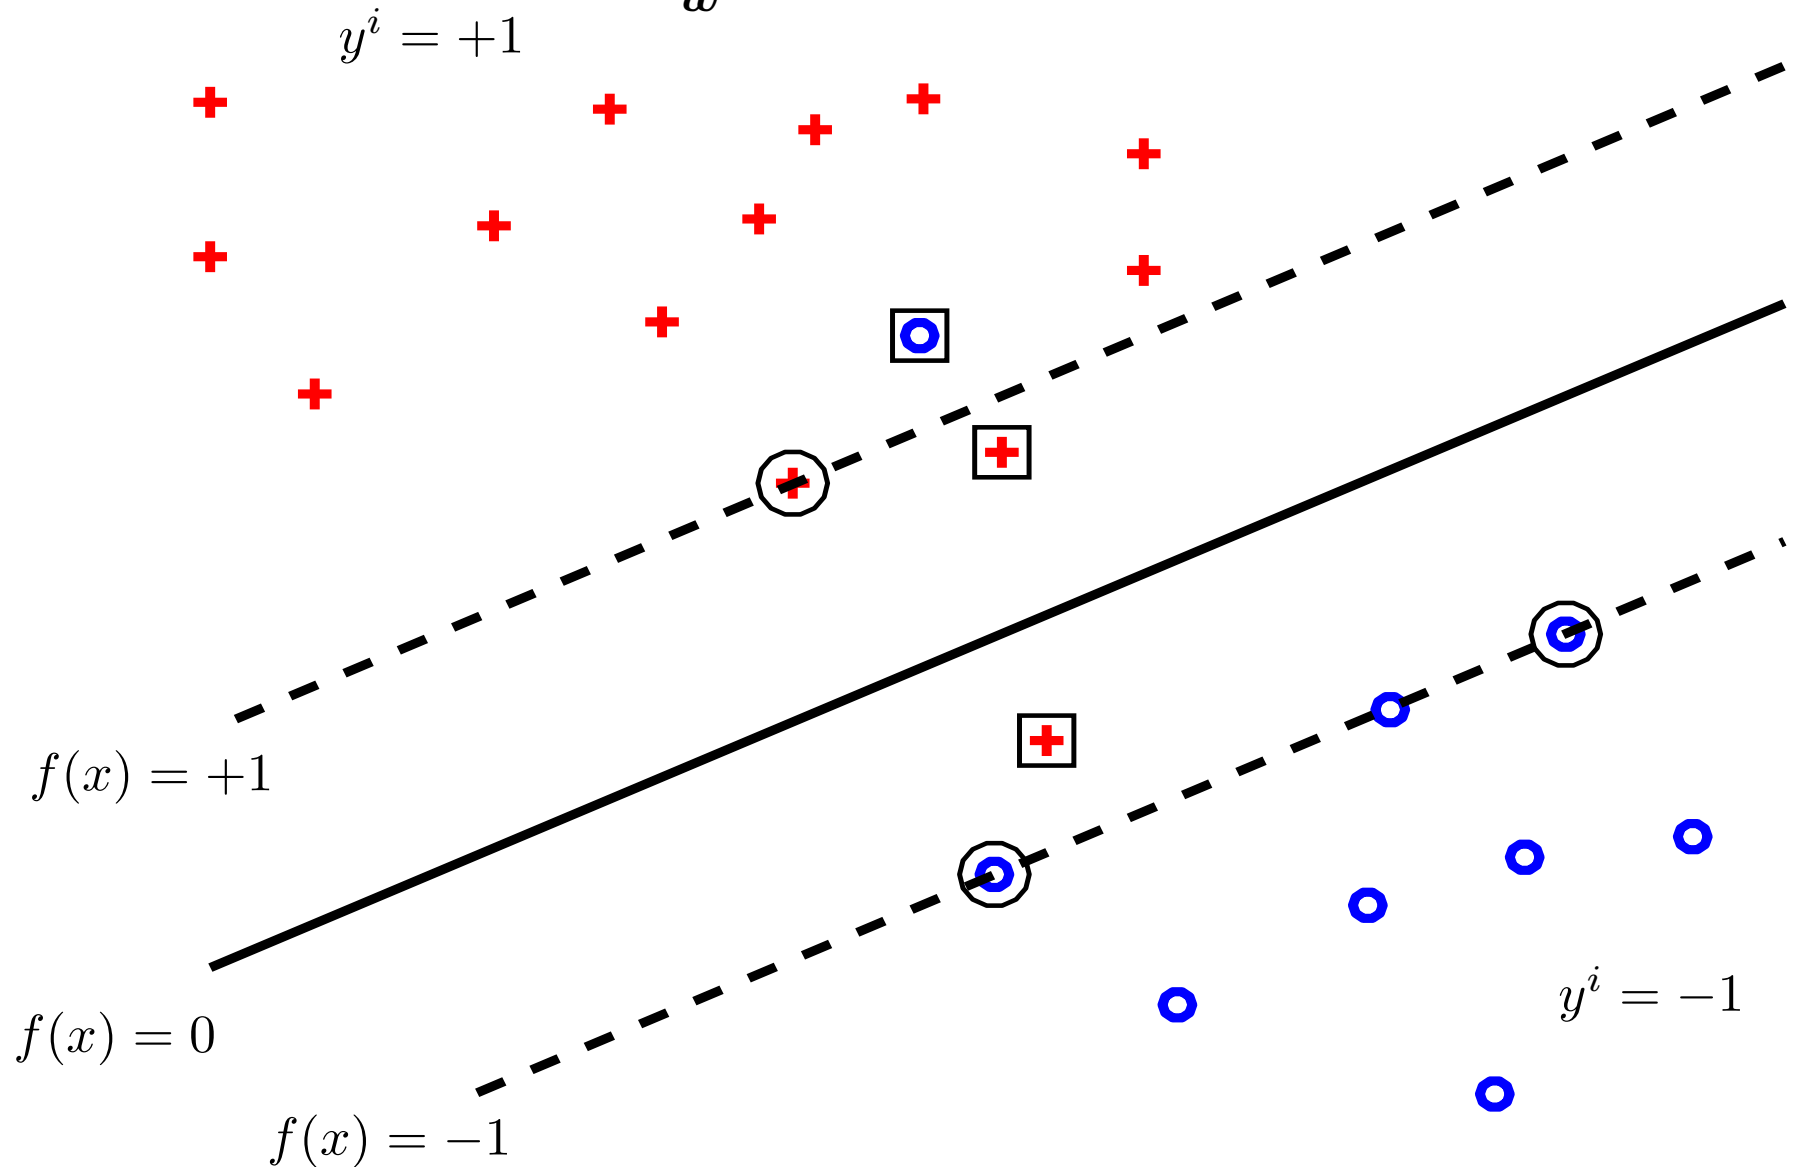


$y^i = -1$



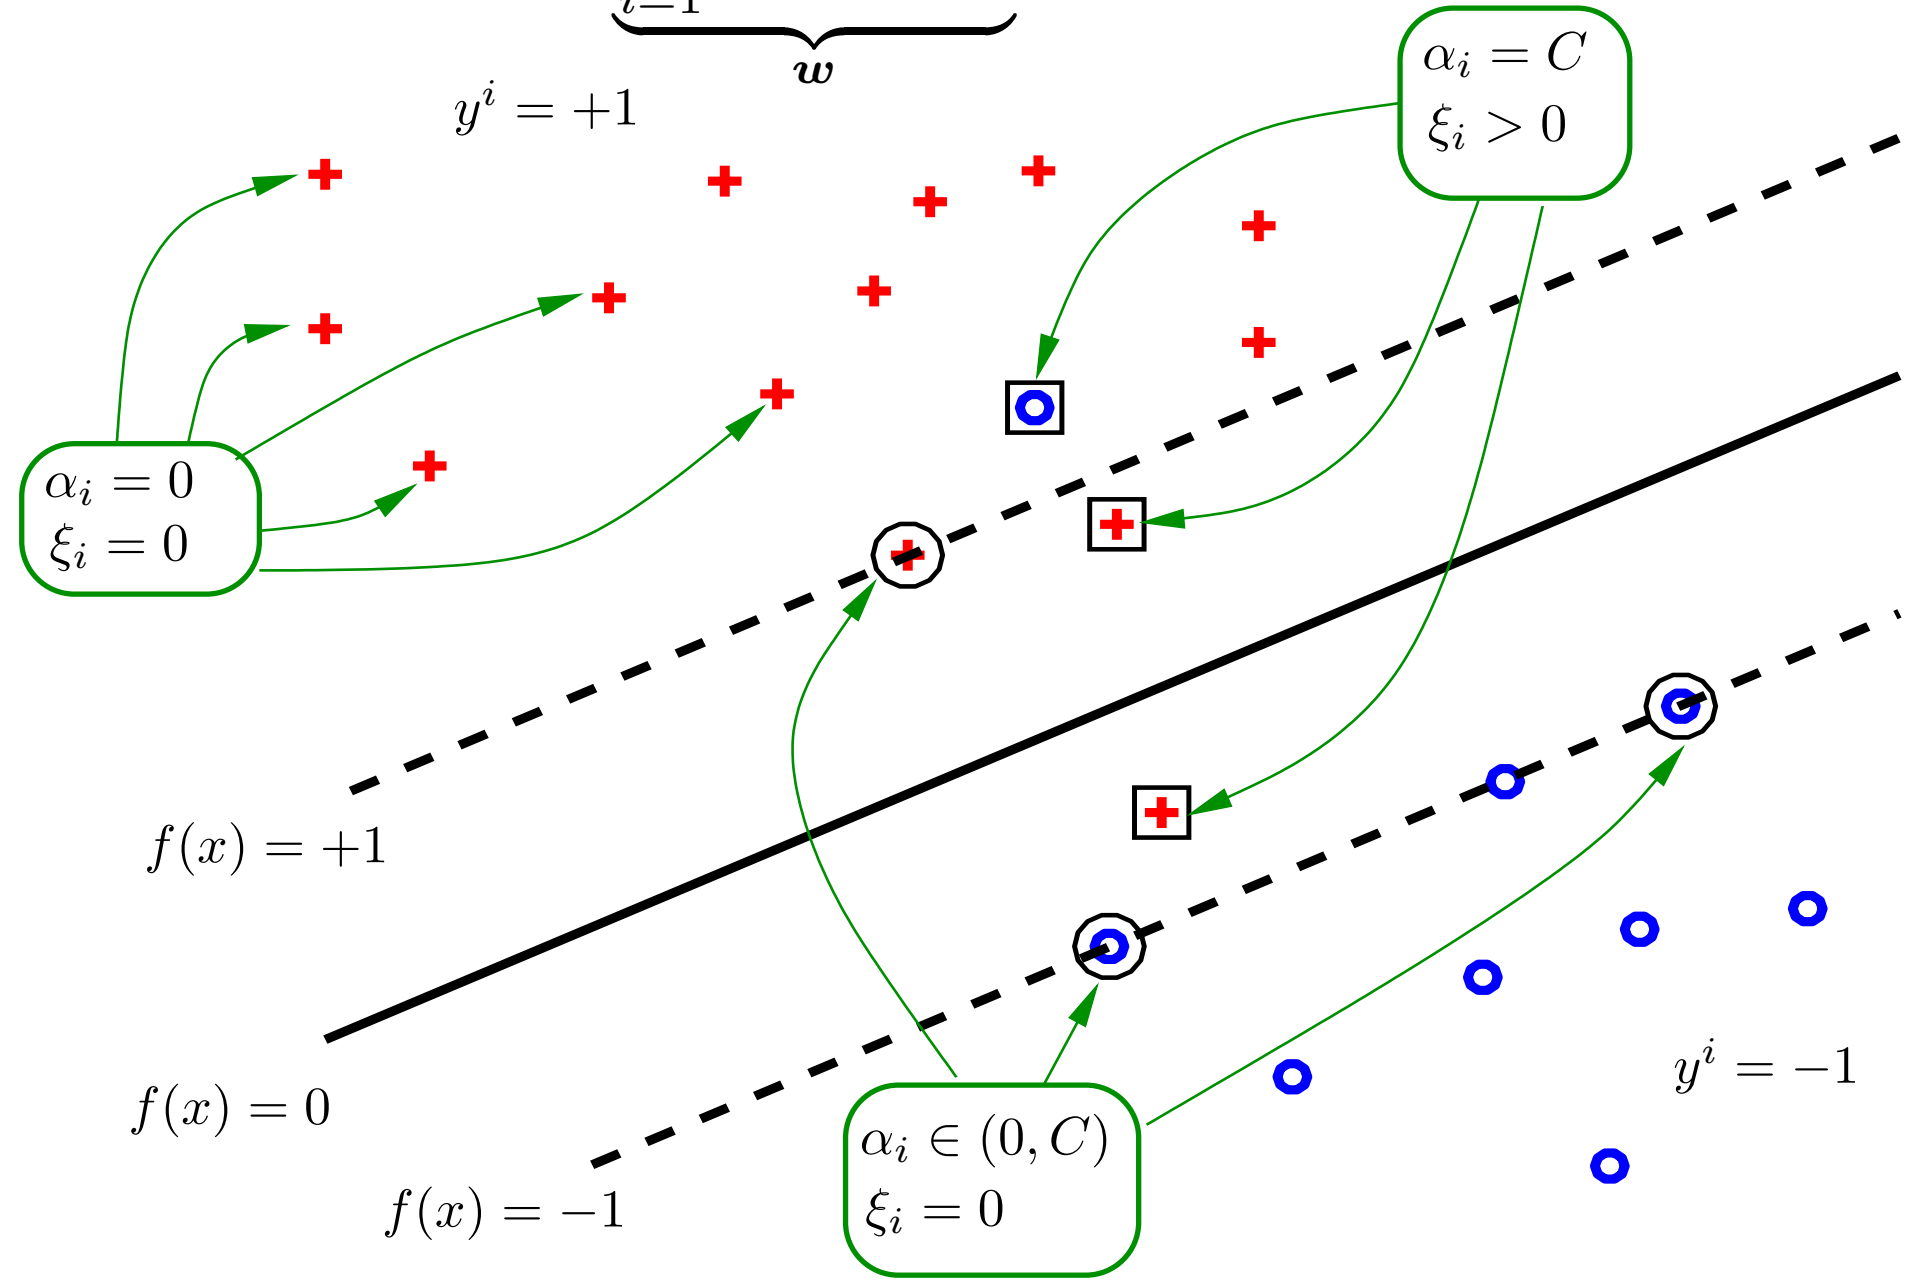
Example: SVM classifier

$$f(x) = \langle \mathbf{w}, \phi(x) \rangle + b = \langle \underbrace{\sum_{i=1}^m y^i \alpha_i \phi(x^i)}_{\mathbf{w}}, \phi(x) \rangle + b$$



Example: SVM classifier

$$f(x) = \langle \mathbf{w}, \phi(x) \rangle + b = \langle \underbrace{\sum_{i=1}^m y^i \alpha_i \phi(x^i)}_{\mathbf{w}}, \phi(x) \rangle + b$$



Kernel SVM

- ◆ The SVM algorithm requires observations in terms of dot products only:

Learning: Given $\mathcal{T}^m = \{(x^i, y^i) \in \mathcal{X} \times \{-1, +1\} \mid i = 1, \dots, m\}$, solve

$$\alpha^* = \operatorname{argmax}_{\alpha \in \mathbb{R}^m} \left(\sum_{i=1}^m \alpha_i - \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m \alpha_i \alpha_j y^i y^j \langle \phi(x^i), \phi(x^j) \rangle \right)$$

s.t. $\sum_{i=1}^m \alpha_i y^i = 0, \quad 0 \leq \alpha_i \leq C, i \in \{1, \dots, m\}$

Prediction:

$$h(x; \alpha, b) = \operatorname{sign}(\langle \mathbf{w}, \phi(x) \rangle + b) = \operatorname{sign} \left(\sum_{i \in \mathcal{I}_{\text{sv}}} y^i \alpha_i \langle \phi(x^i), \phi(x) \rangle + b \right)$$

- ◆ Given a feature map $\phi: \mathcal{X} \rightarrow \mathbb{R}^n$, define kernel function $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$

$$k(x, x') = \langle \phi(x), \phi(x') \rangle$$

Kernel SVM

- ◆ The SVM algorithm requires observations in terms of dot products only:

Learning: Given $\mathcal{T}^m = \{(x, y) \in \mathcal{X} \times \{-1, +1\} \mid i = 1, \dots, m\}$, solve

$$\alpha^* = \operatorname{argmax}_{\alpha \in \mathbb{R}^m} \left(\sum_{i=1}^m \alpha_i - \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m \alpha_i \alpha_j y^i y^j k(x^i, x^j) \right)$$

$$\text{s.t.} \quad \sum_{i=1}^m \alpha_i y^i = 0, \quad 0 \leq \alpha_i \leq C, i \in \{1, \dots, m\}$$

Prediction:

$$h(x; \alpha, b) = \operatorname{sign}(\langle \mathbf{w}, \phi(x) \rangle + b) = \operatorname{sign} \left(\sum_{i \in \mathcal{I}_{\text{SV}}} y^i \alpha_i k(x^i, x) + b \right)$$

- ◆ Given a feature map $\phi: \mathcal{X} \rightarrow \mathbb{R}^n$, define kernel function $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$

$$k(x, x') = \langle \phi(x), \phi(x') \rangle$$

Homogeneous polynomial kernel of degree 2

- ◆ Consider quadratic function of d -dimensional inputs $\mathbf{x} \in \mathcal{X} = \mathbb{R}^d$

$$\begin{aligned}
 f(\mathbf{x}) &= x_1^2 w_1 + \cdots + x_d^2 w_d + \sqrt{2} x_1 x_2 w_{d+1} + \cdots + \sqrt{2} x_{d-1} x_d w_n + b \\
 &= \langle \mathbf{w}, \phi(\mathbf{x}) \rangle + b
 \end{aligned}$$

where $\mathbf{w} \in \mathbb{R}^n$, $n = \frac{(d+1)d}{2}$, and

$$\phi(\mathbf{x}) = (x_1^2, \dots, x_d^2, \sqrt{2}x_1x_2, \dots, \sqrt{2}x_{d-1}x_d)$$

- ◆ The dot product of $\phi(\mathbf{x})$ and $\phi(\mathbf{x}')$ can be via kernel

$$k(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle = \langle \mathbf{x}, \mathbf{x}' \rangle^2$$

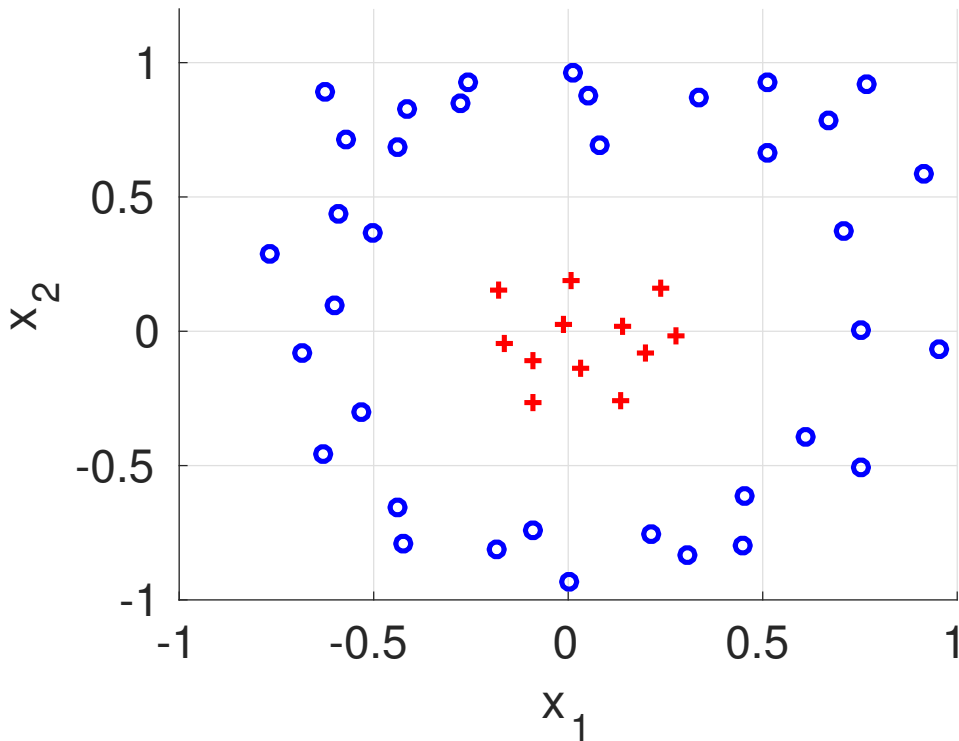
- ◆ In case of $d = 2$, we have

$$\begin{aligned}
 k(\mathbf{x}, \mathbf{x}') &= (x_1 x'_1 + x_2 x'_2)^2 = x_1^2 x_1'^2 + 2 x_1 x_2 x'_1 x'_2 + x_2^2 x_2'^2 \\
 &= \left\langle (x_1^2, \sqrt{2} x_1 x_2, x_2^2), (x_1'^2, \sqrt{2} x'_1 x'_2, x_2'^2) \right\rangle
 \end{aligned}$$

Example: Homogeneous polynomial kernel of degree 2

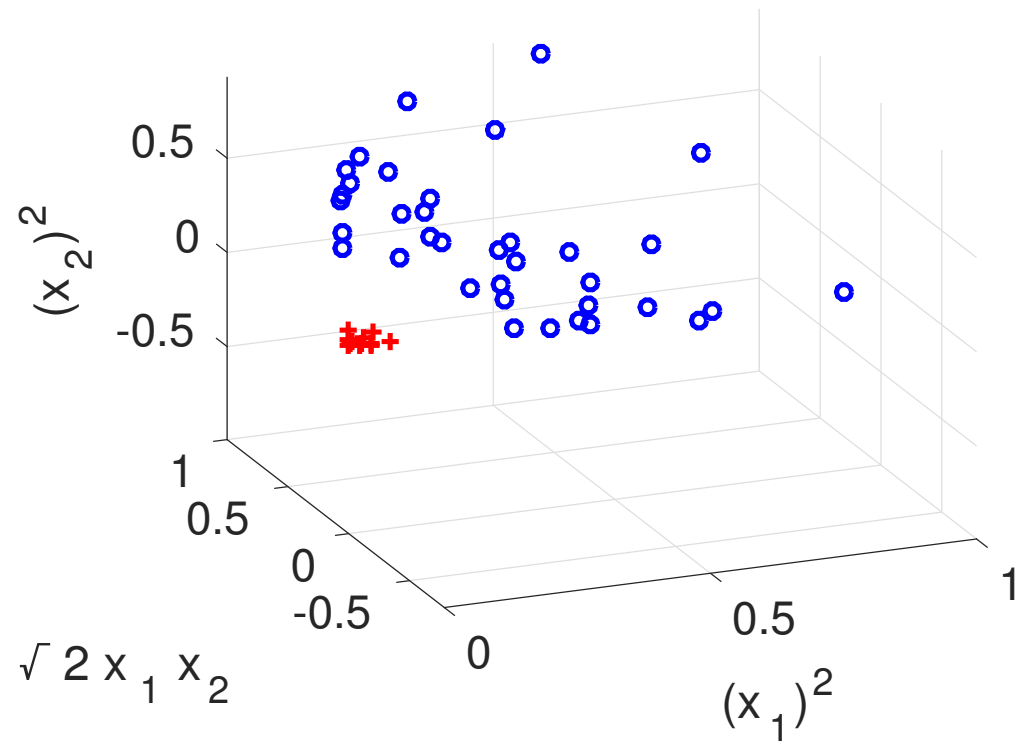
observations

$$\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$$



features

$$\phi(\mathbf{x}) = (x_1^2, \sqrt{2}x_1x_2, x_2^2) \in \mathbb{R}^3$$

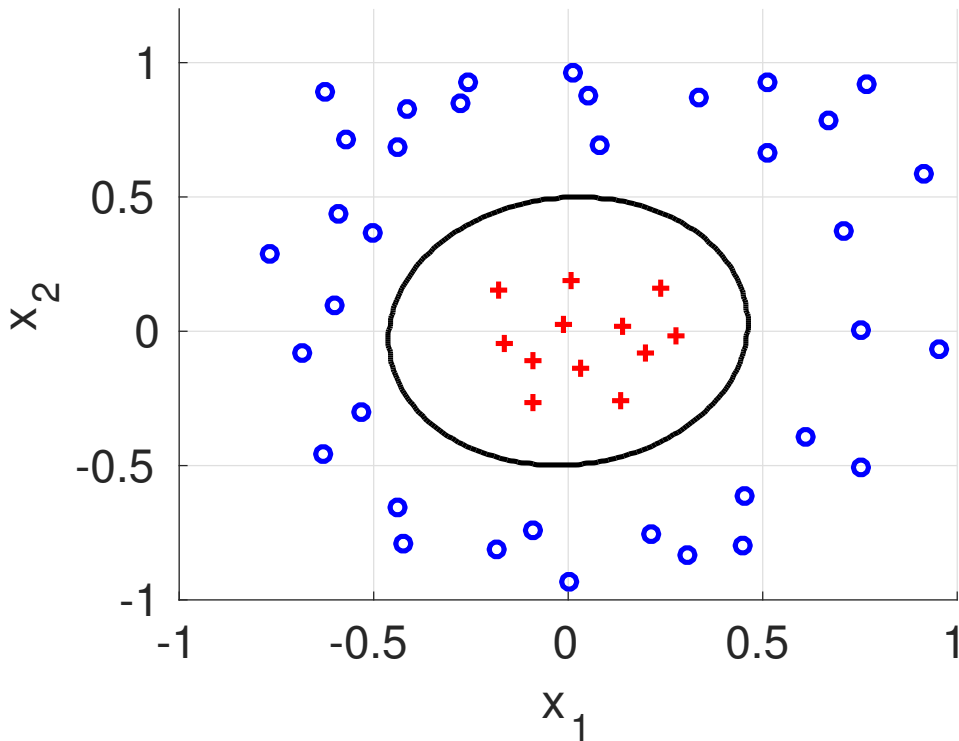


$$f(\mathbf{x}) = w_1\phi_1(\mathbf{x}) + w_2\phi_2(\mathbf{x}) + w_3\phi_3(\mathbf{x}) + b = w_1x_1^2 + w_2\sqrt{2}x_1x_2 + w_3x_2^2 + b$$

Example: Homogeneous polynomial kernel of degree 2

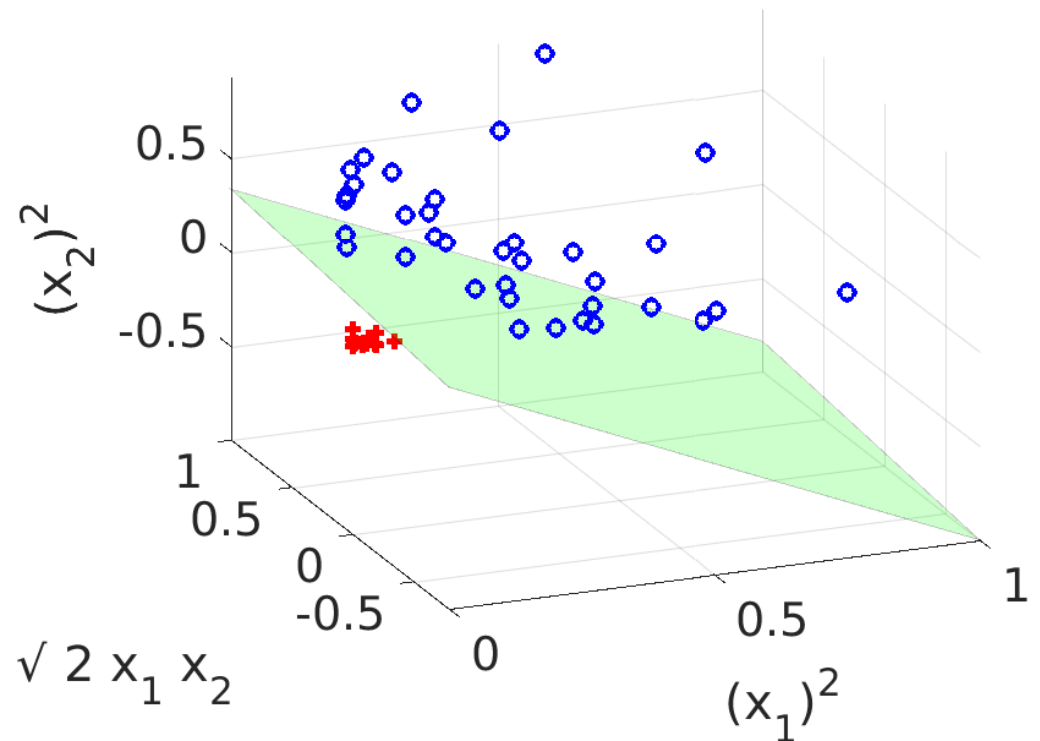
observations

$$\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$$



features

$$\phi(\mathbf{x}) = (x_1^2, \sqrt{2}x_1x_2, x_2^2) \in \mathbb{R}^3$$



$$f(\mathbf{x}) = w_1\phi_1(\mathbf{x}) + w_2\phi_2(\mathbf{x}) + w_3\phi_3(\mathbf{x}) + b = w_1x_1^2 + w_2\sqrt{2}x_1x_2 + w_3x_2^2 + b$$

Radial basis function kernel

- ◆ Assume the observations are real-valued features: $\mathcal{X} = \mathbb{R}^d$
- ◆ The RBF kernel

$$k(\mathbf{x}, \mathbf{x}') = \exp(-\gamma \|\mathbf{x} - \mathbf{x}'\|^2)$$

corresponds to dot product $\langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$ in infinite dimensional space.

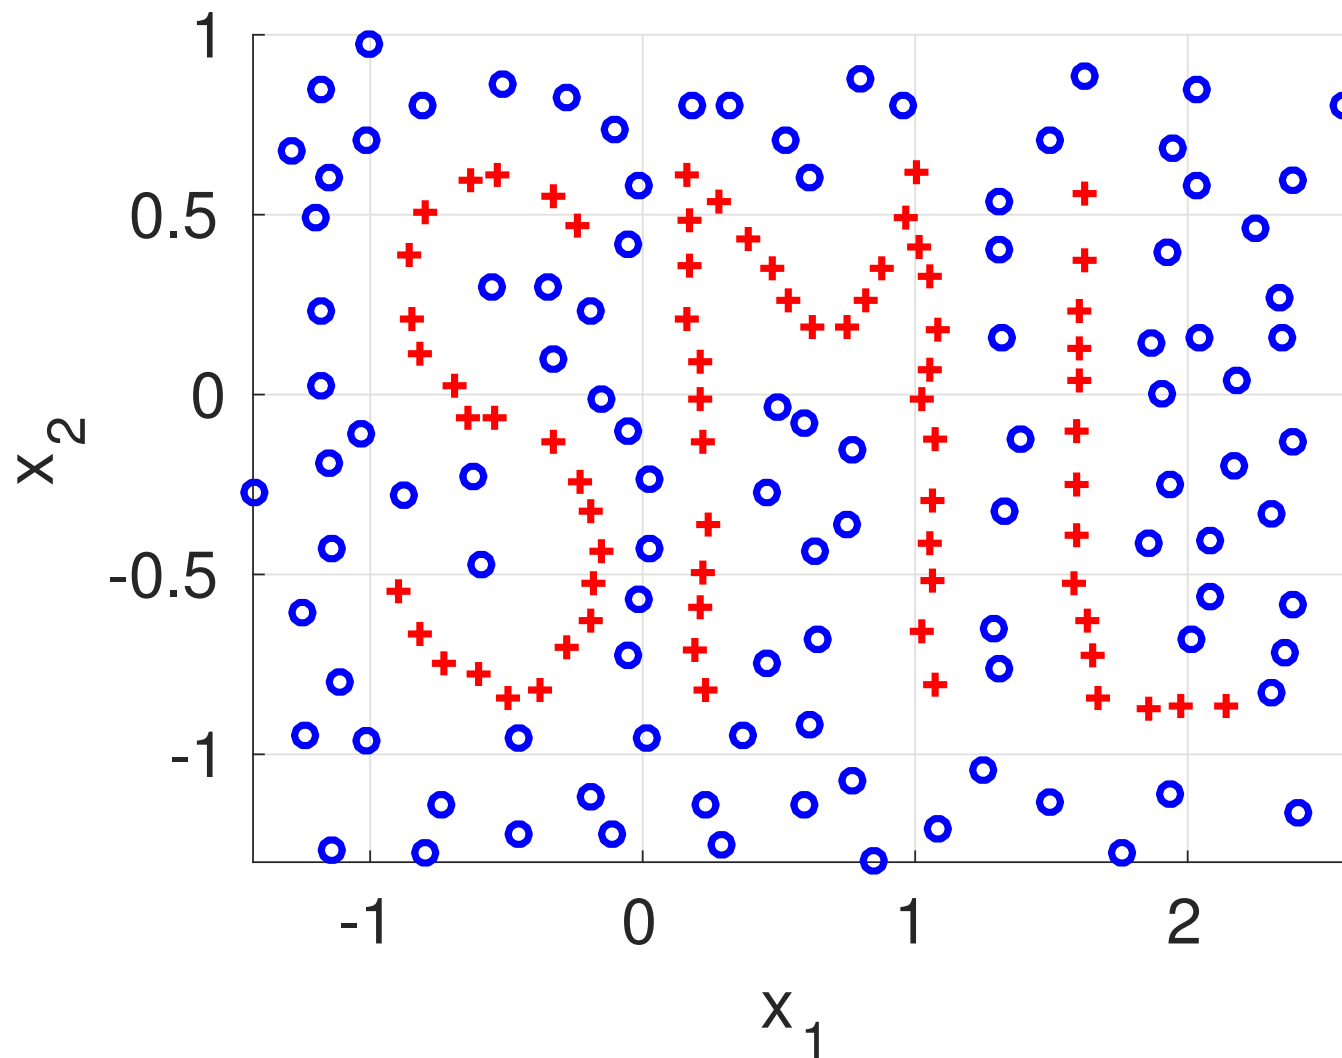
- ◆ In one dimensional case, $\mathcal{X} = \mathbb{R}$ the RBF kernel reads

$$\begin{aligned} k(x, x') &= \exp(-\gamma(x - x')^2) \\ &= \exp(-\gamma x^2) \exp(-\gamma x'^2) \exp(2\gamma x x') \\ &= \exp(-\gamma x^2) \exp(-\gamma x'^2) \sum_{n=0}^{\infty} \frac{(2 x x' \gamma)^n}{n!} \\ &= \exp(-\gamma x^2) \exp(-\gamma x'^2) \sum_{n=0}^{\infty} \frac{(\sqrt{2\gamma} x)^n}{\sqrt{n!}} \frac{(\sqrt{2\gamma} x')^n}{\sqrt{n!}} \end{aligned}$$

where we used the Taylor expansion $\exp(a) = \sum_{n=0}^{\infty} \frac{a^n}{n!}$

Example: SVM with RBF kernel

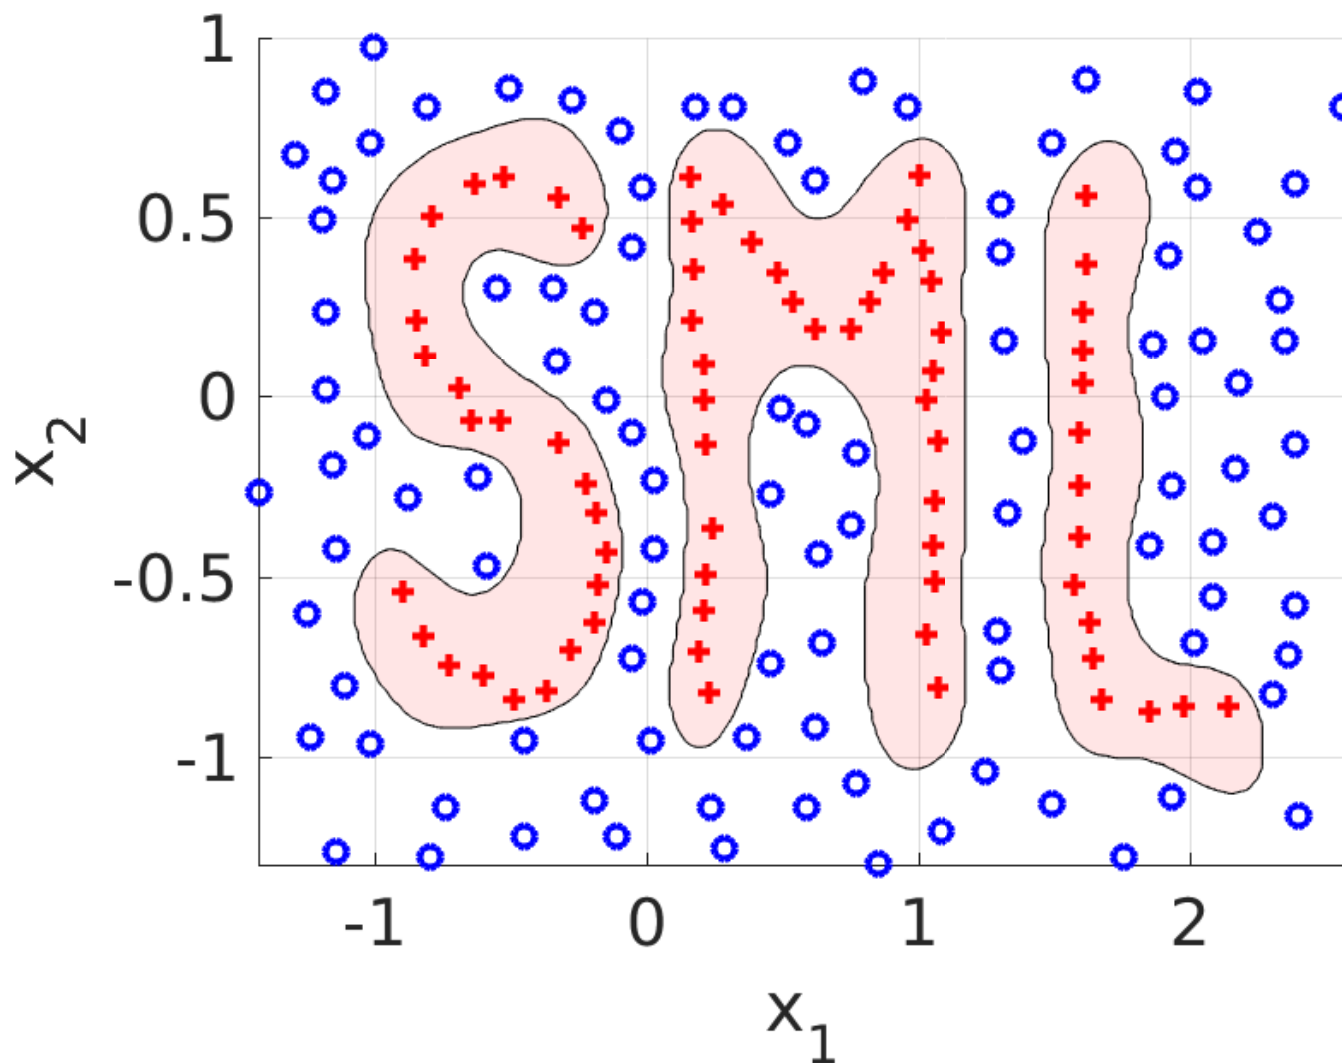
$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$



Example: SVM with RBF kernel

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

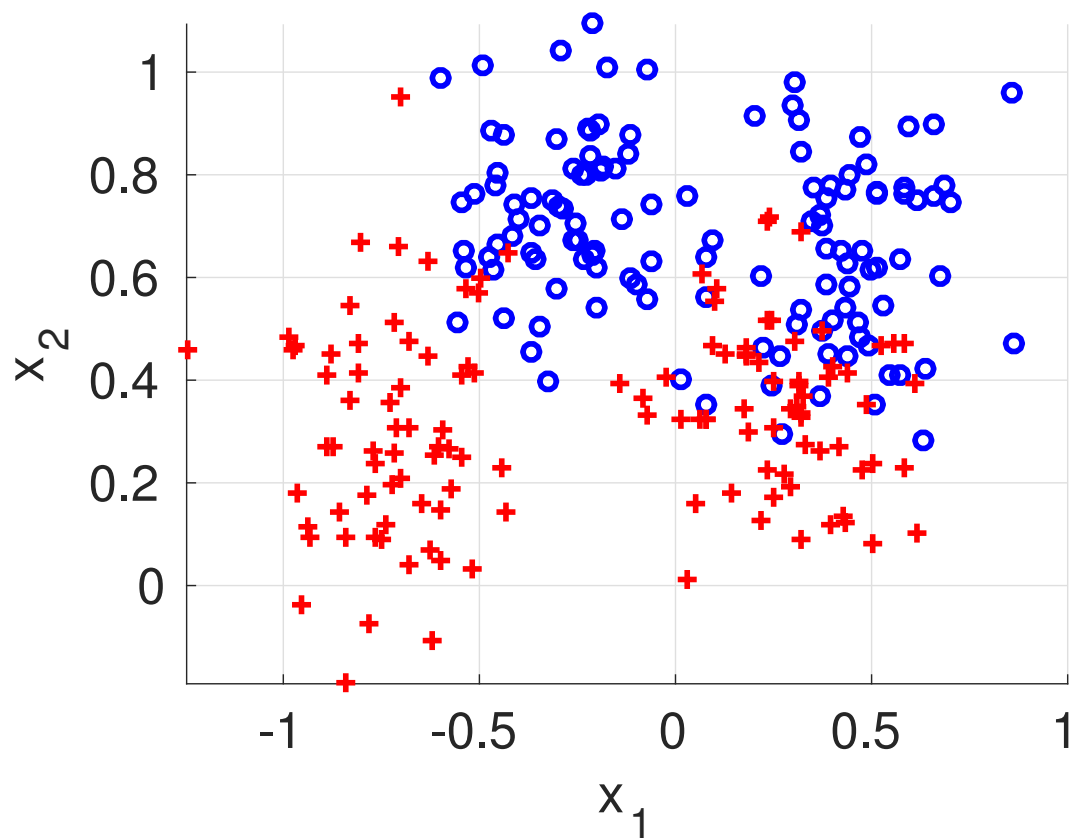
$$\gamma = 10, C = 1000$$



Example: SVM with RBF kernel

training error vs. γ

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

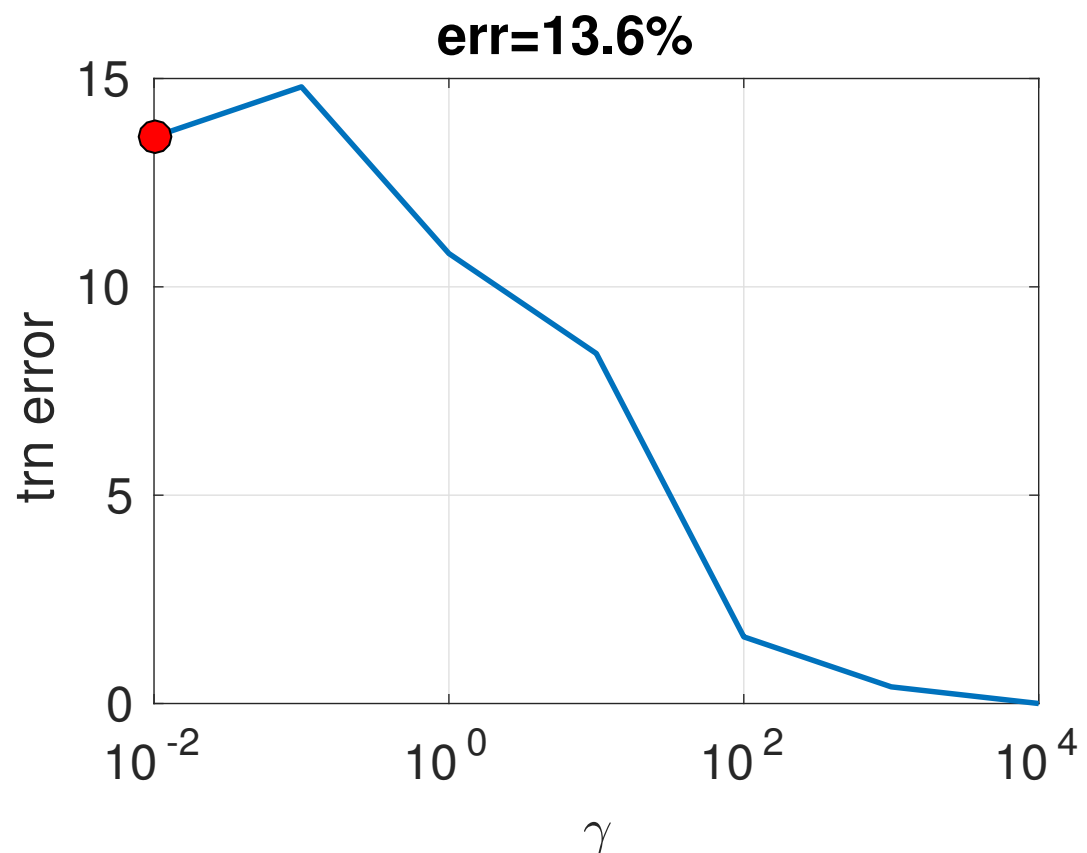
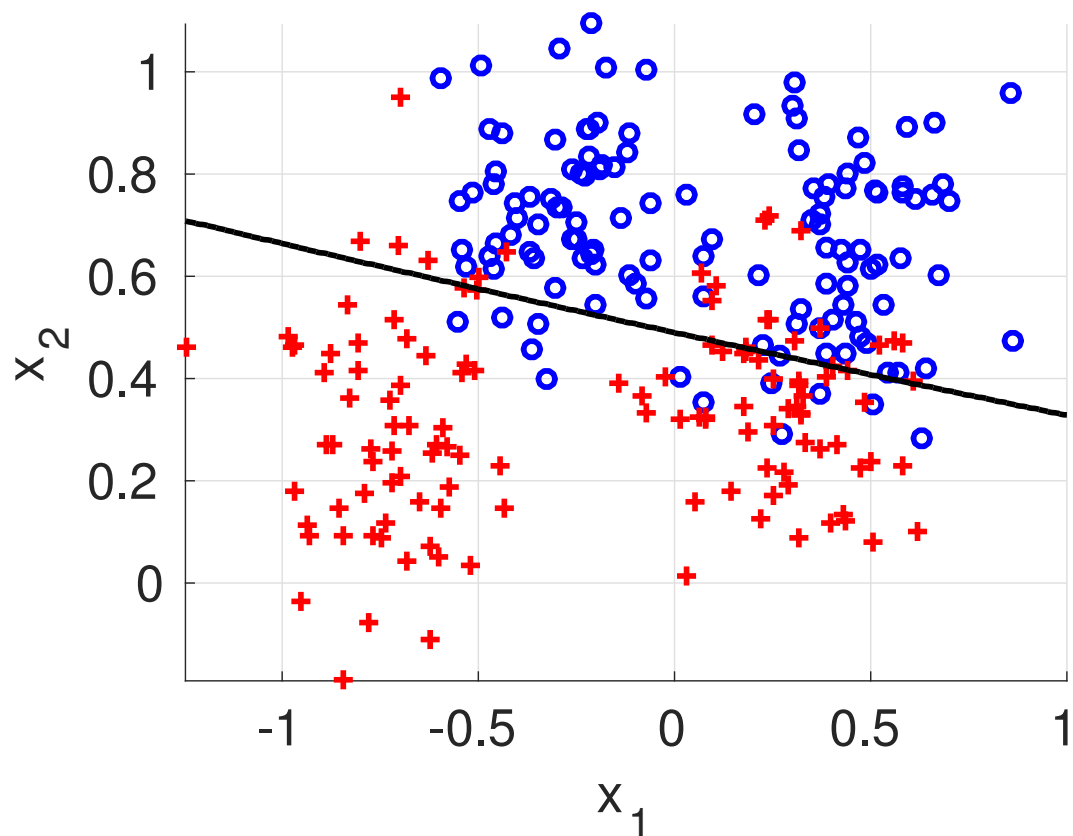


Example: SVM with RBF kernel

training error vs. γ

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 0.01, C = 100$$

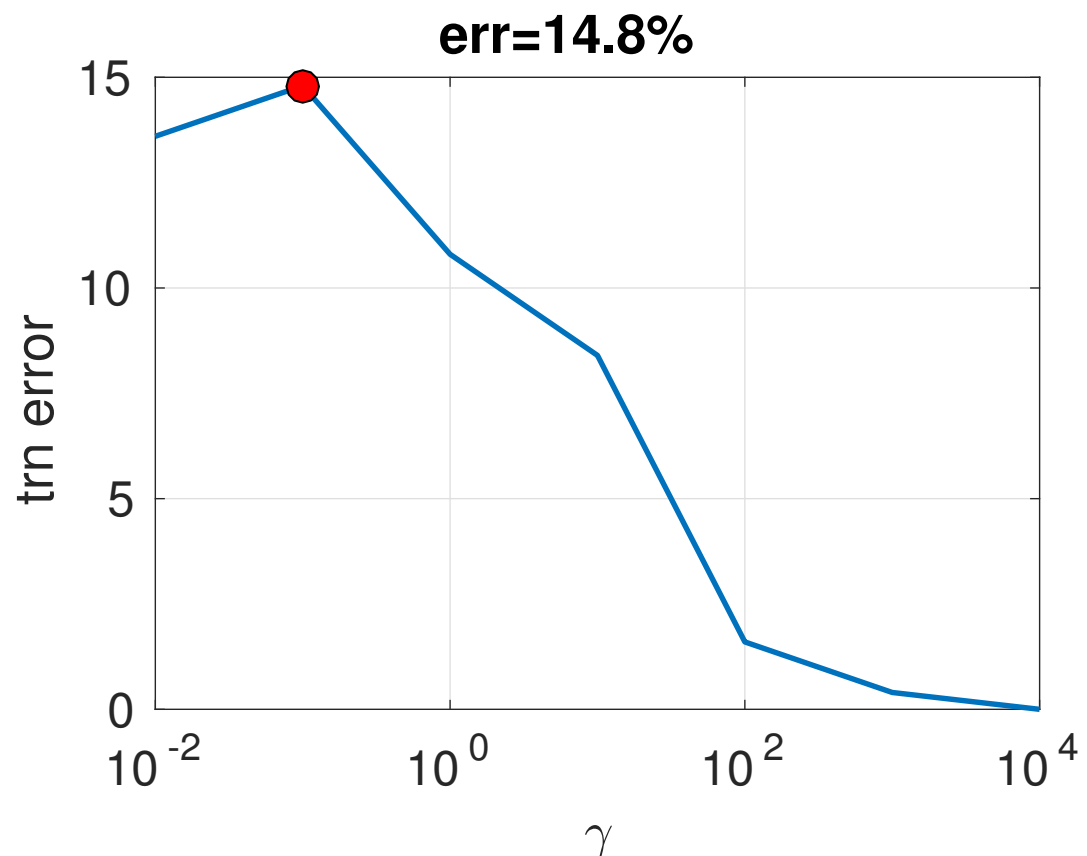
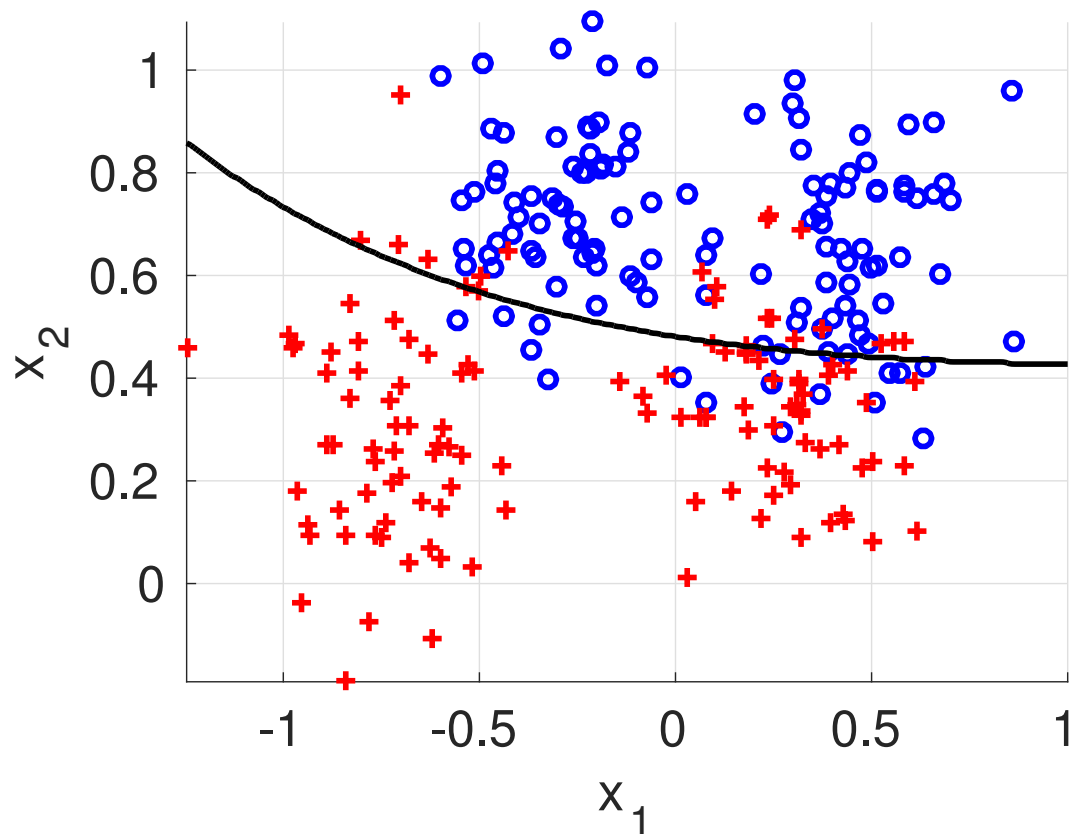


Example: SVM with RBF kernel

training error vs. γ

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 0.1, C = 100$$

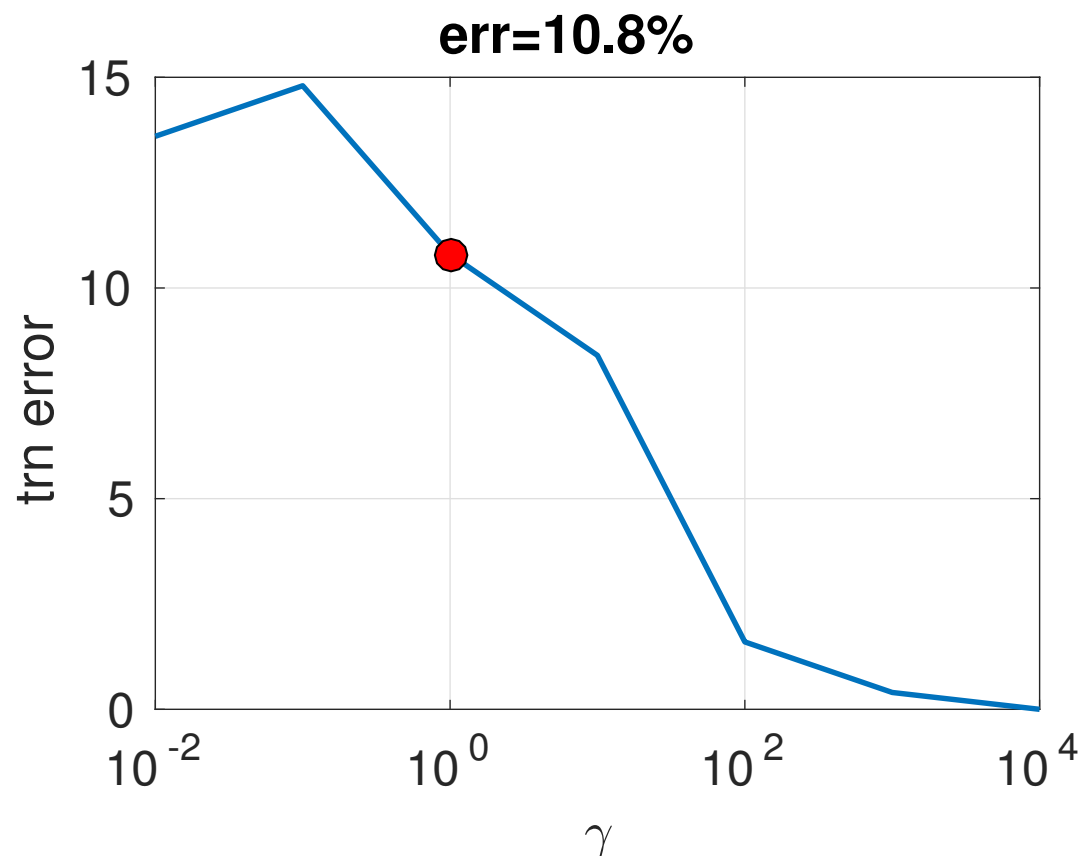
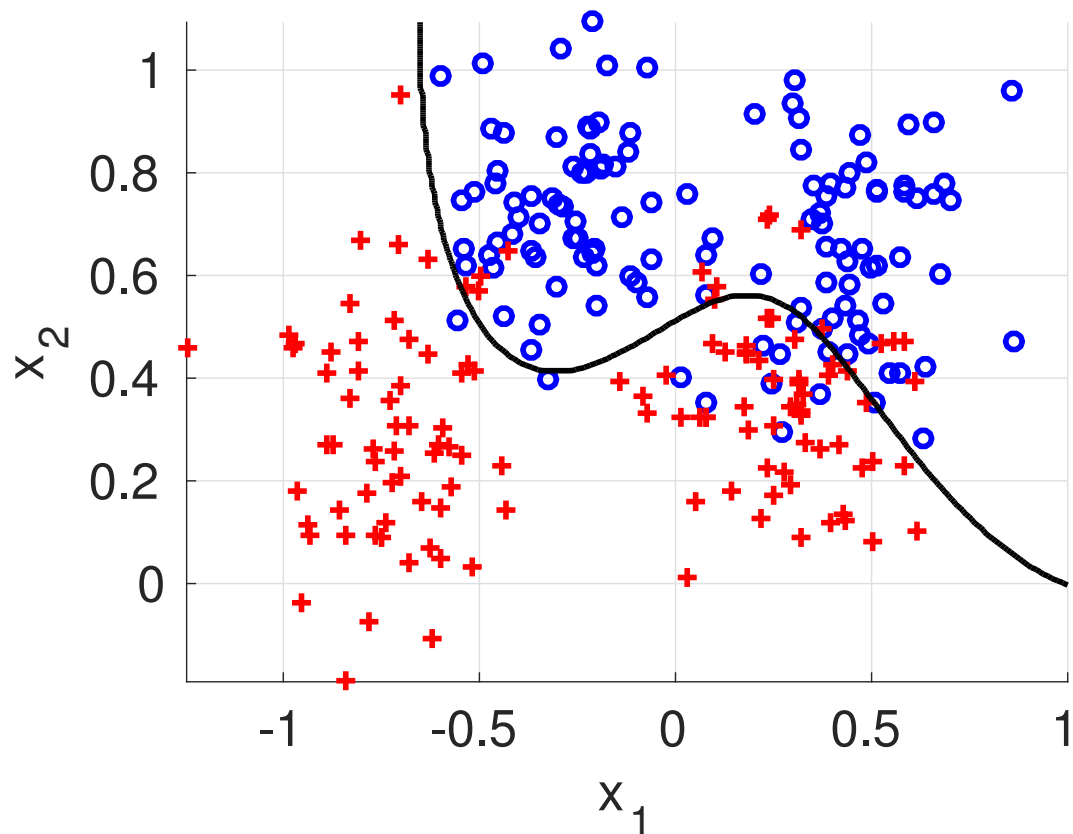


Example: SVM with RBF kernel

training error vs. γ

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 1, C = 100$$

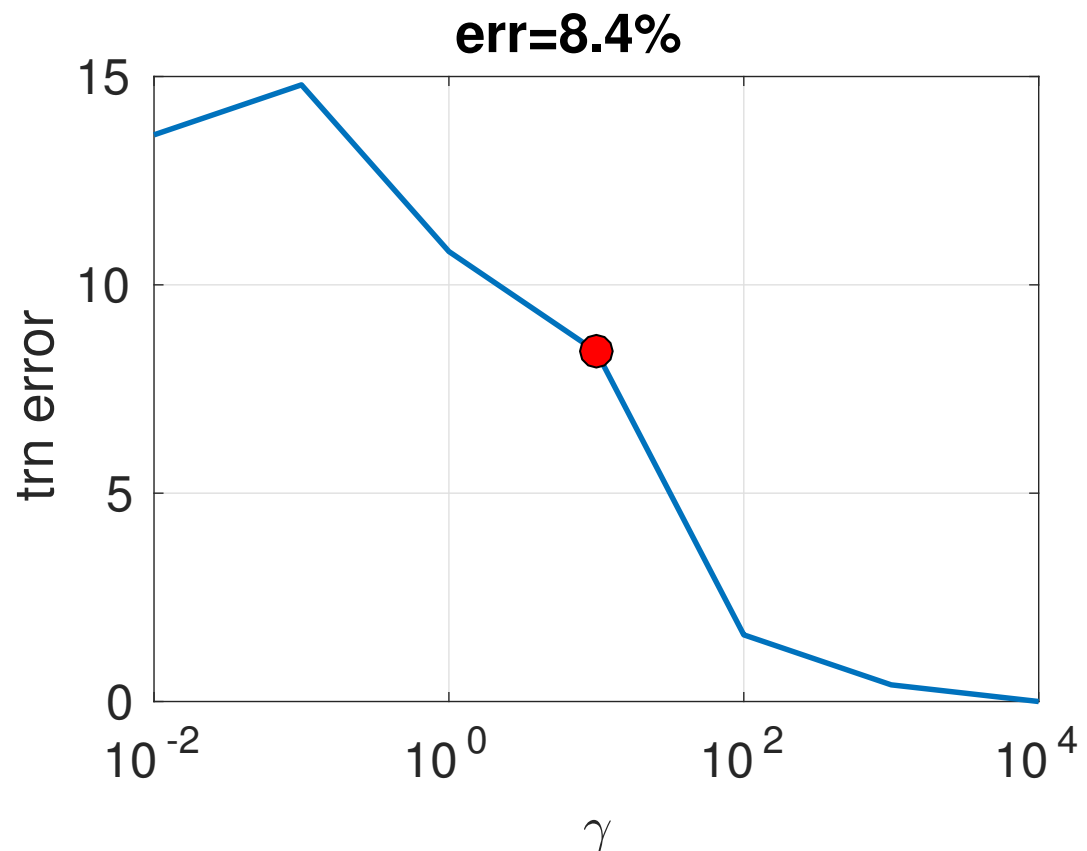
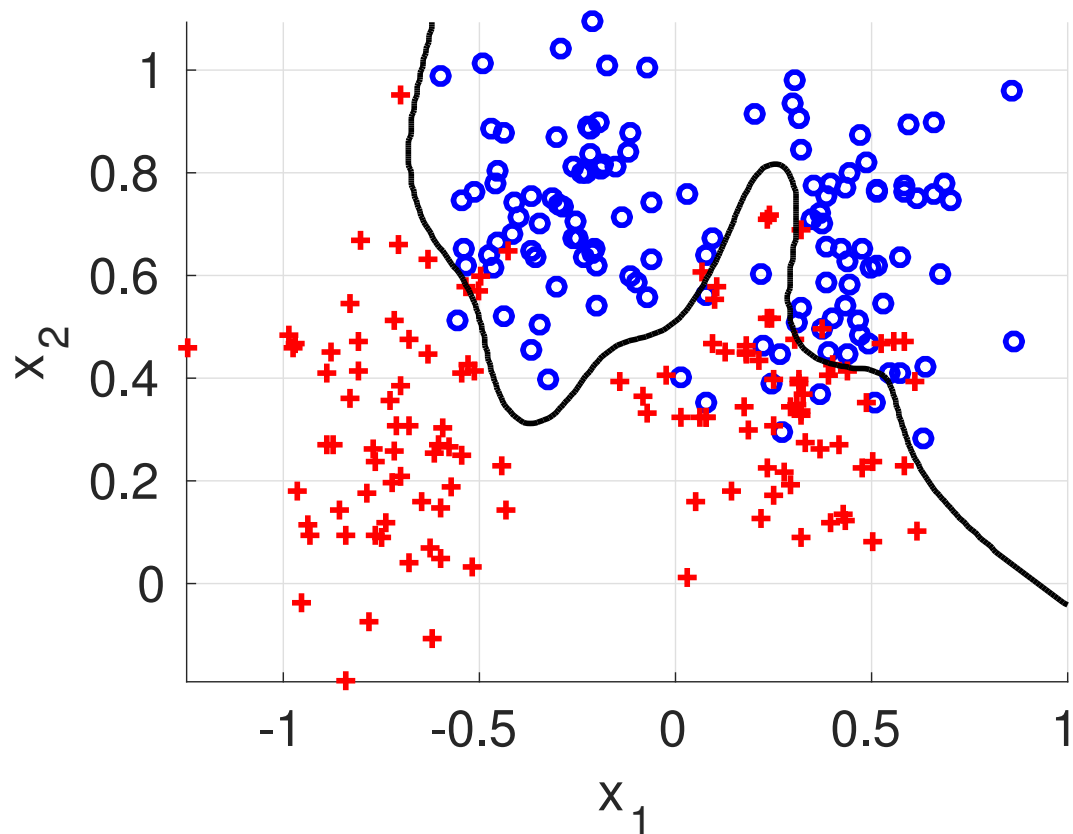


Example: SVM with RBF kernel

training error vs. γ

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 10, C = 100$$

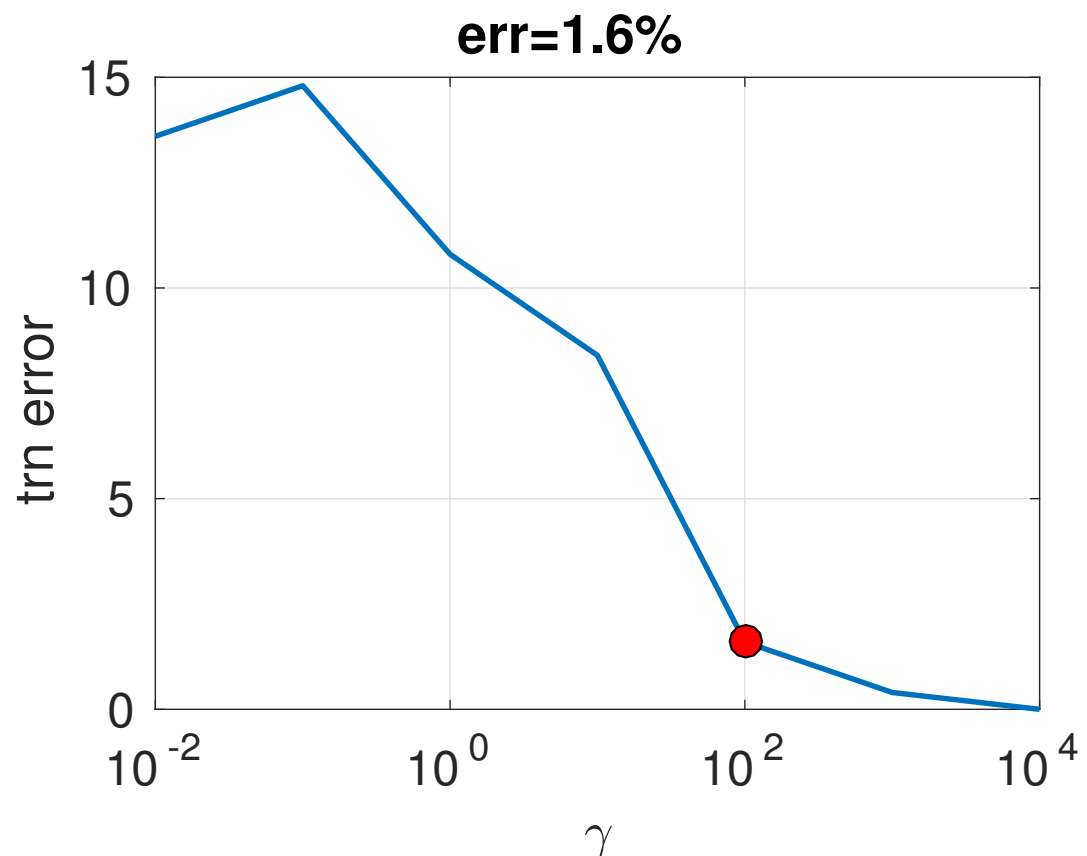
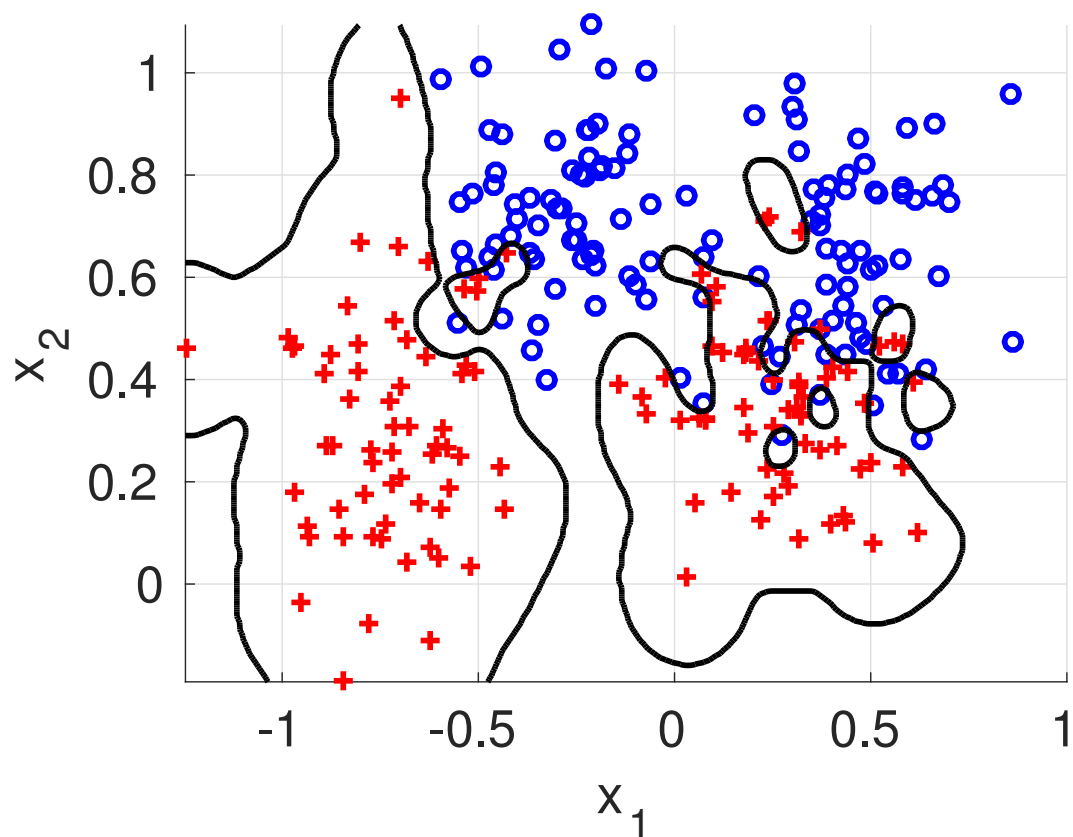


Example: SVM with RBF kernel

training error vs. γ

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 100, C = 100$$

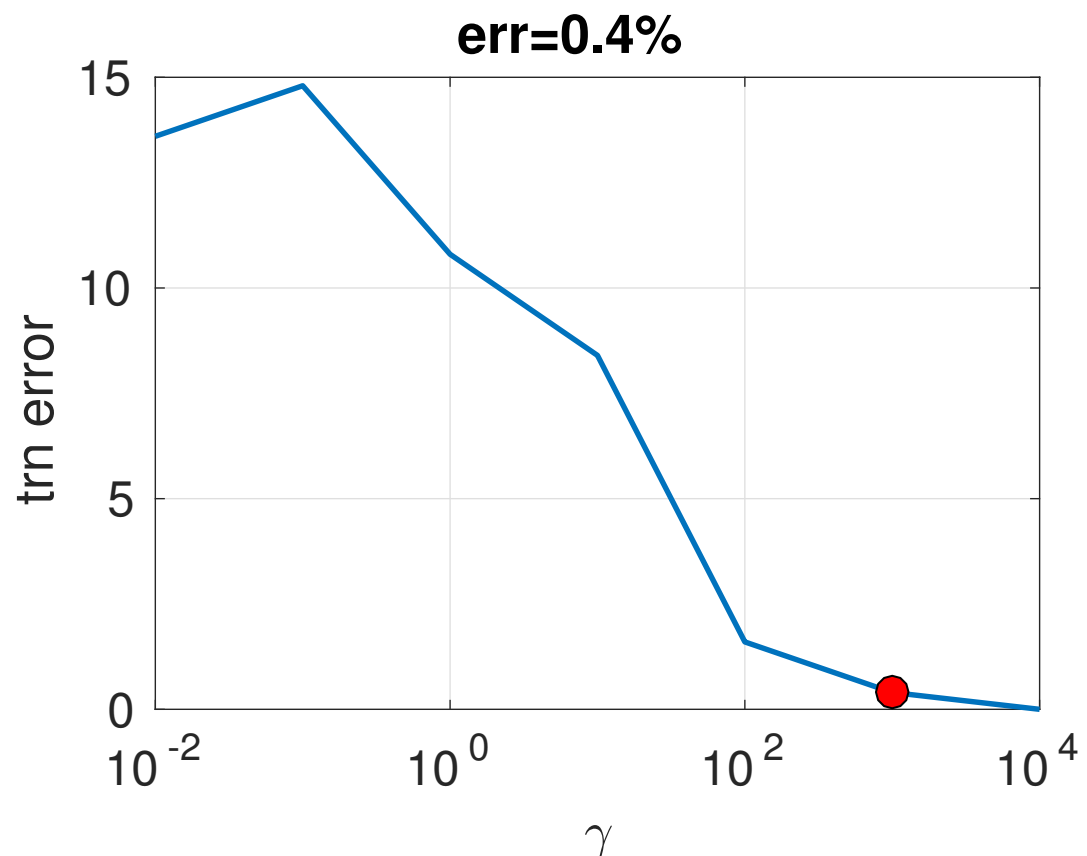
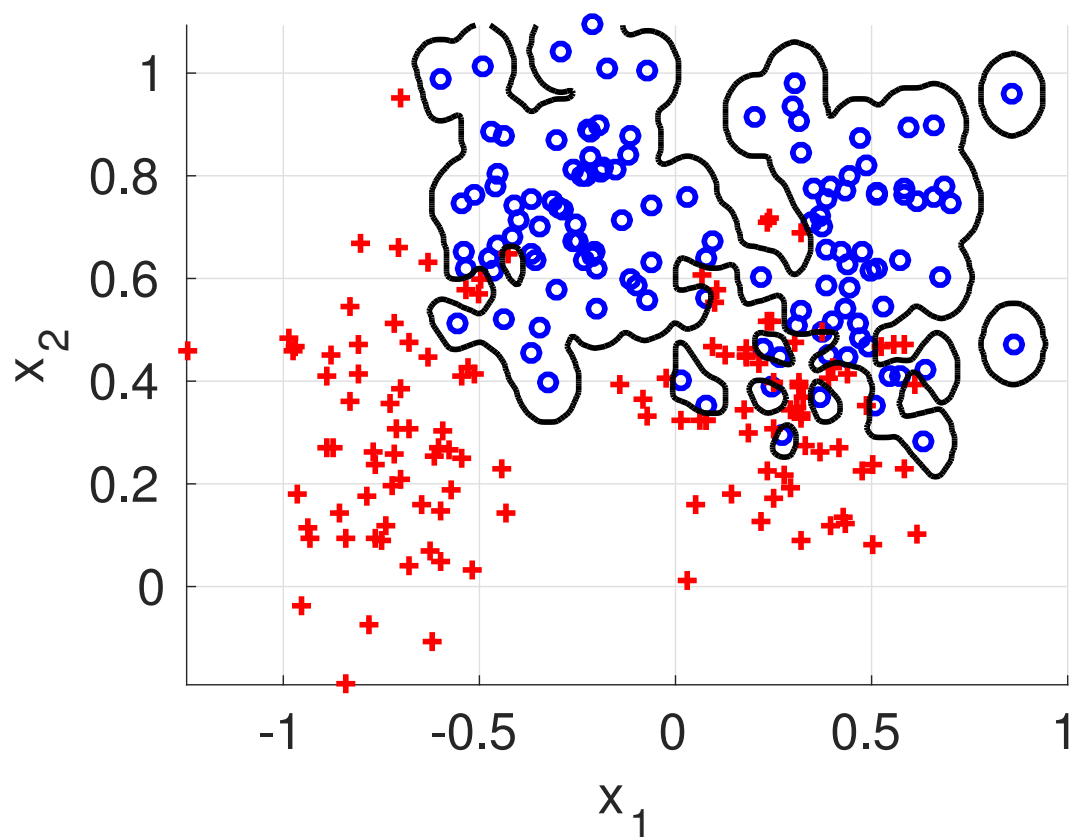


Example: SVM with RBF kernel

training error vs. γ

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 1000, C = 100$$

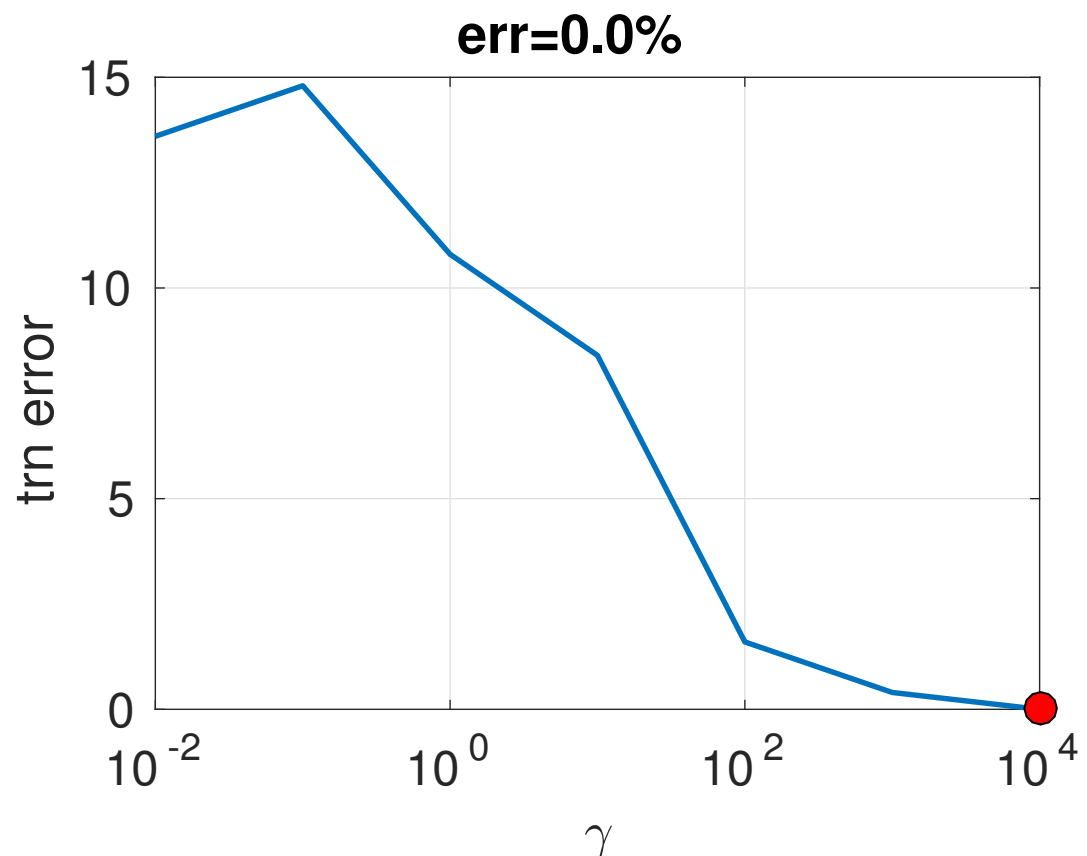
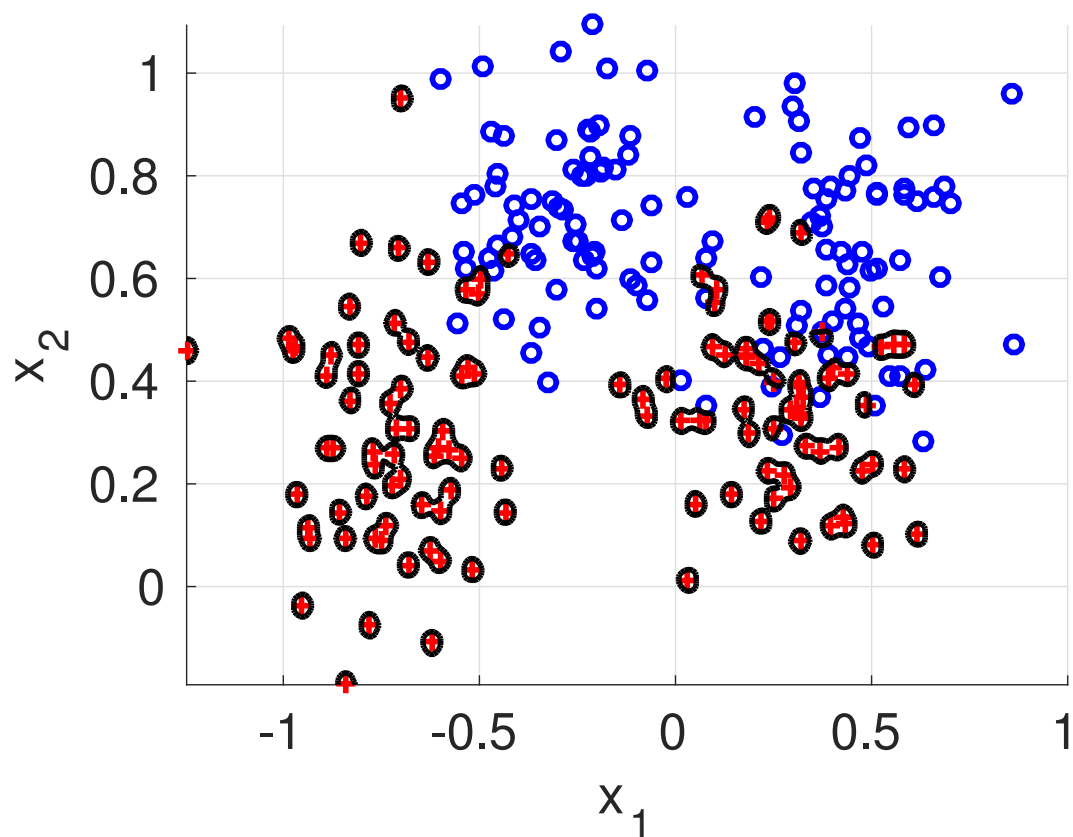


Example: SVM with RBF kernel

training error vs. γ

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

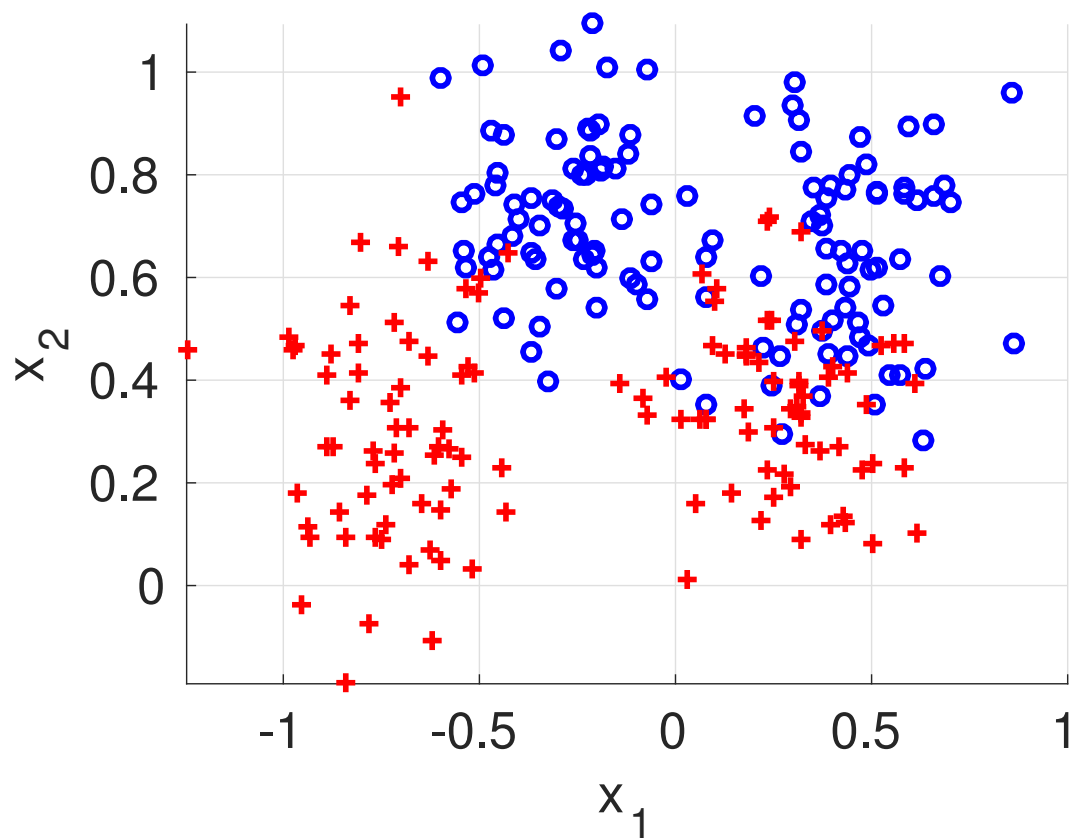
$$\gamma = 10000, C = 100$$



Example: SVM with RBF kernel

training error vs. C

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

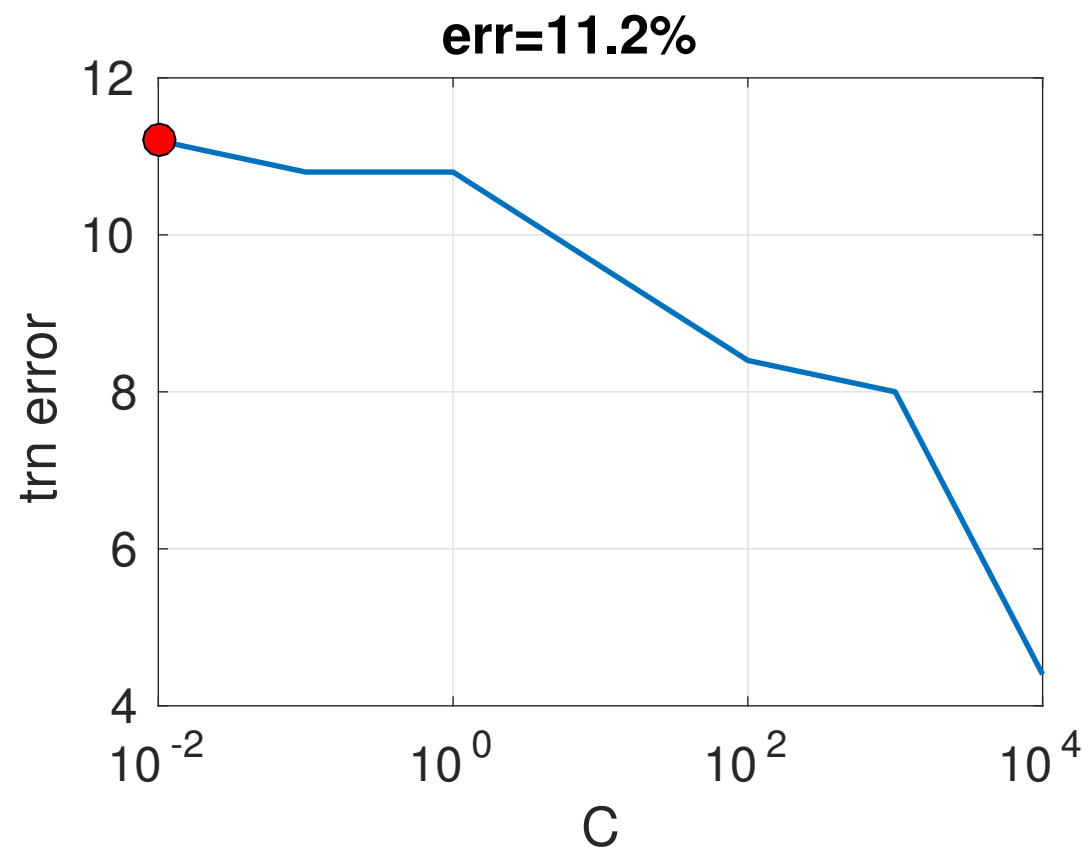
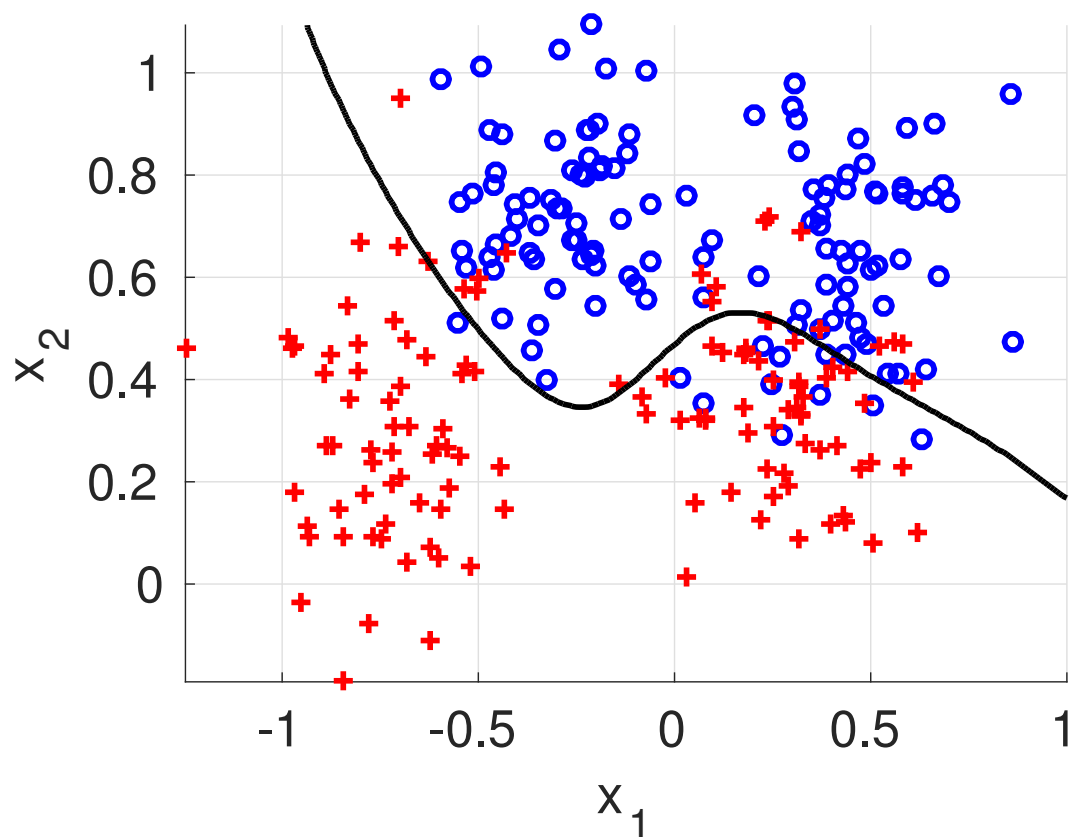


Example: SVM with RBF kernel

training error vs. C

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 10, C = 0.01$$

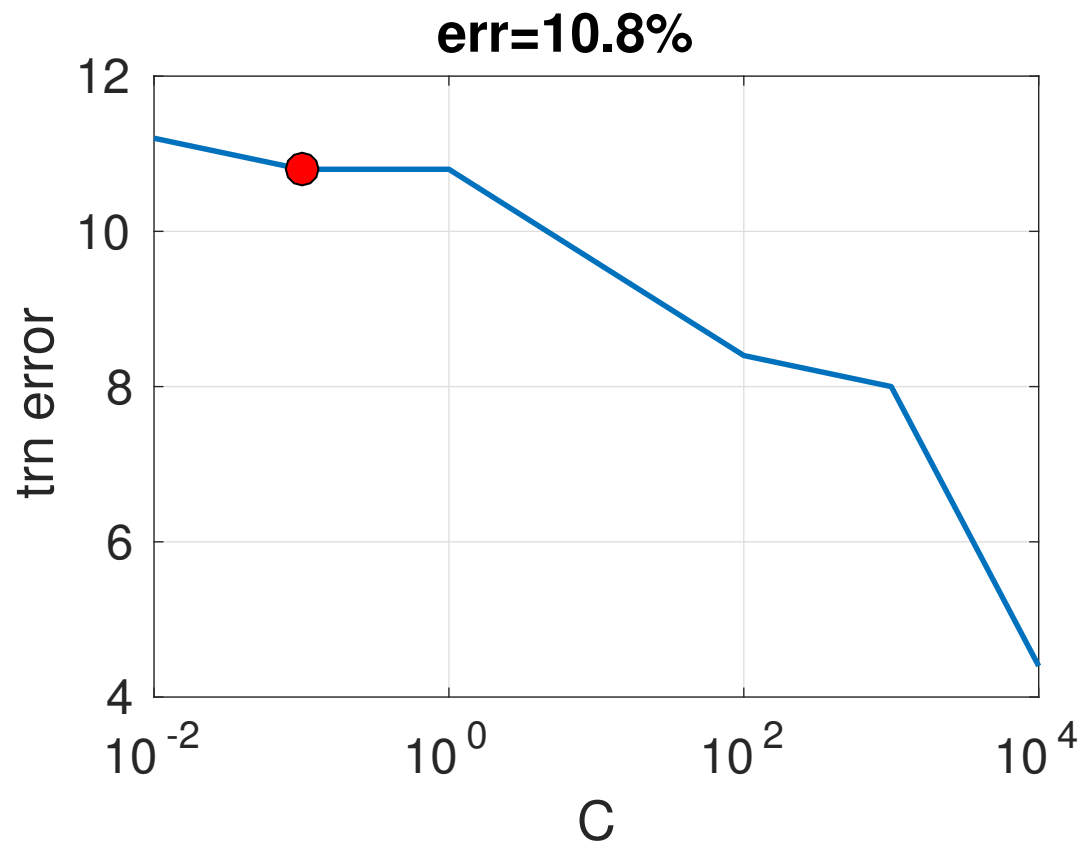
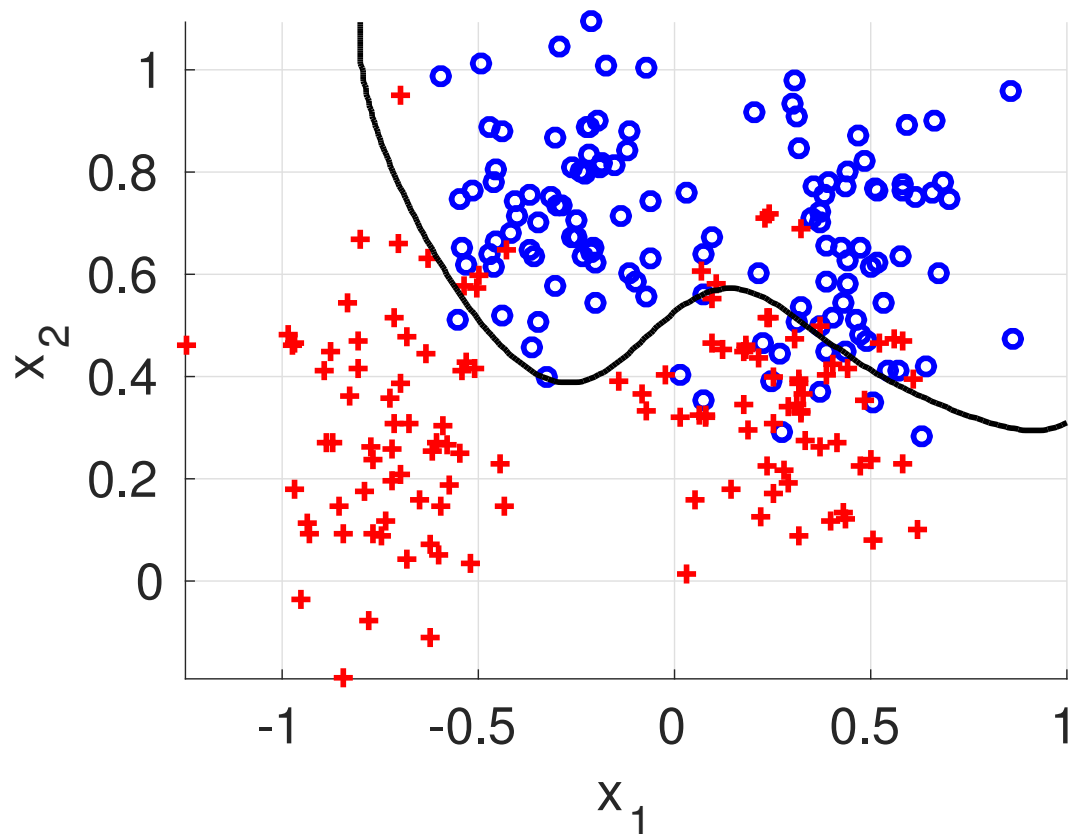


Example: SVM with RBF kernel

training error vs. C

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 10, C = 0.1$$

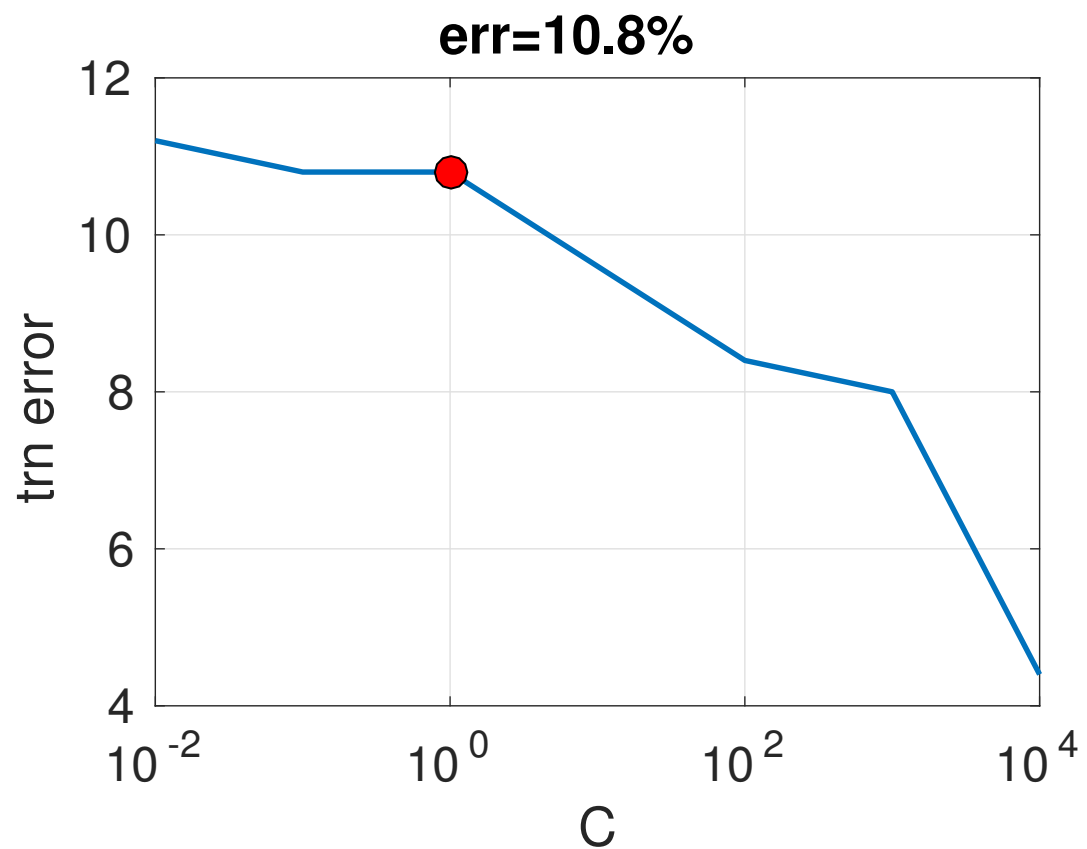
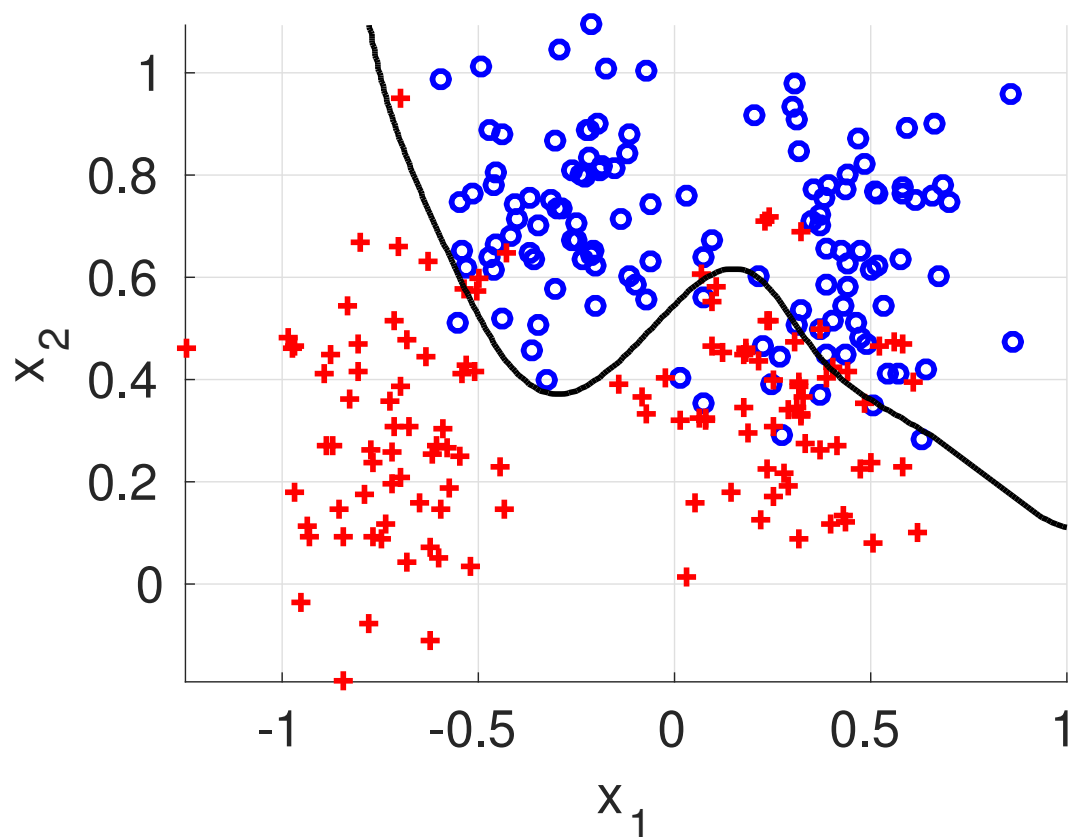


Example: SVM with RBF kernel

training error vs. C

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 10, C = 1$$

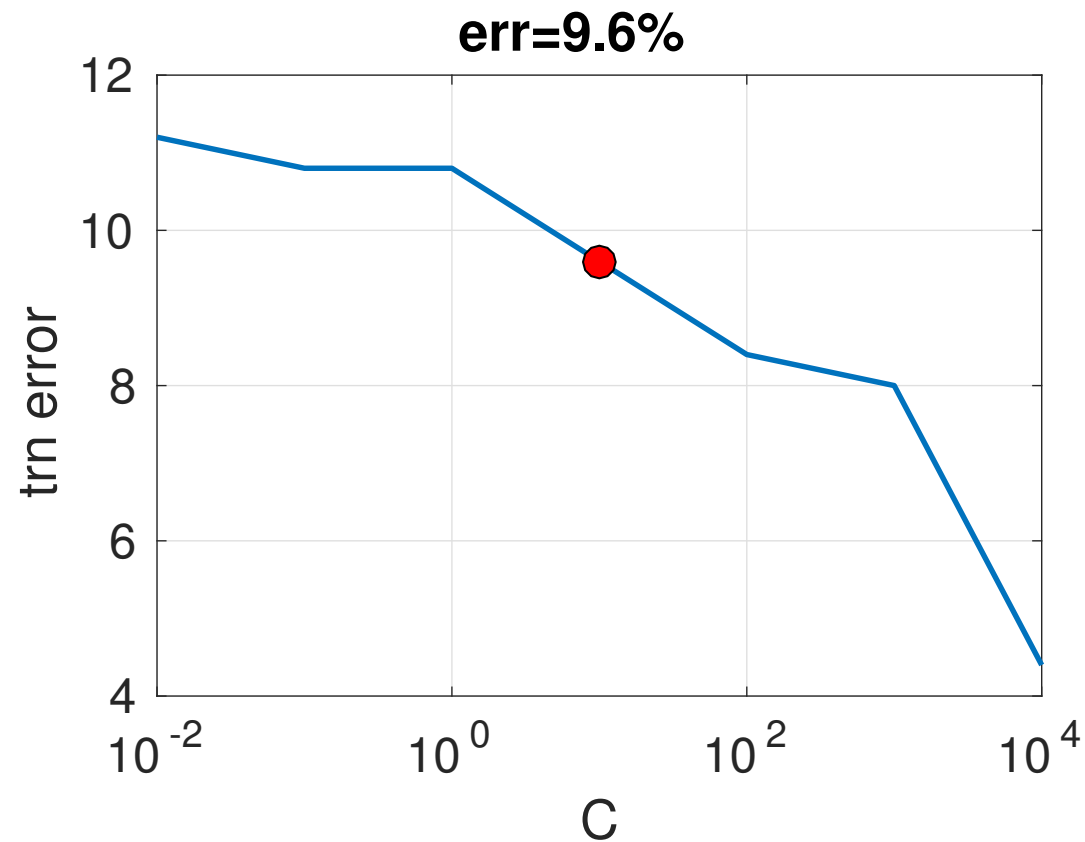
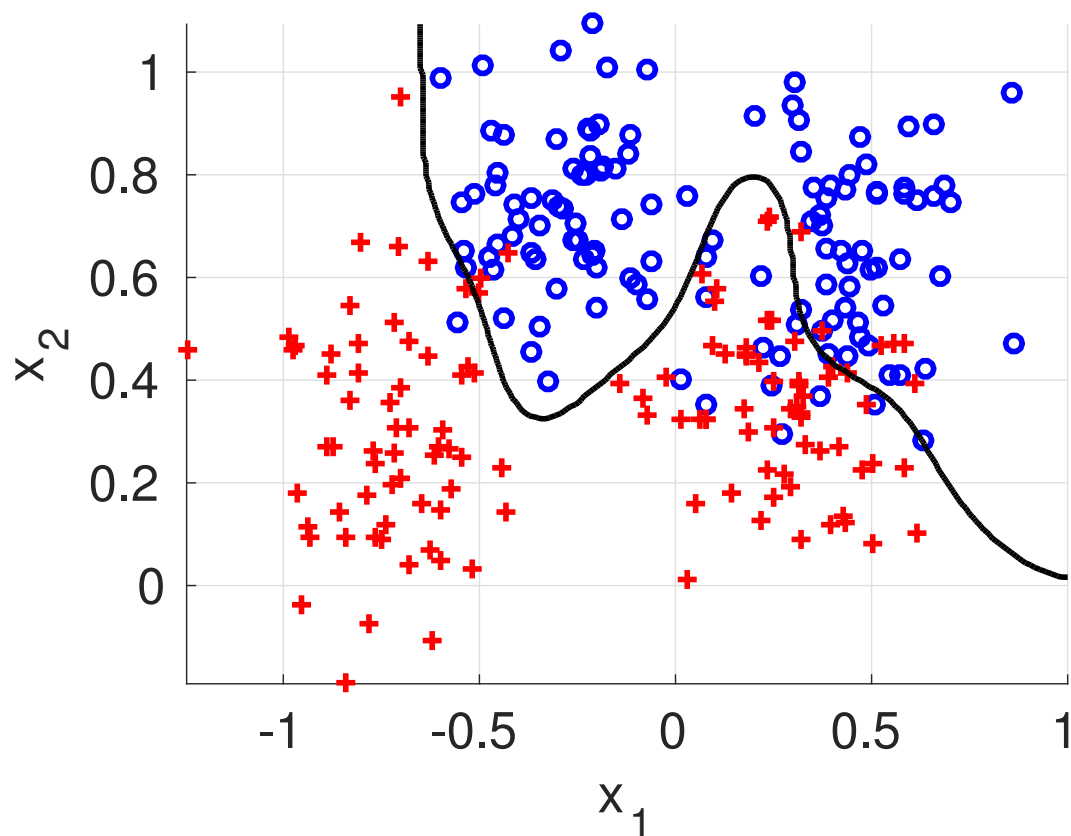


Example: SVM with RBF kernel

training error vs. C

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 10, C = 10$$

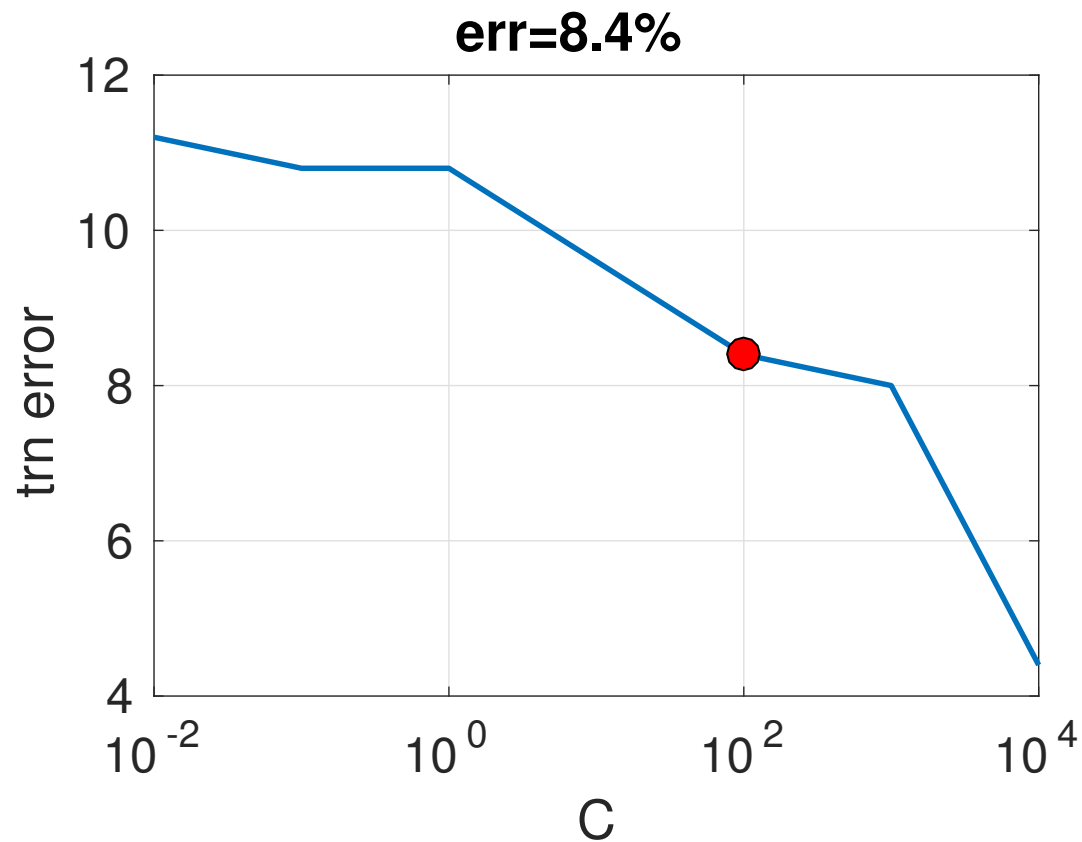
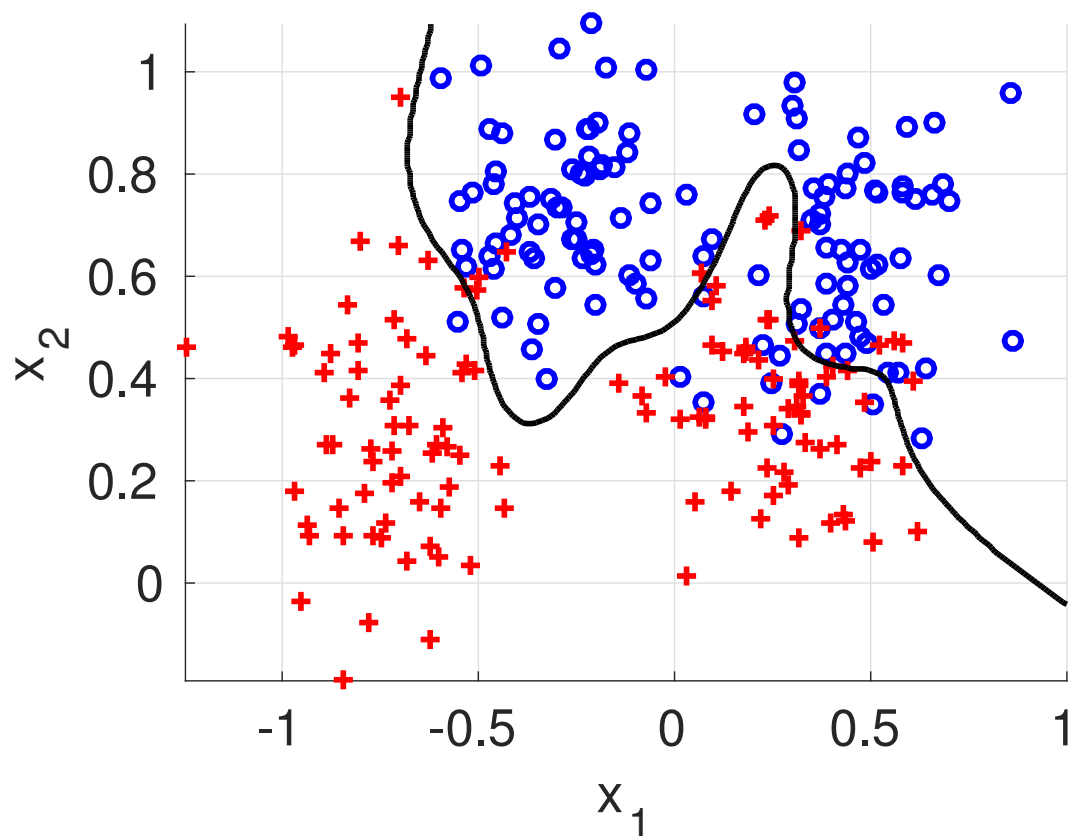


Example: SVM with RBF kernel

training error vs. C

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 10, C = 100$$

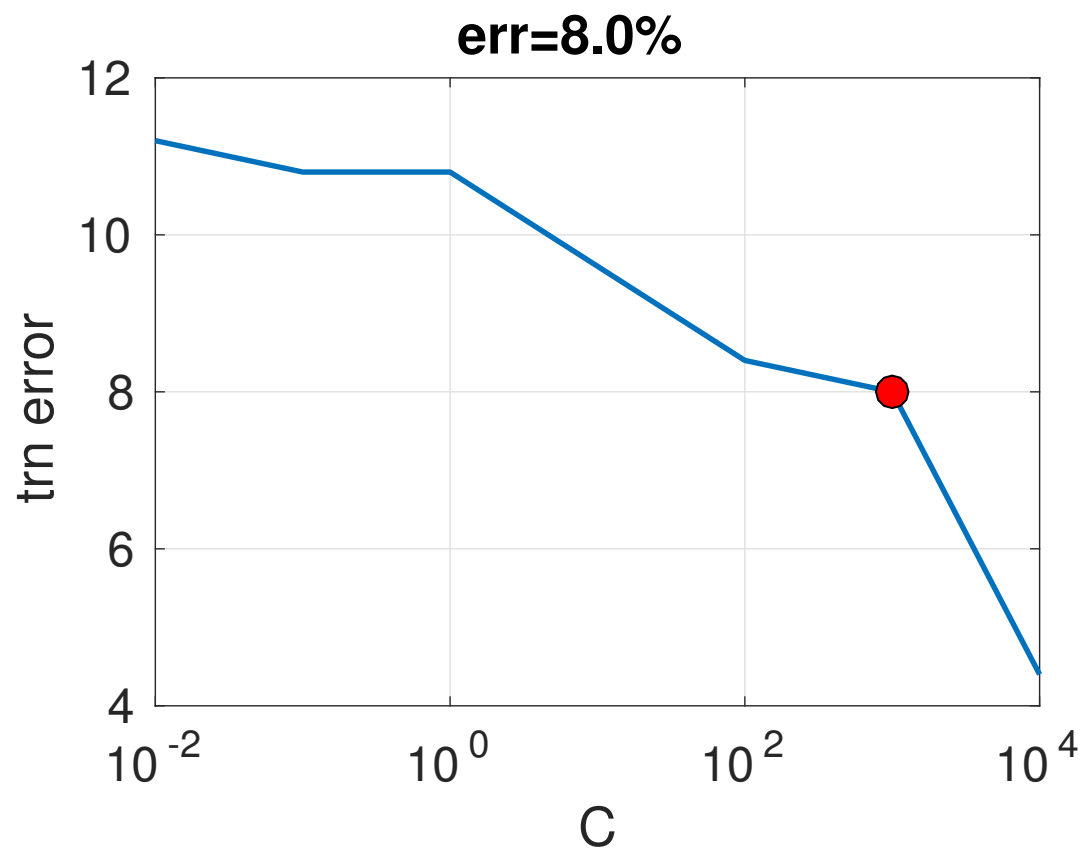
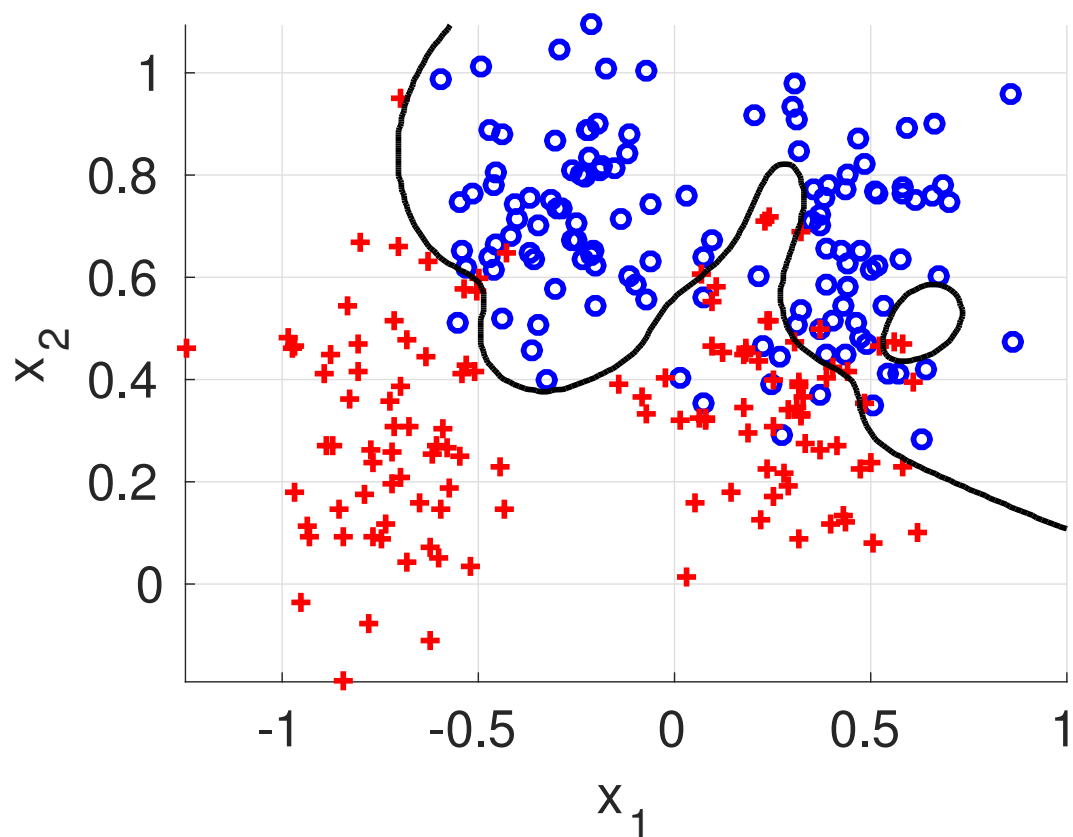


Example: SVM with RBF kernel

training error vs. C

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 10, C = 1000$$

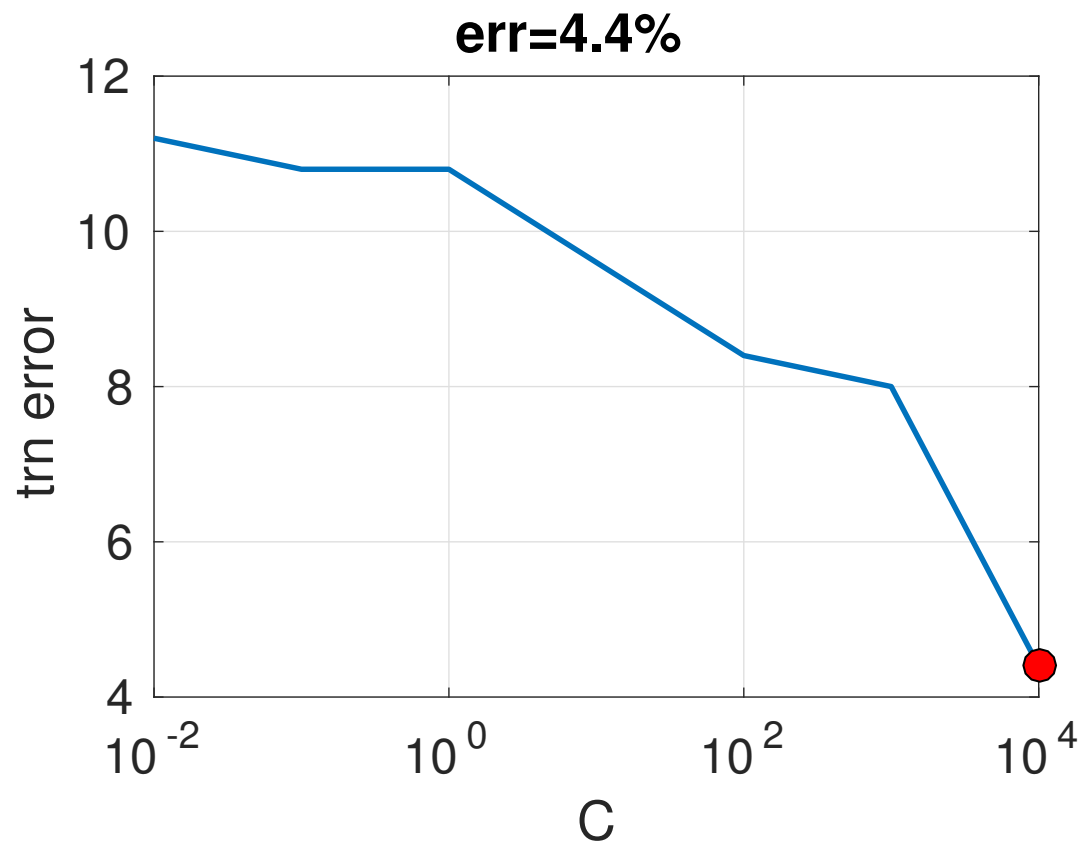
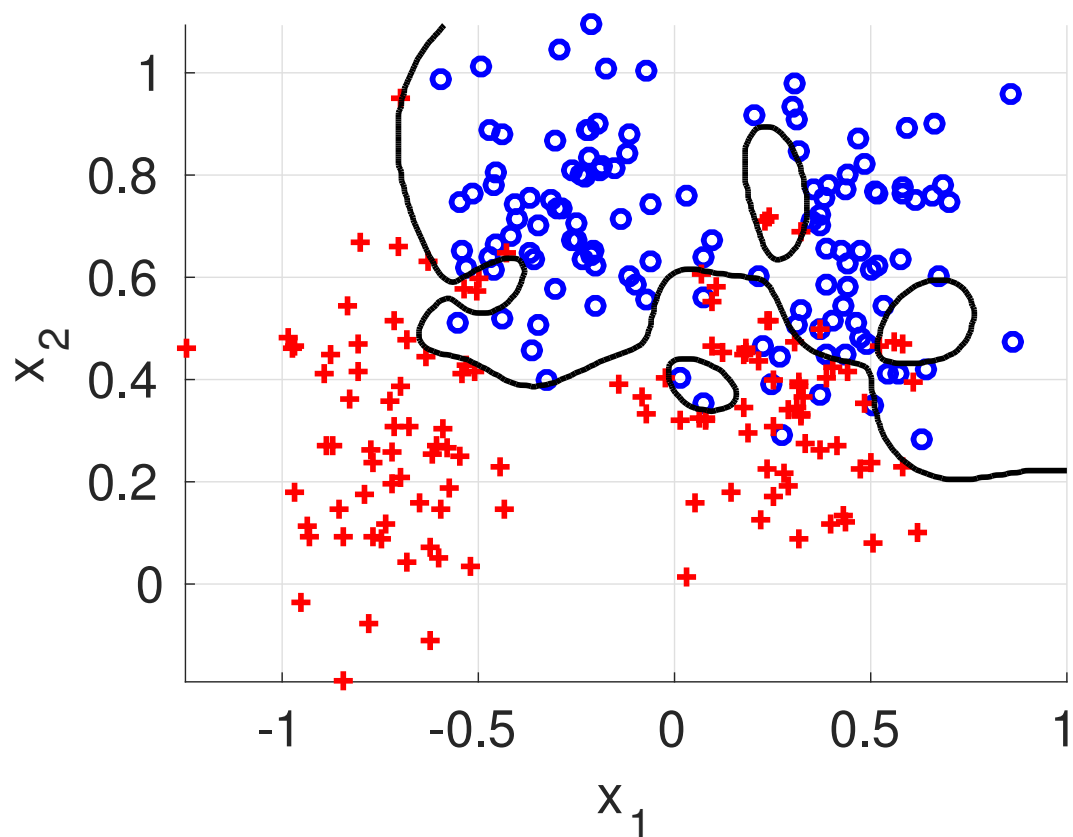


Example: SVM with RBF kernel

training error vs. C

$$f(\mathbf{x}) = \sum_{i \in \mathcal{I}_{SV}} y^i \alpha_i \exp(-\gamma \|\mathbf{x} - \mathbf{x}^i\|^2) + b$$

$$\gamma = 10, C = 10000$$



String Subsequence kernel

- ◆ Input space $\mathcal{X} = \cup_{d=1}^{\infty} \Sigma^d$ contains all strings on a finite alphabet Σ
- ◆ The features: occurrence+continuity of sub-sequences of length q :

$$\phi: \mathcal{X} \rightarrow \mathbb{R}^{|\Sigma|^q} \quad \text{and} \quad \phi_u(s) = \sum_{i: u=s[i]} \lambda^{l(i)}, \quad \forall u \in \Sigma^q$$

where $s[i]$ denotes a sub-sequence of s and $\lambda \in (0, 1]$ is a decay factor.

- ◆ Example for strings "cat", "car", "bat" and "bar" and $q = 2$:

	ca	ct	at	ba	bt	cr	ar	br
$\phi(\text{"cat"})$	λ^2	λ^3	λ^2	0	0	0	0	0
$\phi(\text{"car"})$	λ^2	0	0	0	0	λ^3	λ^2	0
$\phi(\text{"bat"})$	0	0	λ^2	λ^2	λ^3	0	0	0
$\phi(\text{"bar"})$	0	0	0	λ^2	0	0	λ^2	λ^3

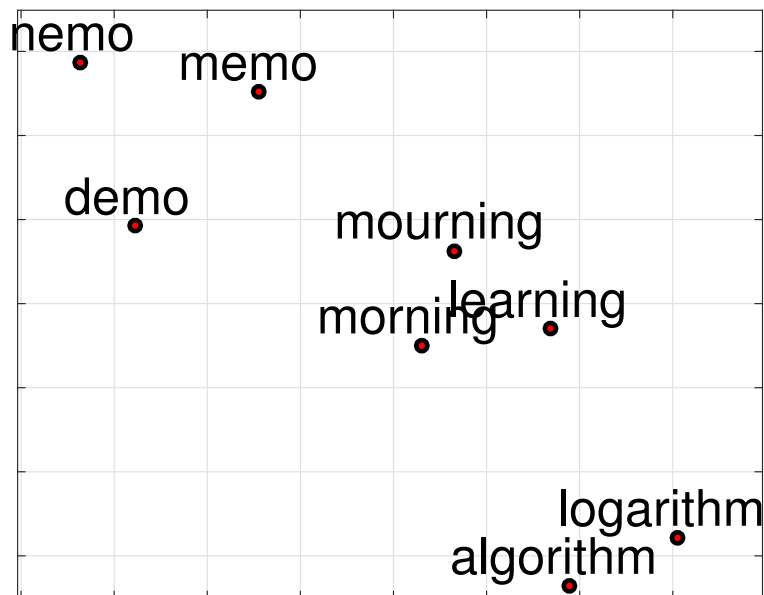
$k(\text{"cat"}, \text{"car"}) = \lambda^4$, $k(\text{"cat"}, \text{"bar"}) = \lambda^4$, $k(\text{"cat"}, \text{"bar"}) = 0$, ...

- ◆ The kernel value $k(s, t)$ can be computed by dynamic programming in time $\mathcal{O}(q |s| |t|)$.

Example: String Subsequence kernel

kernel parameters: $q = 2, \lambda = 0.4$

Feature space
(2D embedding via tSNE)



Kernel matrix

algorithm	0.24	0.15	0.01	0.04	0.02	0.00	0.00	0.00
logarithm	0.15	0.24	0.04	0.02	0.01	0.00	0.00	0.00
learning	0.01	0.04	0.23	0.13	0.13	0.00	0.00	0.00
morning	0.04	0.02	0.13	0.19	0.17	0.03	0.03	0.03
mourning	0.02	0.01	0.13	0.17	0.23	0.03	0.03	0.03
demo	0.00	0.00	0.00	0.03	0.03	0.09	0.06	0.06
memo	0.00	0.00	0.00	0.03	0.03	0.06	0.09	0.06
nemo	0.00	0.00	0.00	0.03	0.03	0.06	0.06	0.09
	algorithm	logarithm	learning	morning	mourning	demo	memo	nemo

Primal vs. kernel SVM in terms of memory

$\phi: \mathcal{X} \rightarrow \mathbb{R}^n$.. feature map; m .. number of training examples; $\mathcal{O}(d)$.. size of the observation $x \in \mathcal{X}$ in bytes

Memory requirements: (assuming $x \in \mathcal{X}$ requires $\mathcal{O}(d)$ bytes)

	training set \mathcal{T}^m	prediction rule
primal formulation	$\mathcal{O}(m \cdot n)$	$\mathcal{O}(n + 1)$
kernel formulation	$\mathcal{O}(m^2)$	$\mathcal{O}((d + 1) \cdot \mathcal{I}_{\text{SV}} + 1)$

Examples of kernel functions:

	Input space \mathcal{X}	kernel $k(\mathbf{x}, \mathbf{x}')$	dim of $ \Phi(\mathbf{x}) $
2nd polynomial	\mathbb{R}^d	$\langle \mathbf{x}, \mathbf{x}' \rangle^2$	$n = \frac{d(d+1)}{2}$
RBF	\mathbb{R}^d	$\exp\left(-\gamma \ \mathbf{x} - \mathbf{x}'\ ^2\right)$	$n = \infty$
string sub-sequence	$\cup_{d=0}^{\infty} \Sigma^d$	dynamic programming	$n = \Sigma ^q$

Positive definite kernel

Definition 1. Let \mathcal{X} be a non-empty set. The function $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a *positive definite kernel* if it is *symmetric* and for any finite set of inputs x^1, \dots, x^m , the *kernel matrix* $\mathbf{K} \in \mathbb{R}^{m \times m}$ with elements $K_{i,j} = k(x^i, x^j)$ is *positive semi-definite*.

- ◆ The kernel matrix $\mathbf{K} \in \mathbb{R}^{m \times m}$ represents similarities between each pair of inputs $\{x^1, \dots, x^m\}$:

$$\mathbf{K} = \begin{pmatrix} k(x^1, x^1), & k(x^1, x^2), & \dots, & k(x^1, x^m) \\ k(x^2, x^1), & k(x^2, x^2), & \dots, & k(x^2, x^m) \\ \vdots & & & \\ k(x^m, x^1), & k(x^m, x^2), & \dots, & k(x^m, x^m) \end{pmatrix}$$

- ◆ A matrix $\mathbf{K} \in \mathbb{R}^{m \times m}$ is PSD if for every $\alpha \in \mathbb{R}^m$, $\alpha^T \mathbf{K} \alpha \geq 0$.

Every kernel defines a feature space

Theorem 1. *For every positive definite kernel $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ there exists a Hilbert space \mathcal{H} and a feature map $\phi: \mathcal{X} \rightarrow \mathcal{H}$ such that $k(x, x') = \langle \phi(x), \phi(x') \rangle$.*

Proof for a kernel defined on a finite input space \mathcal{X} :

The kernel matrix $\mathbf{K} \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$ is symmetric and PSD hence the spectral decomposition exists $\mathbf{K} = \mathbf{V}\mathbf{D}\mathbf{V}^T$ where $\mathbf{V} \in \mathbb{R}^{m \times m}$ is orthogonal and $\mathbf{D} = \text{diag}(\lambda_1, \dots, \lambda_{|\mathcal{X}|})$ is diagonal matrix of non-negative eigenvalues.

Therefore $\mathbf{K} = \Phi^T \Phi$ where $\Phi^T = \mathbf{V}\mathbf{D}^{\frac{1}{2}}$.

Construction of valid kernels

Let $k_1: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $k_2: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be p.d. kernels, $\lambda \in [0, 1]$, $\alpha \geq 0$, $\phi: \mathcal{X} \rightarrow \mathbb{R}^n$ and $k_3: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ a p.d. kernel, \mathbf{K} a symmetric positive definite matrix. Then the following functions are also p.d. kernels:

$$k(x, z) = (1 - \lambda) k_1(x, z) + \lambda k_2(x, z)$$

$$k(x, z) = \alpha k_1(x, z)$$

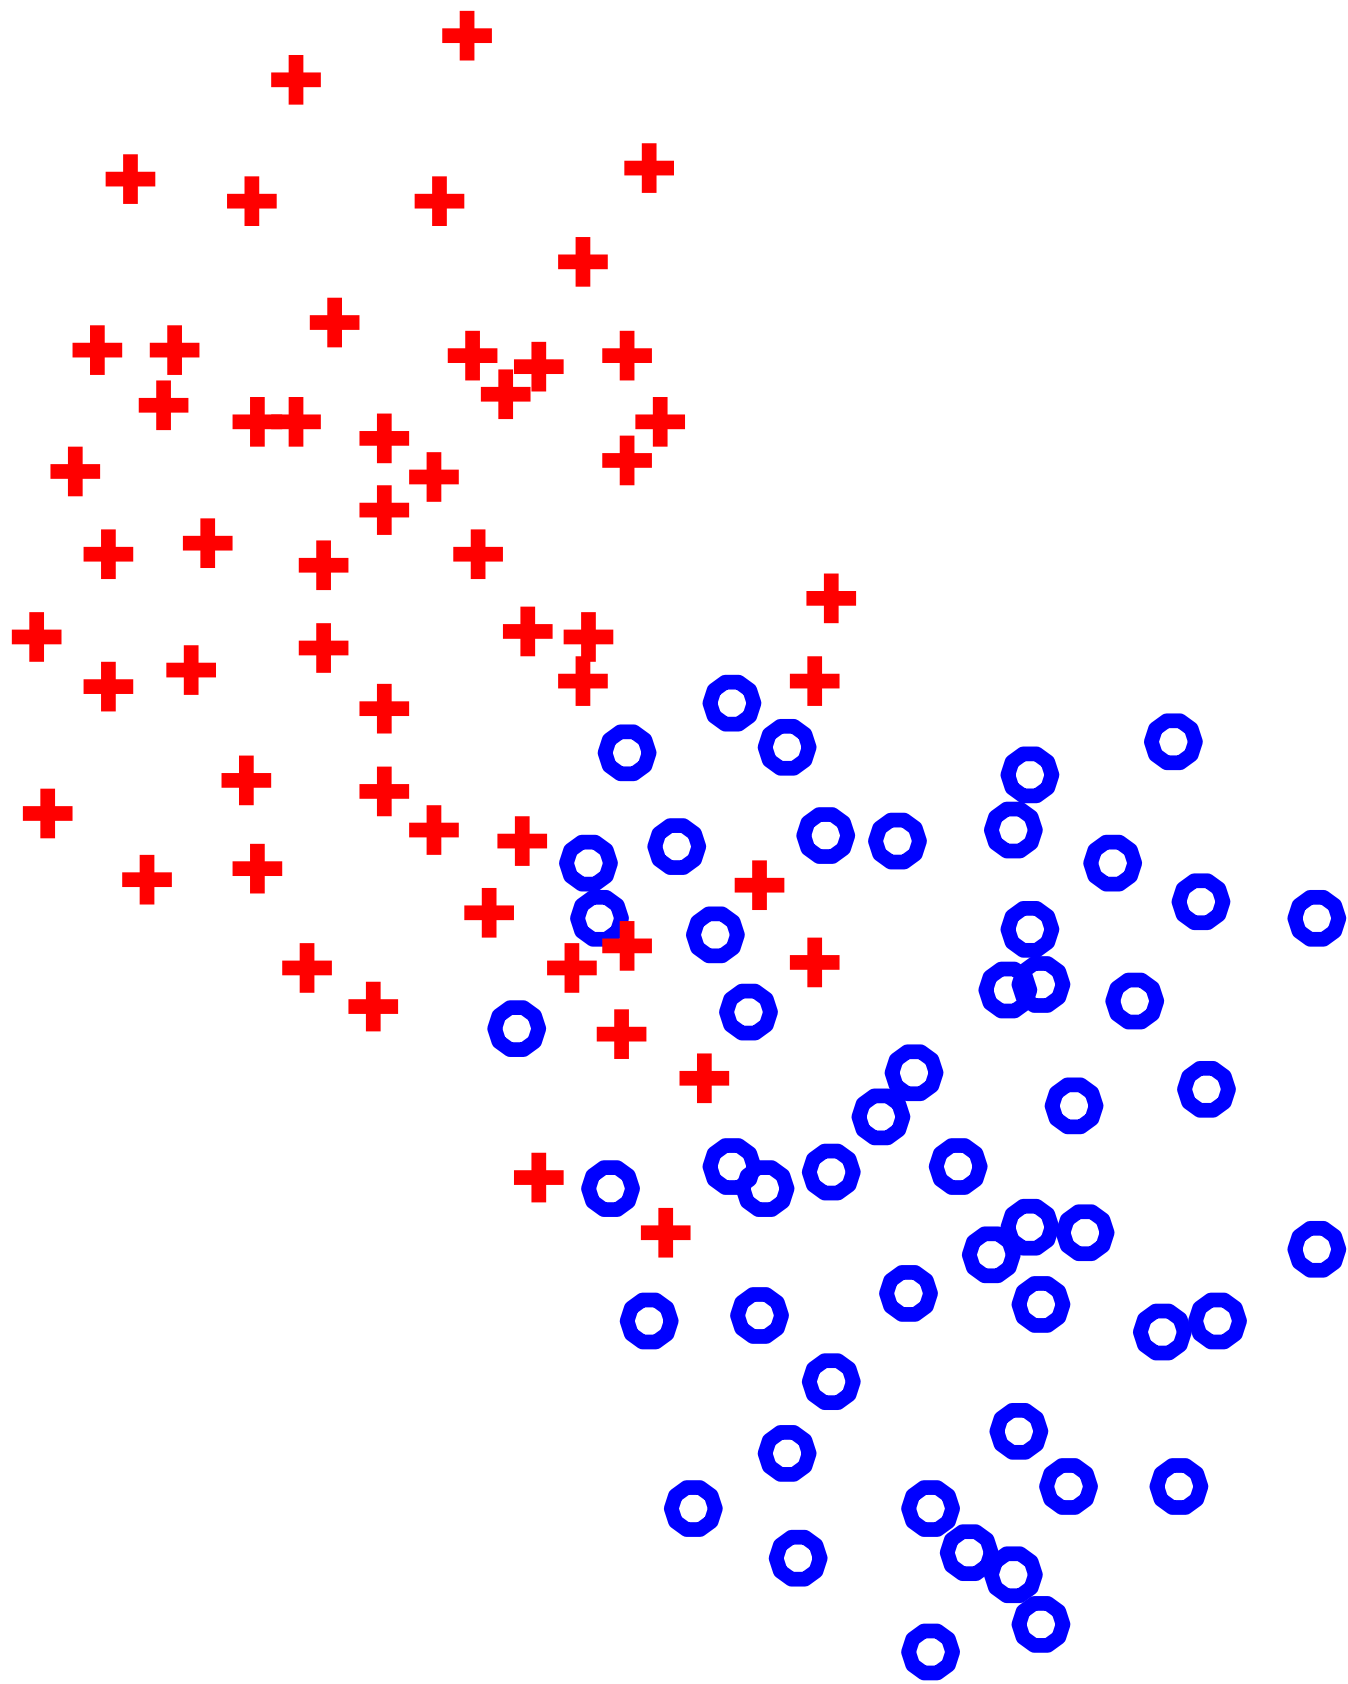
$$k(x, z) = k_1(x, z) k_2(x, z)$$

$$k(x, z) = k_3(\phi(x), \phi(z))$$

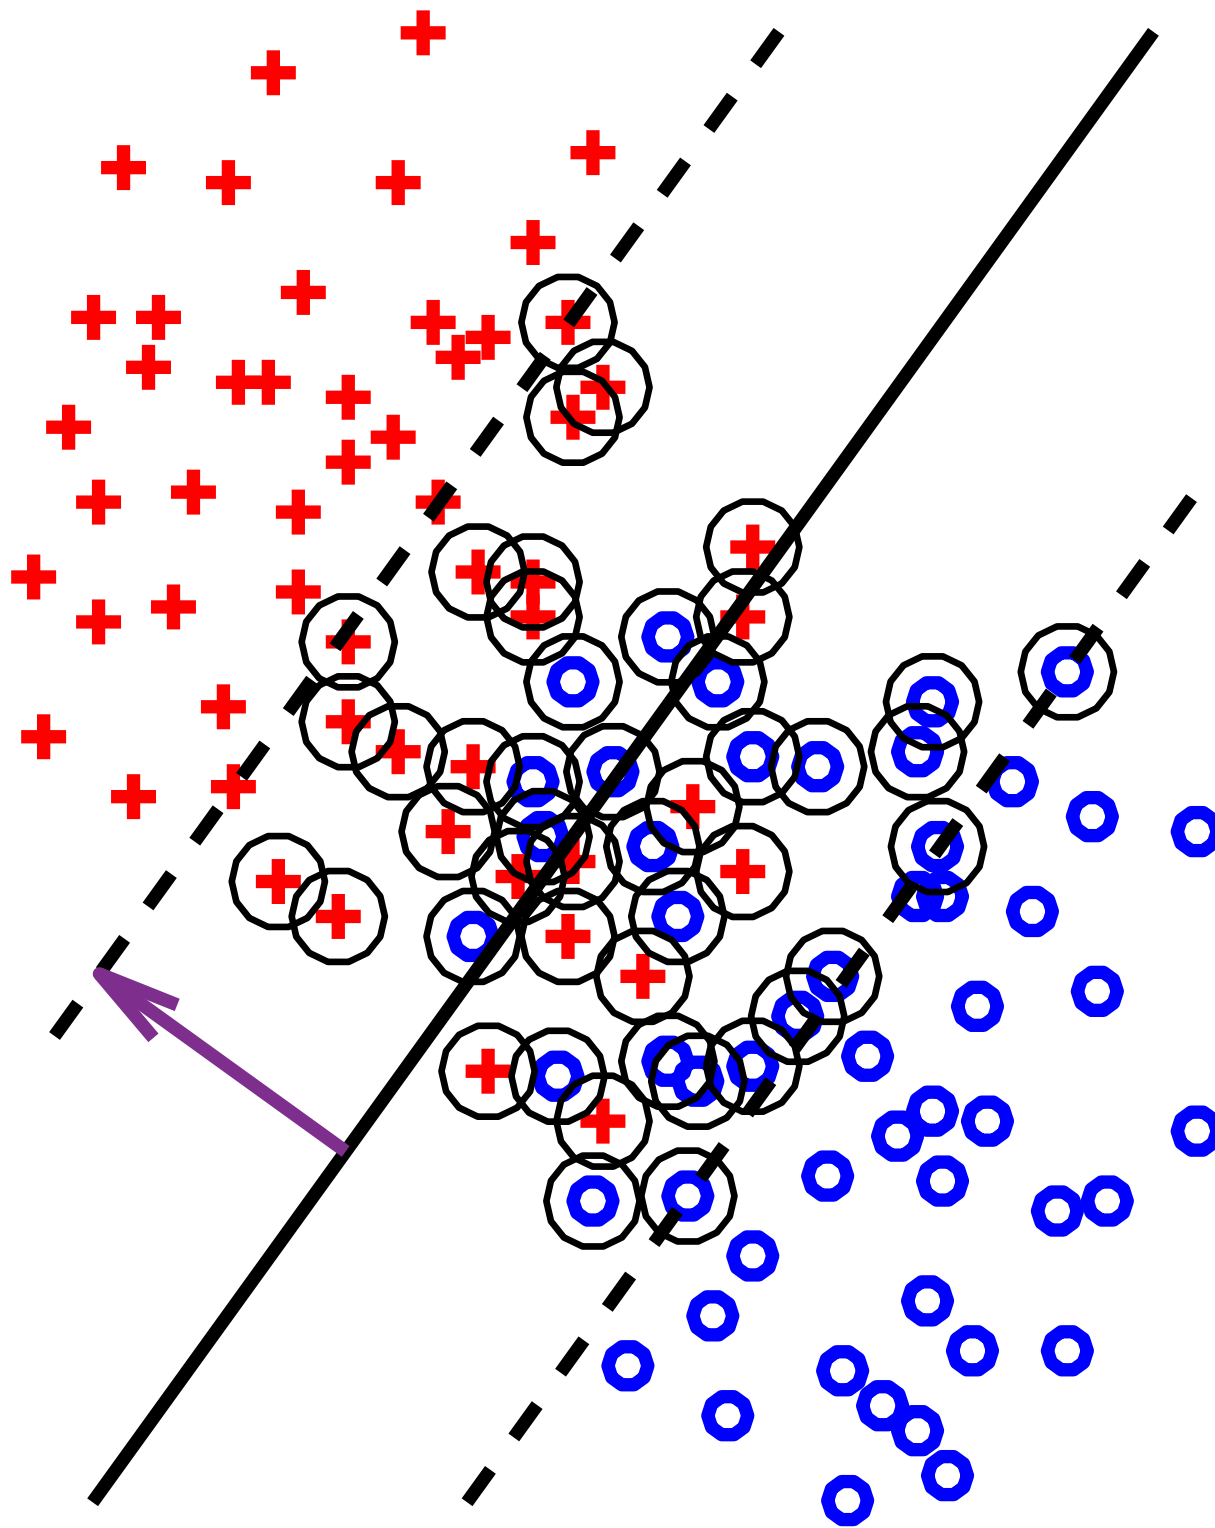
$$k(\mathbf{x}, \mathbf{z}) = \mathbf{x}^T \mathbf{K} \mathbf{z}$$

Summary

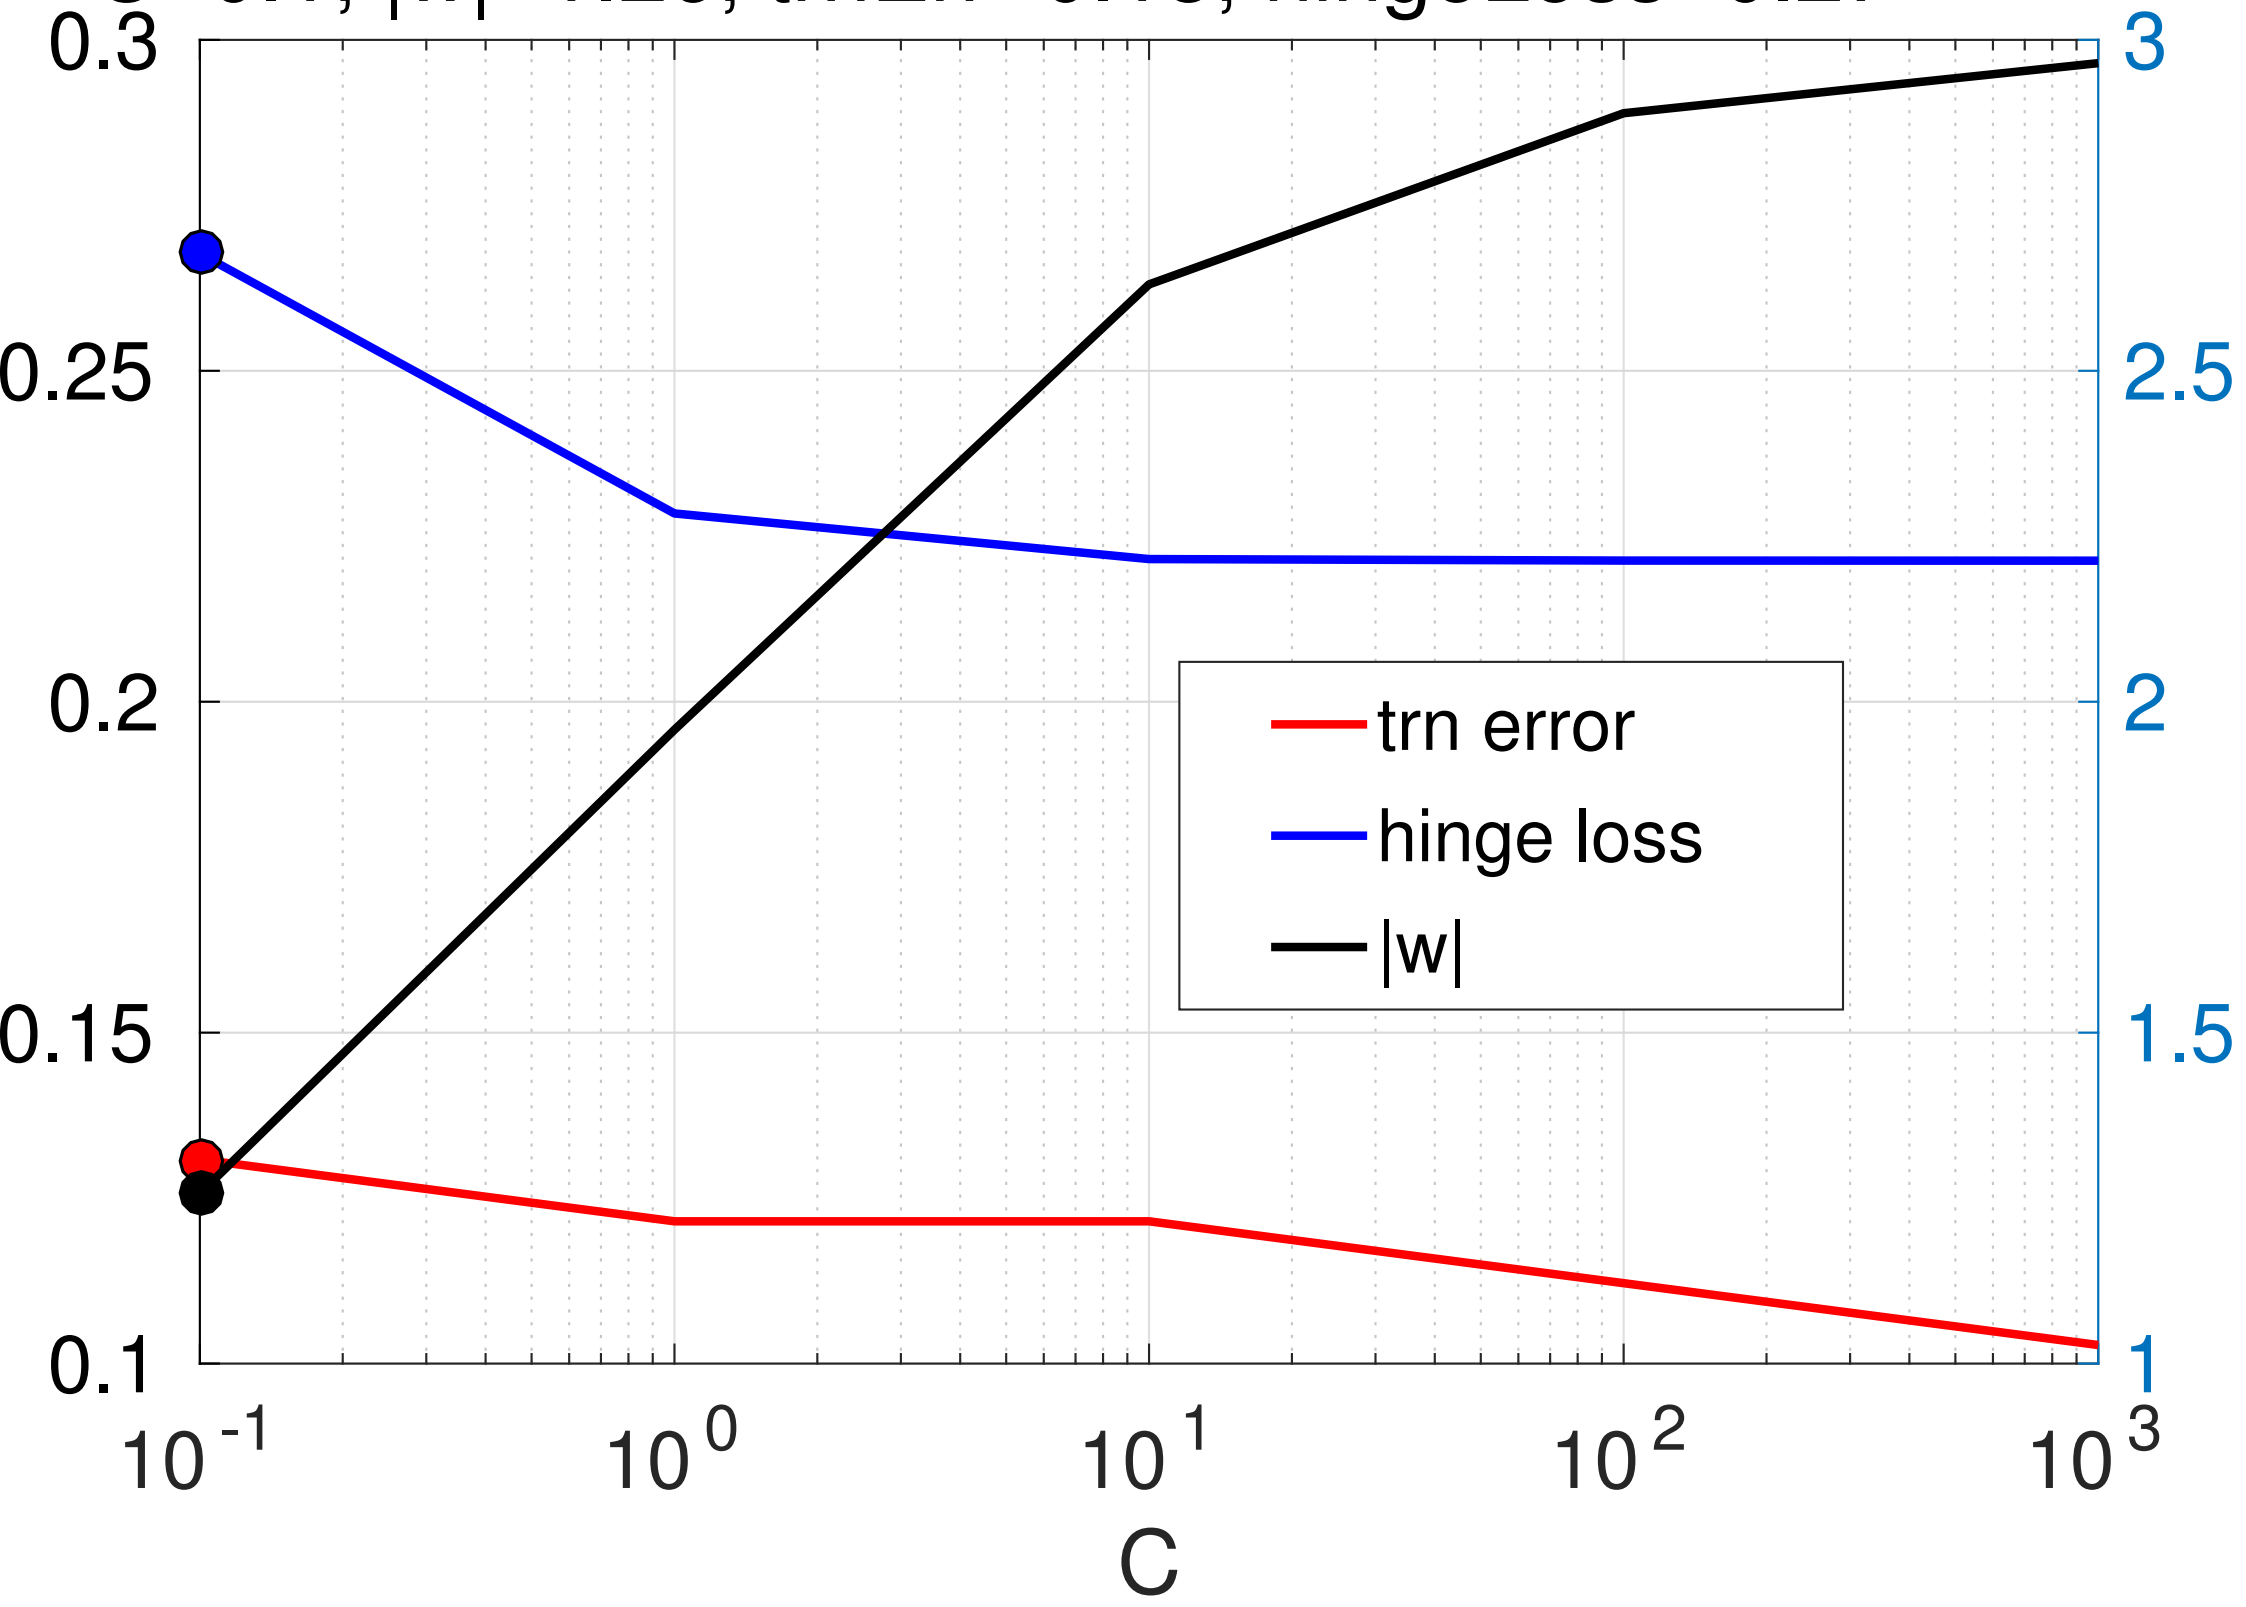
- ◆ Formulation of SVM as a convex Quadratic Program
- ◆ Dual formulation of SVM that uses only dot products of data
- ◆ Kernel trick: replace dot product by kernel function
- ◆ Examples of kernel functions: polynomial, RBF, strings
- ◆ Positive definite kernels



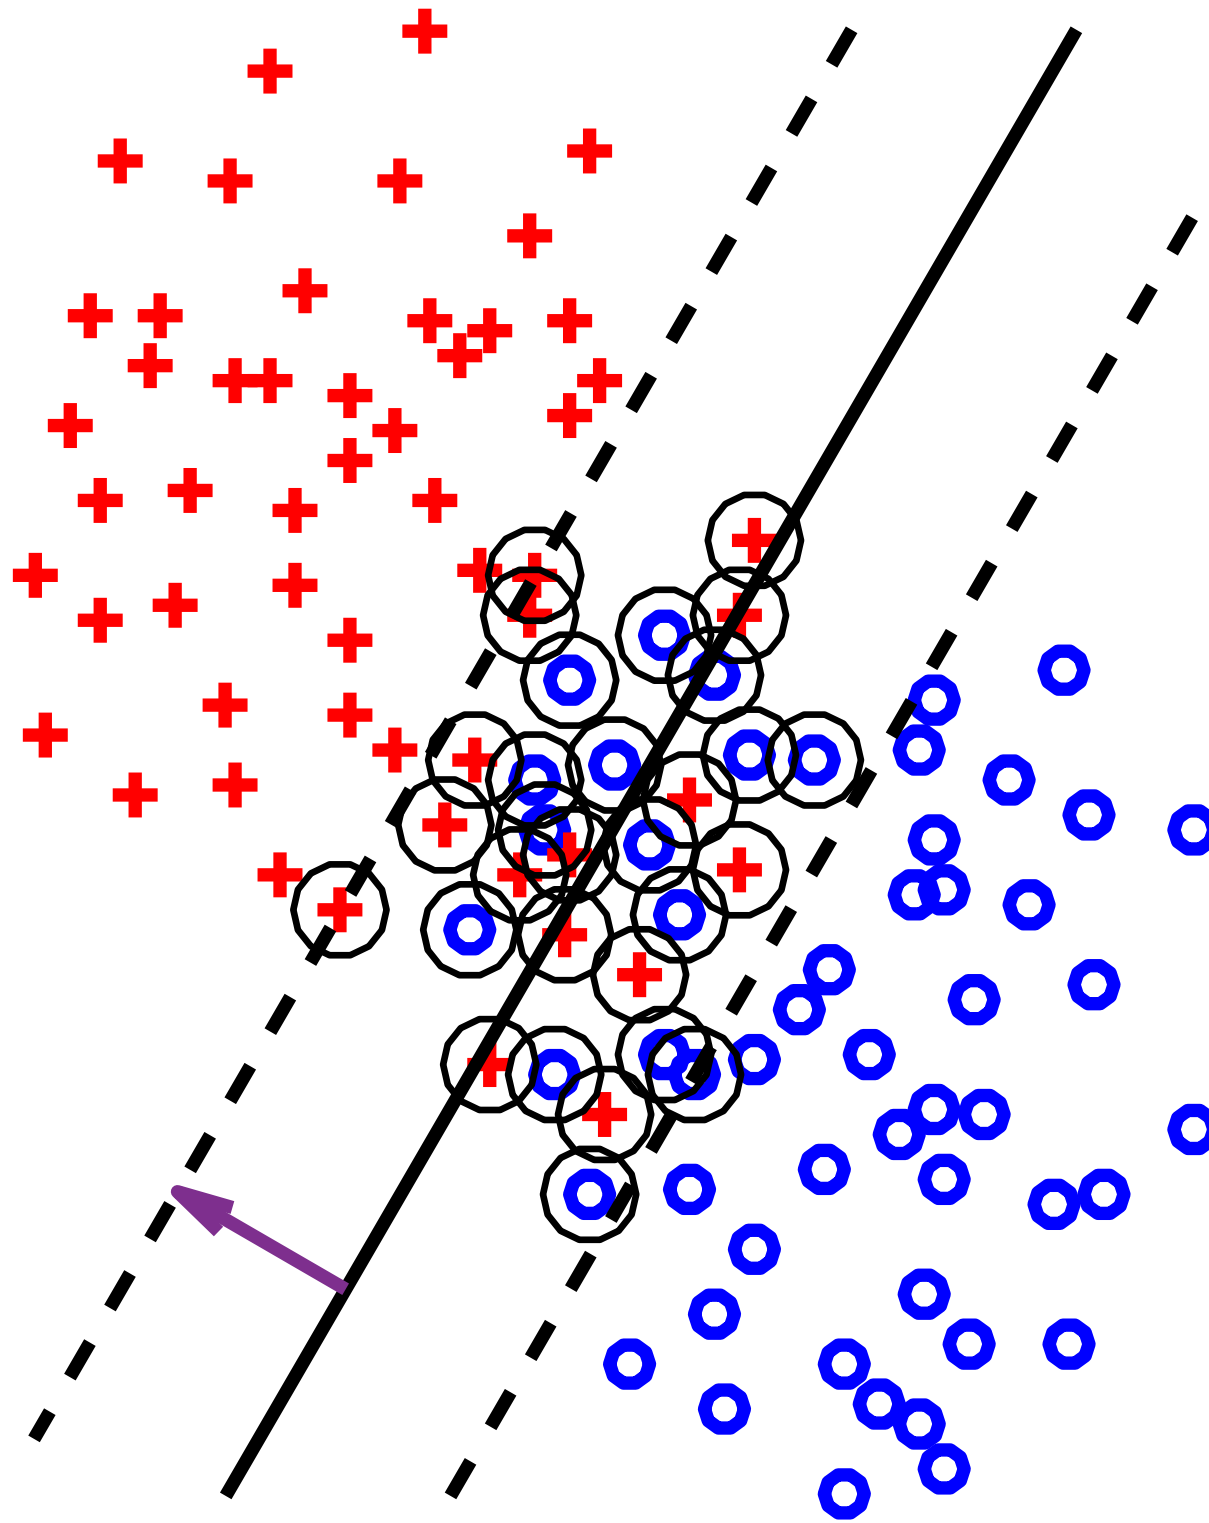
$1/|w|=0.80$



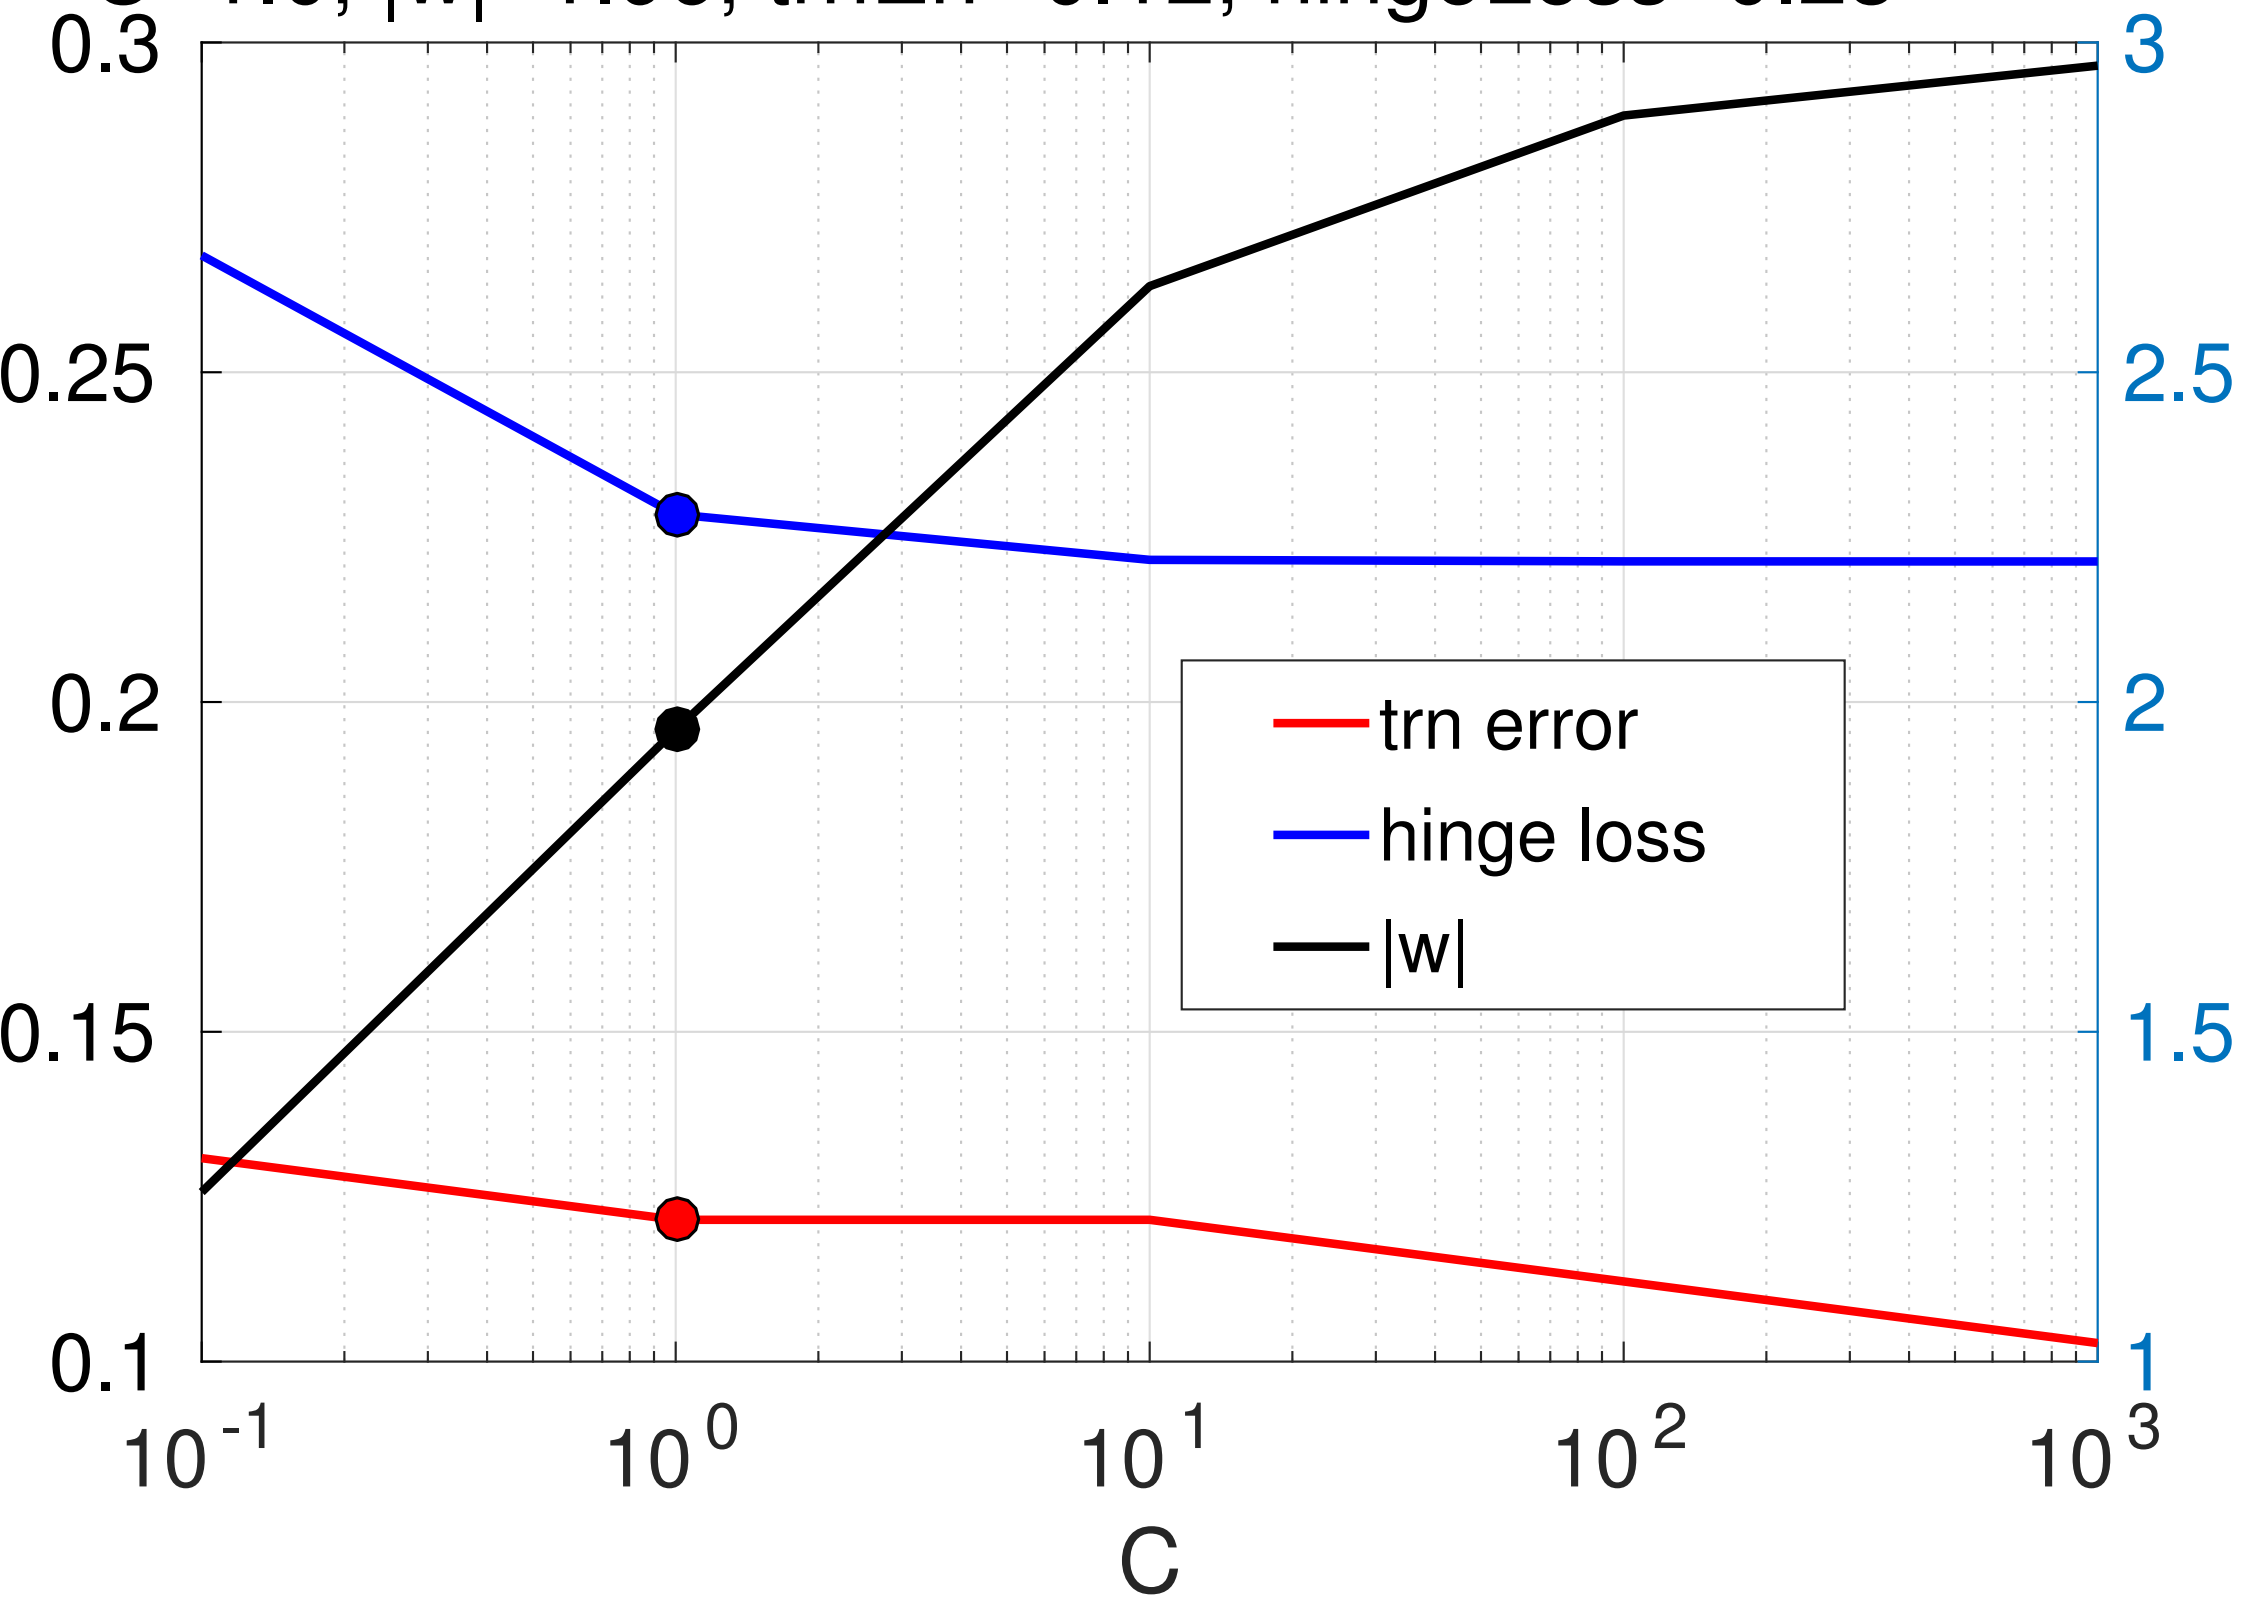
$C=0.1$, $|w|=1.26$, $\text{trnErr}=0.13$, $\text{hingeLoss}=0.27$



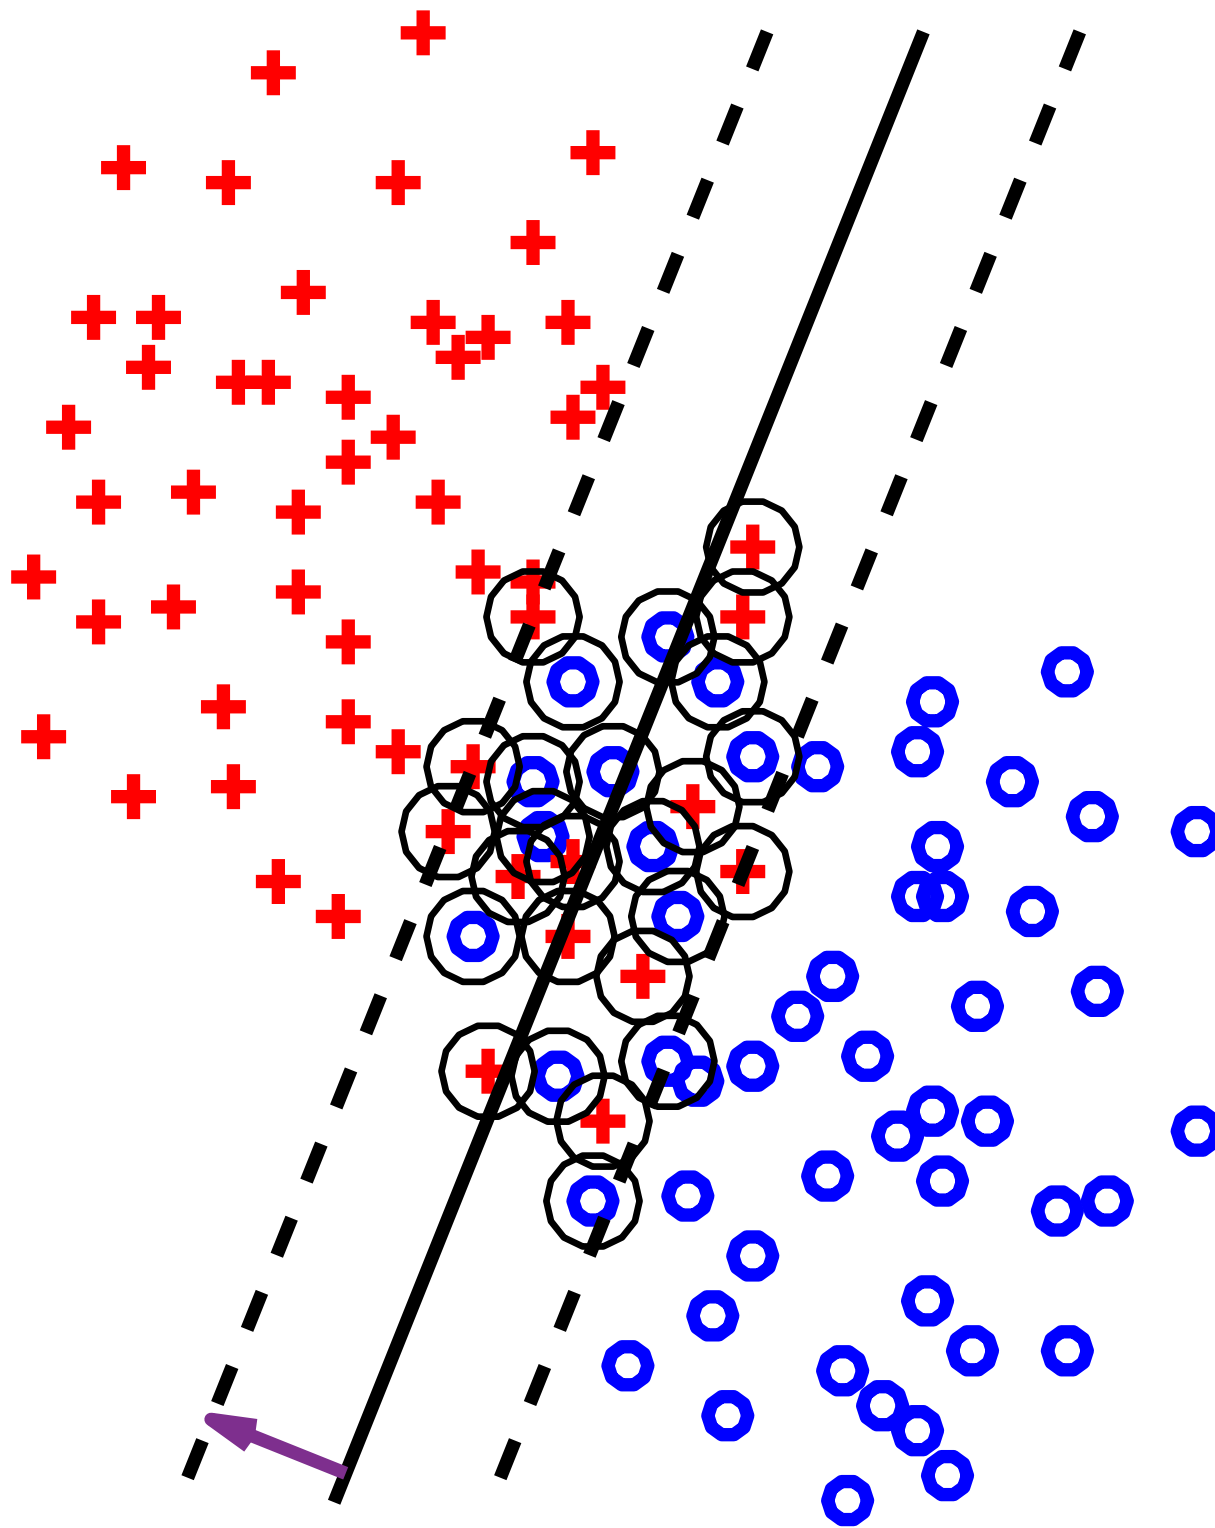
$$1/|w|=0.51$$



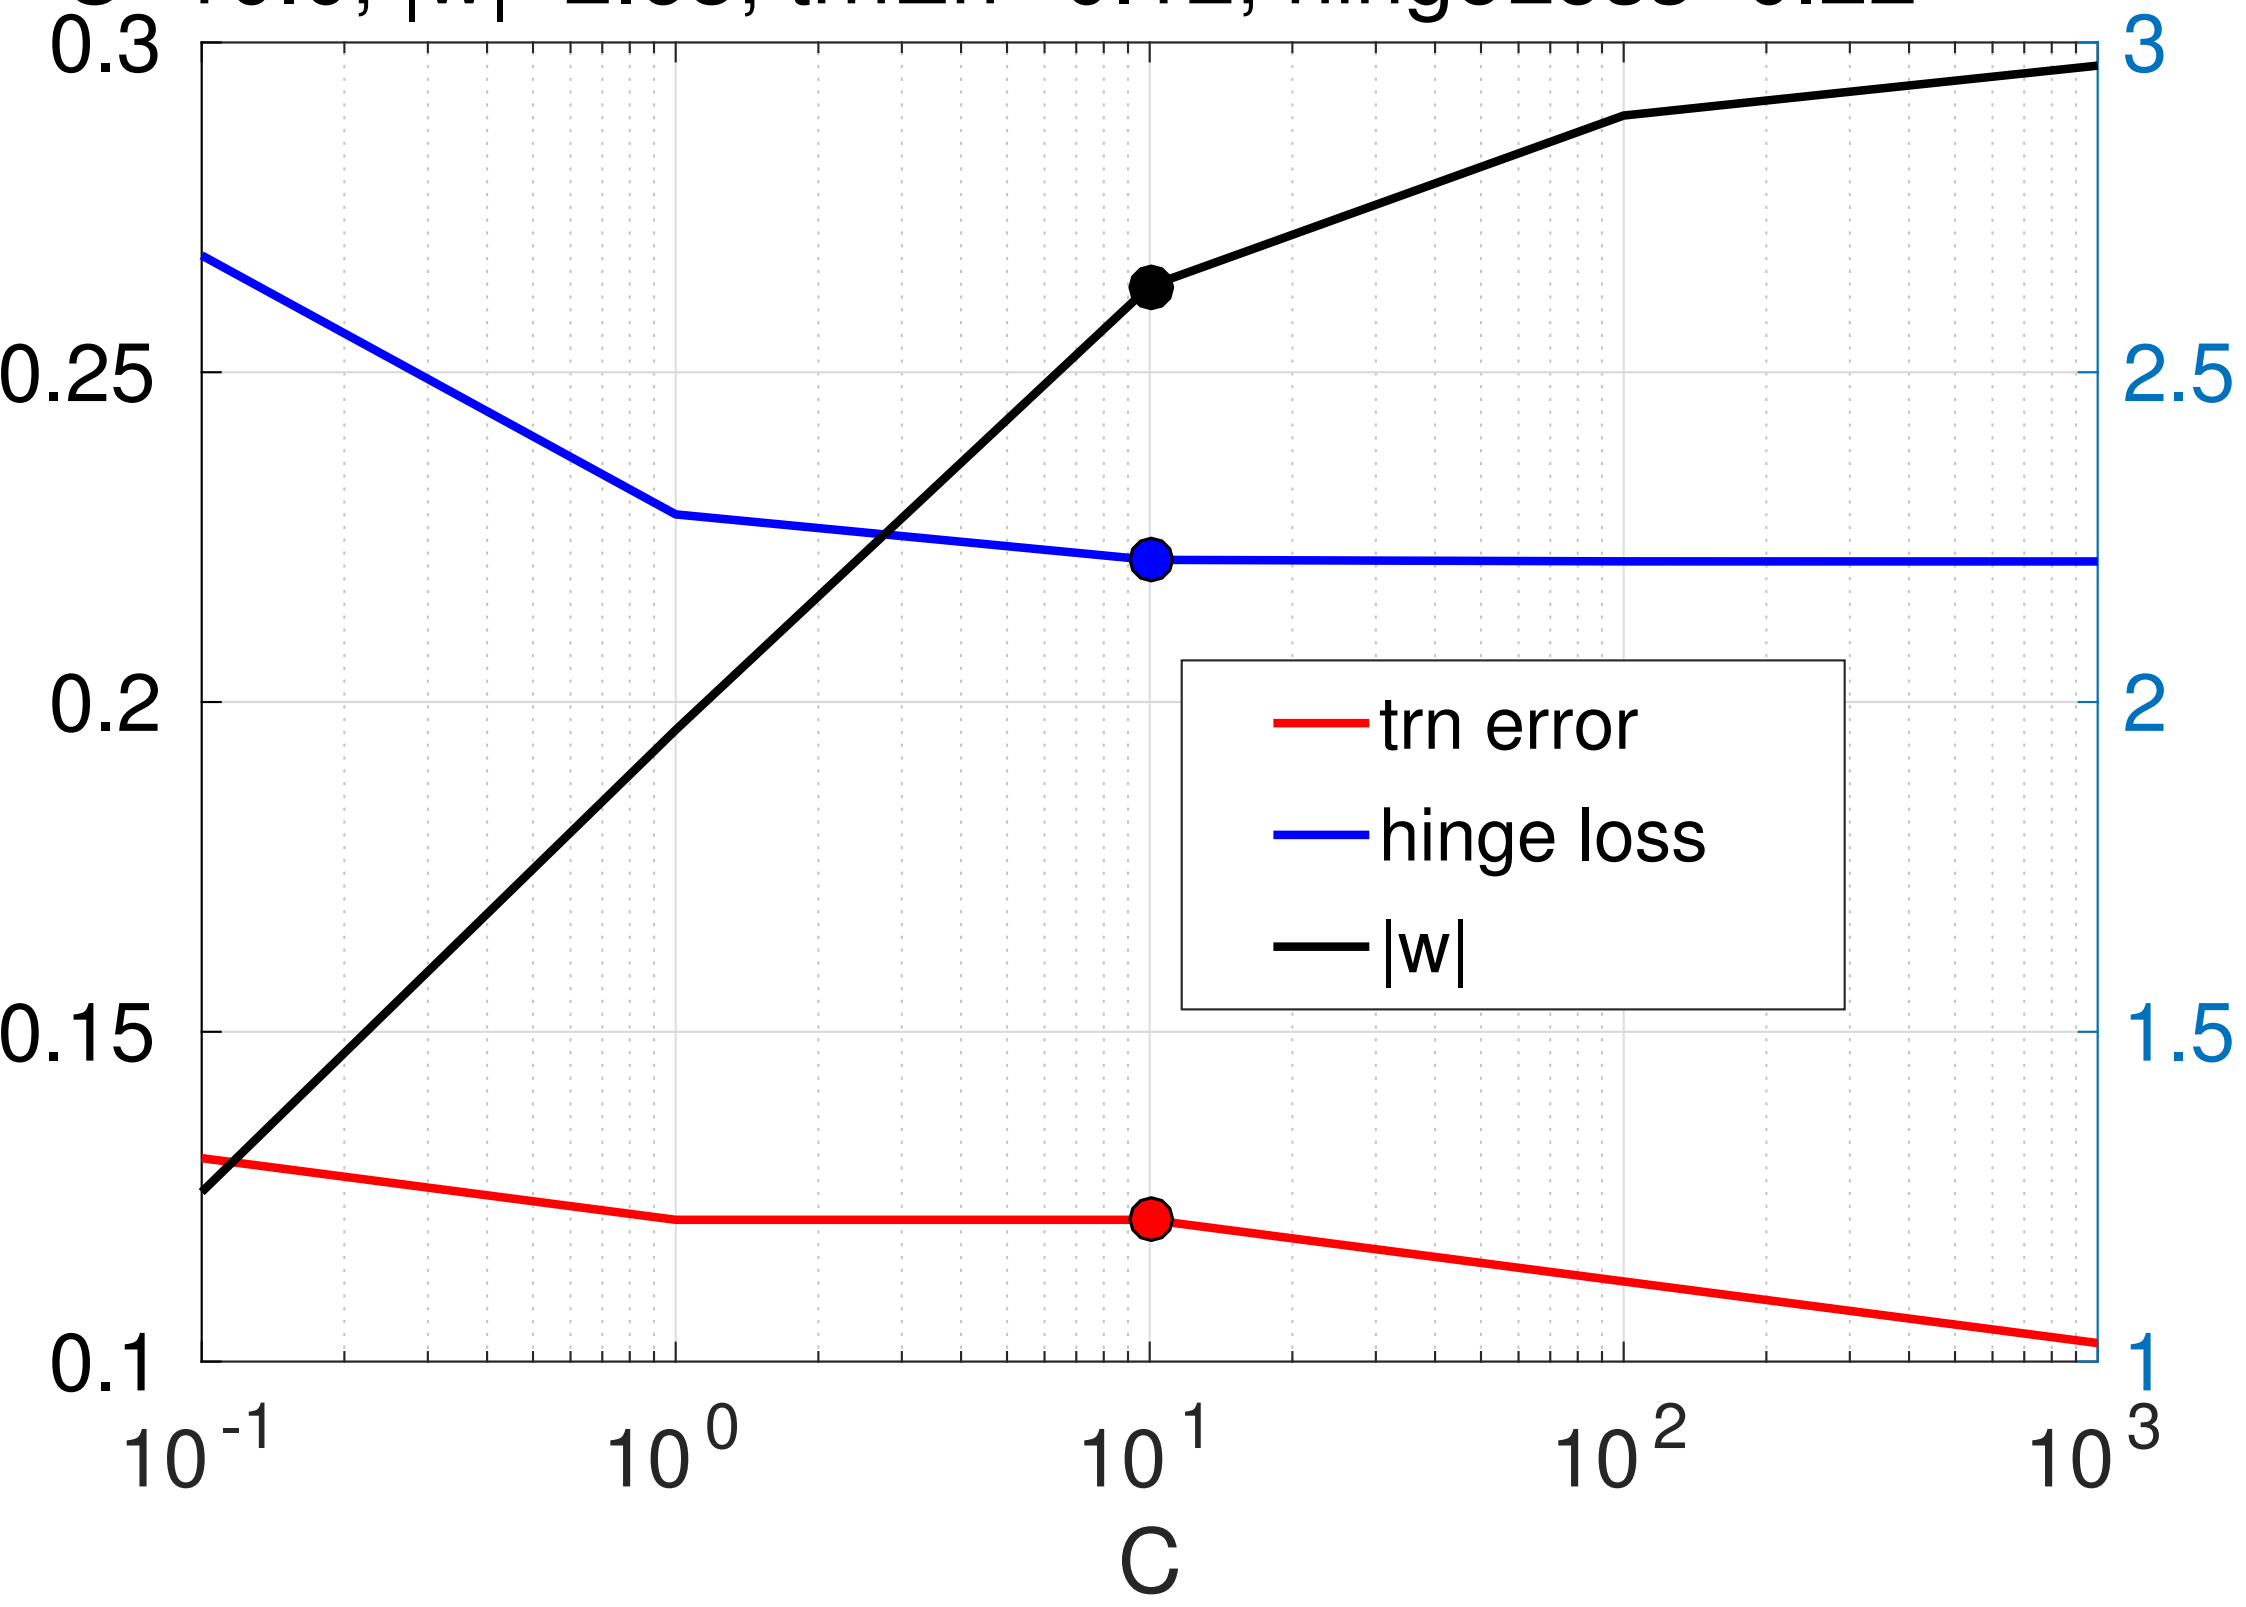
$C=1.0$, $|w|=1.96$, $\text{trnErr}=0.12$, $\text{hingeLoss}=0.23$



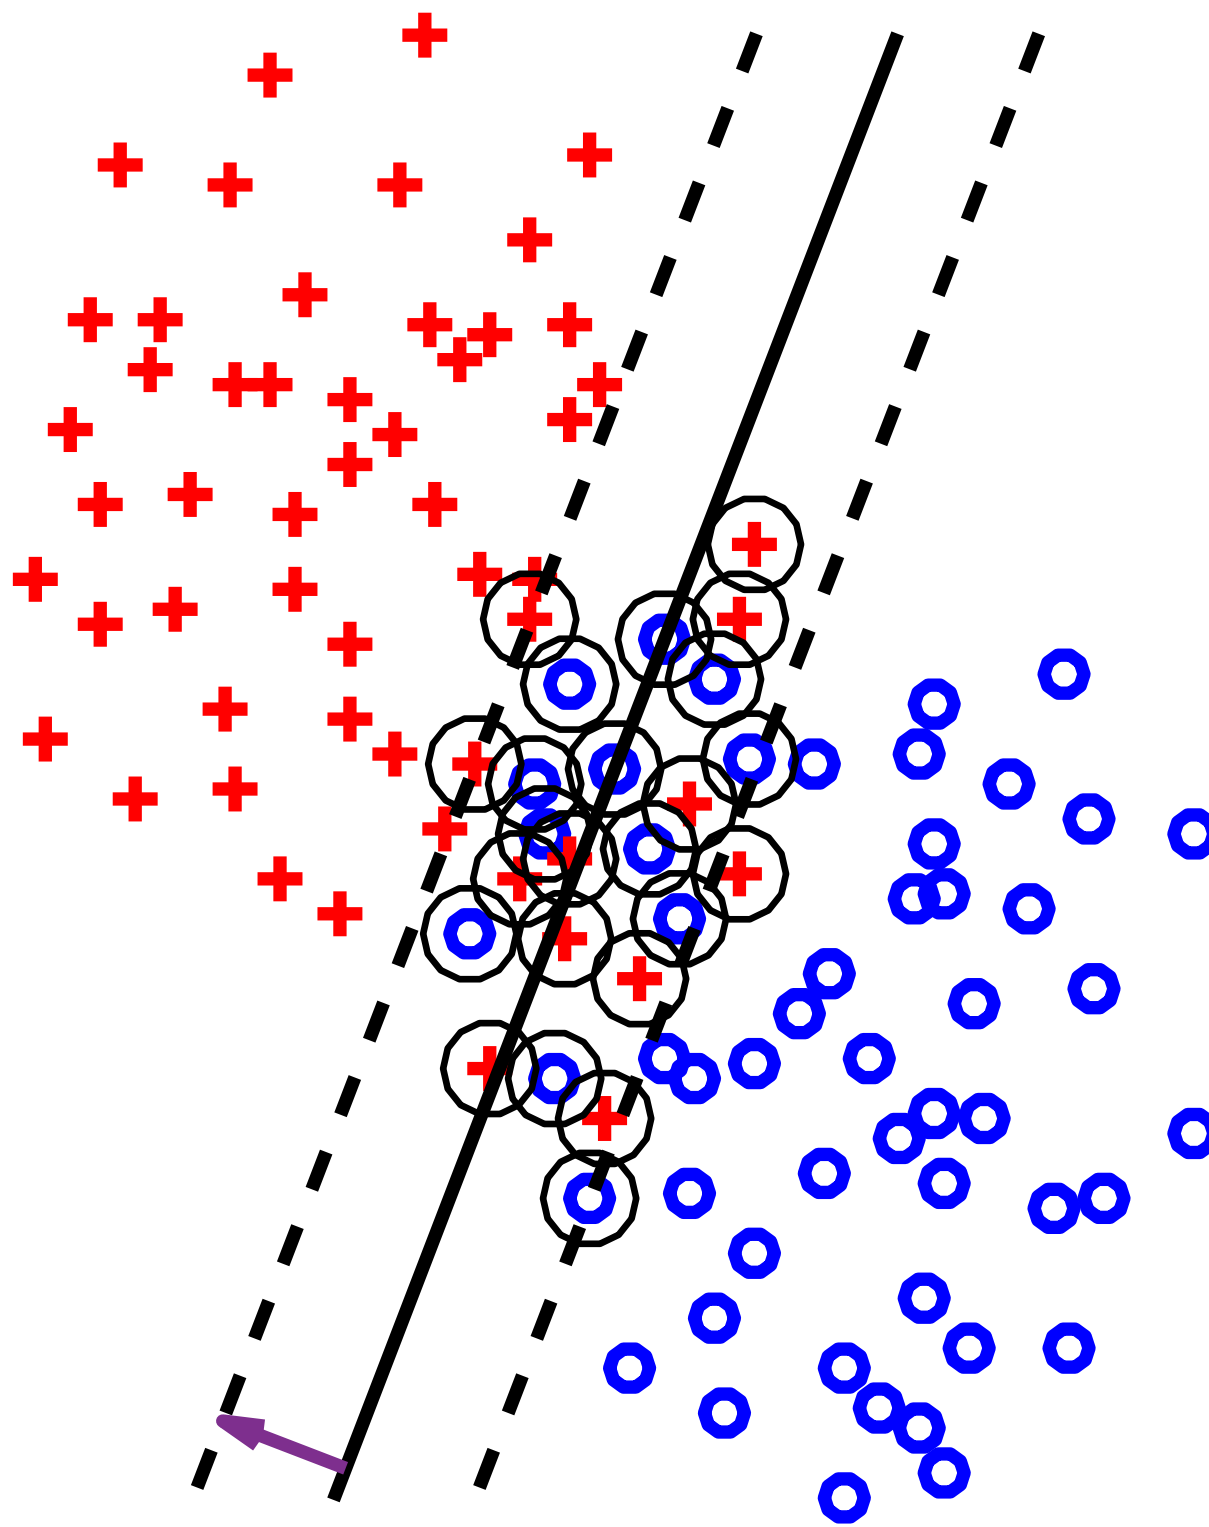
$$1/|w|=0.38$$



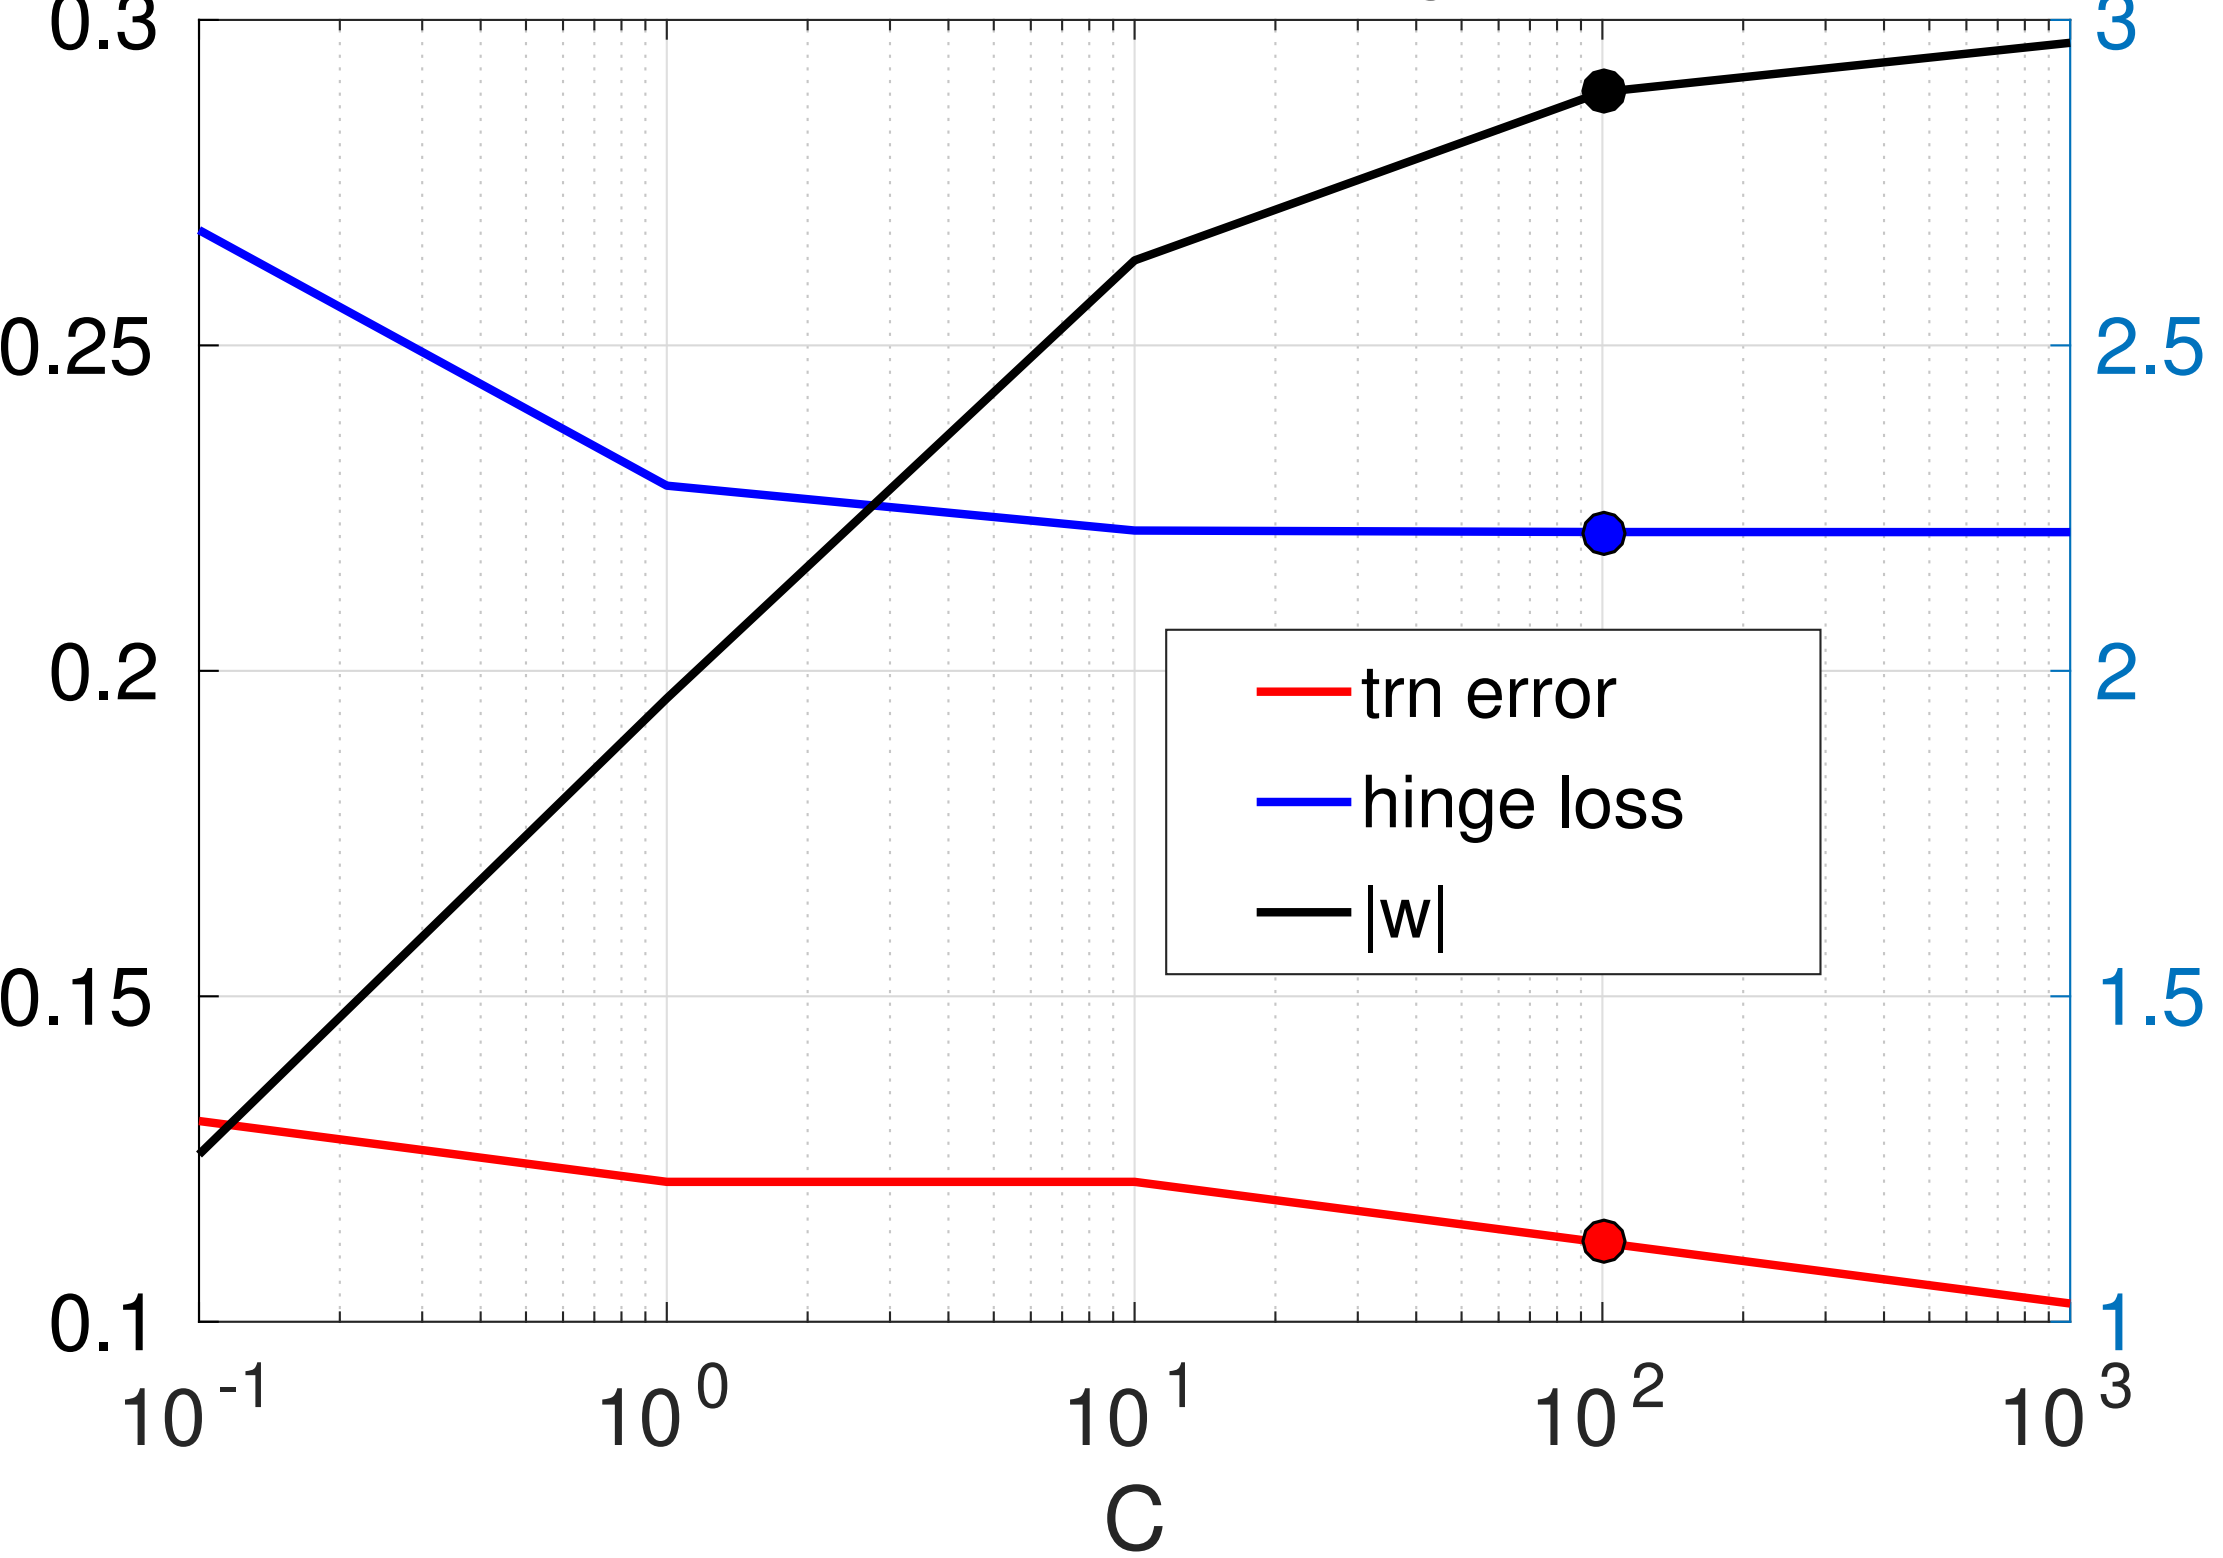
$C=10.0$, $|w|=2.63$, $\text{trnErr}=0.12$, $\text{hingeLoss}=0.22$



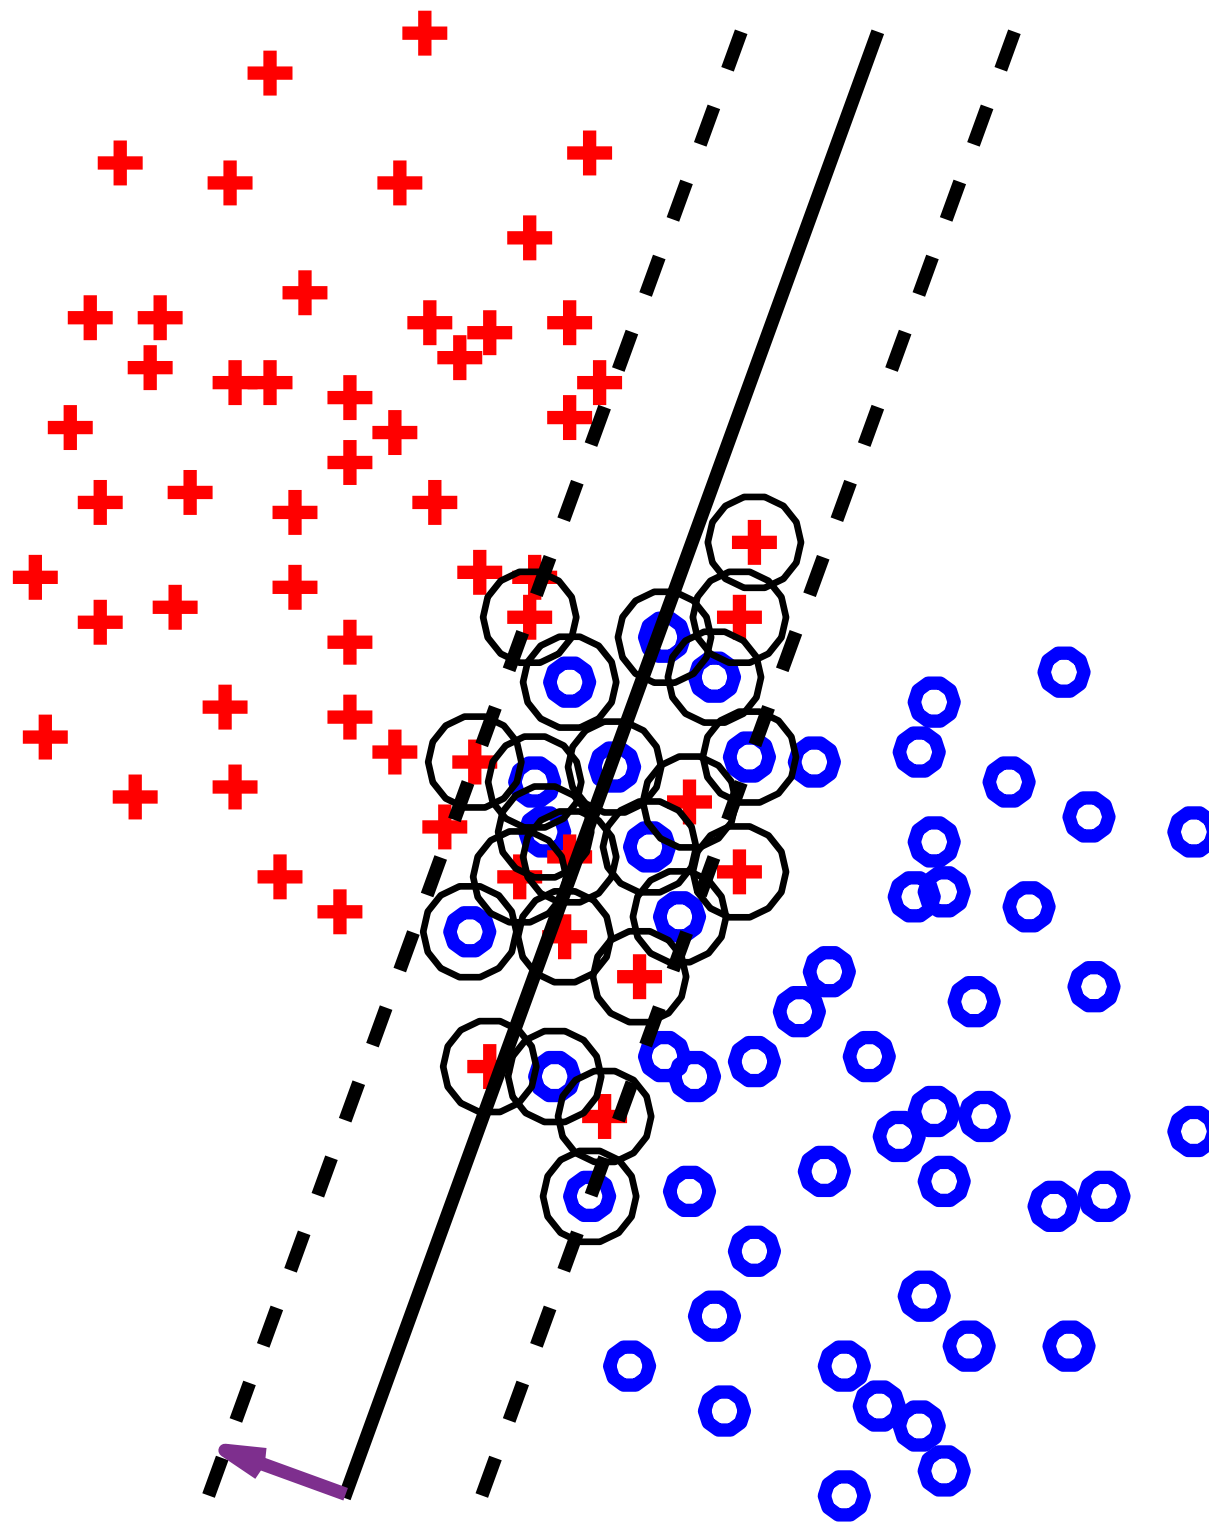
$$1/|w|=0.35$$



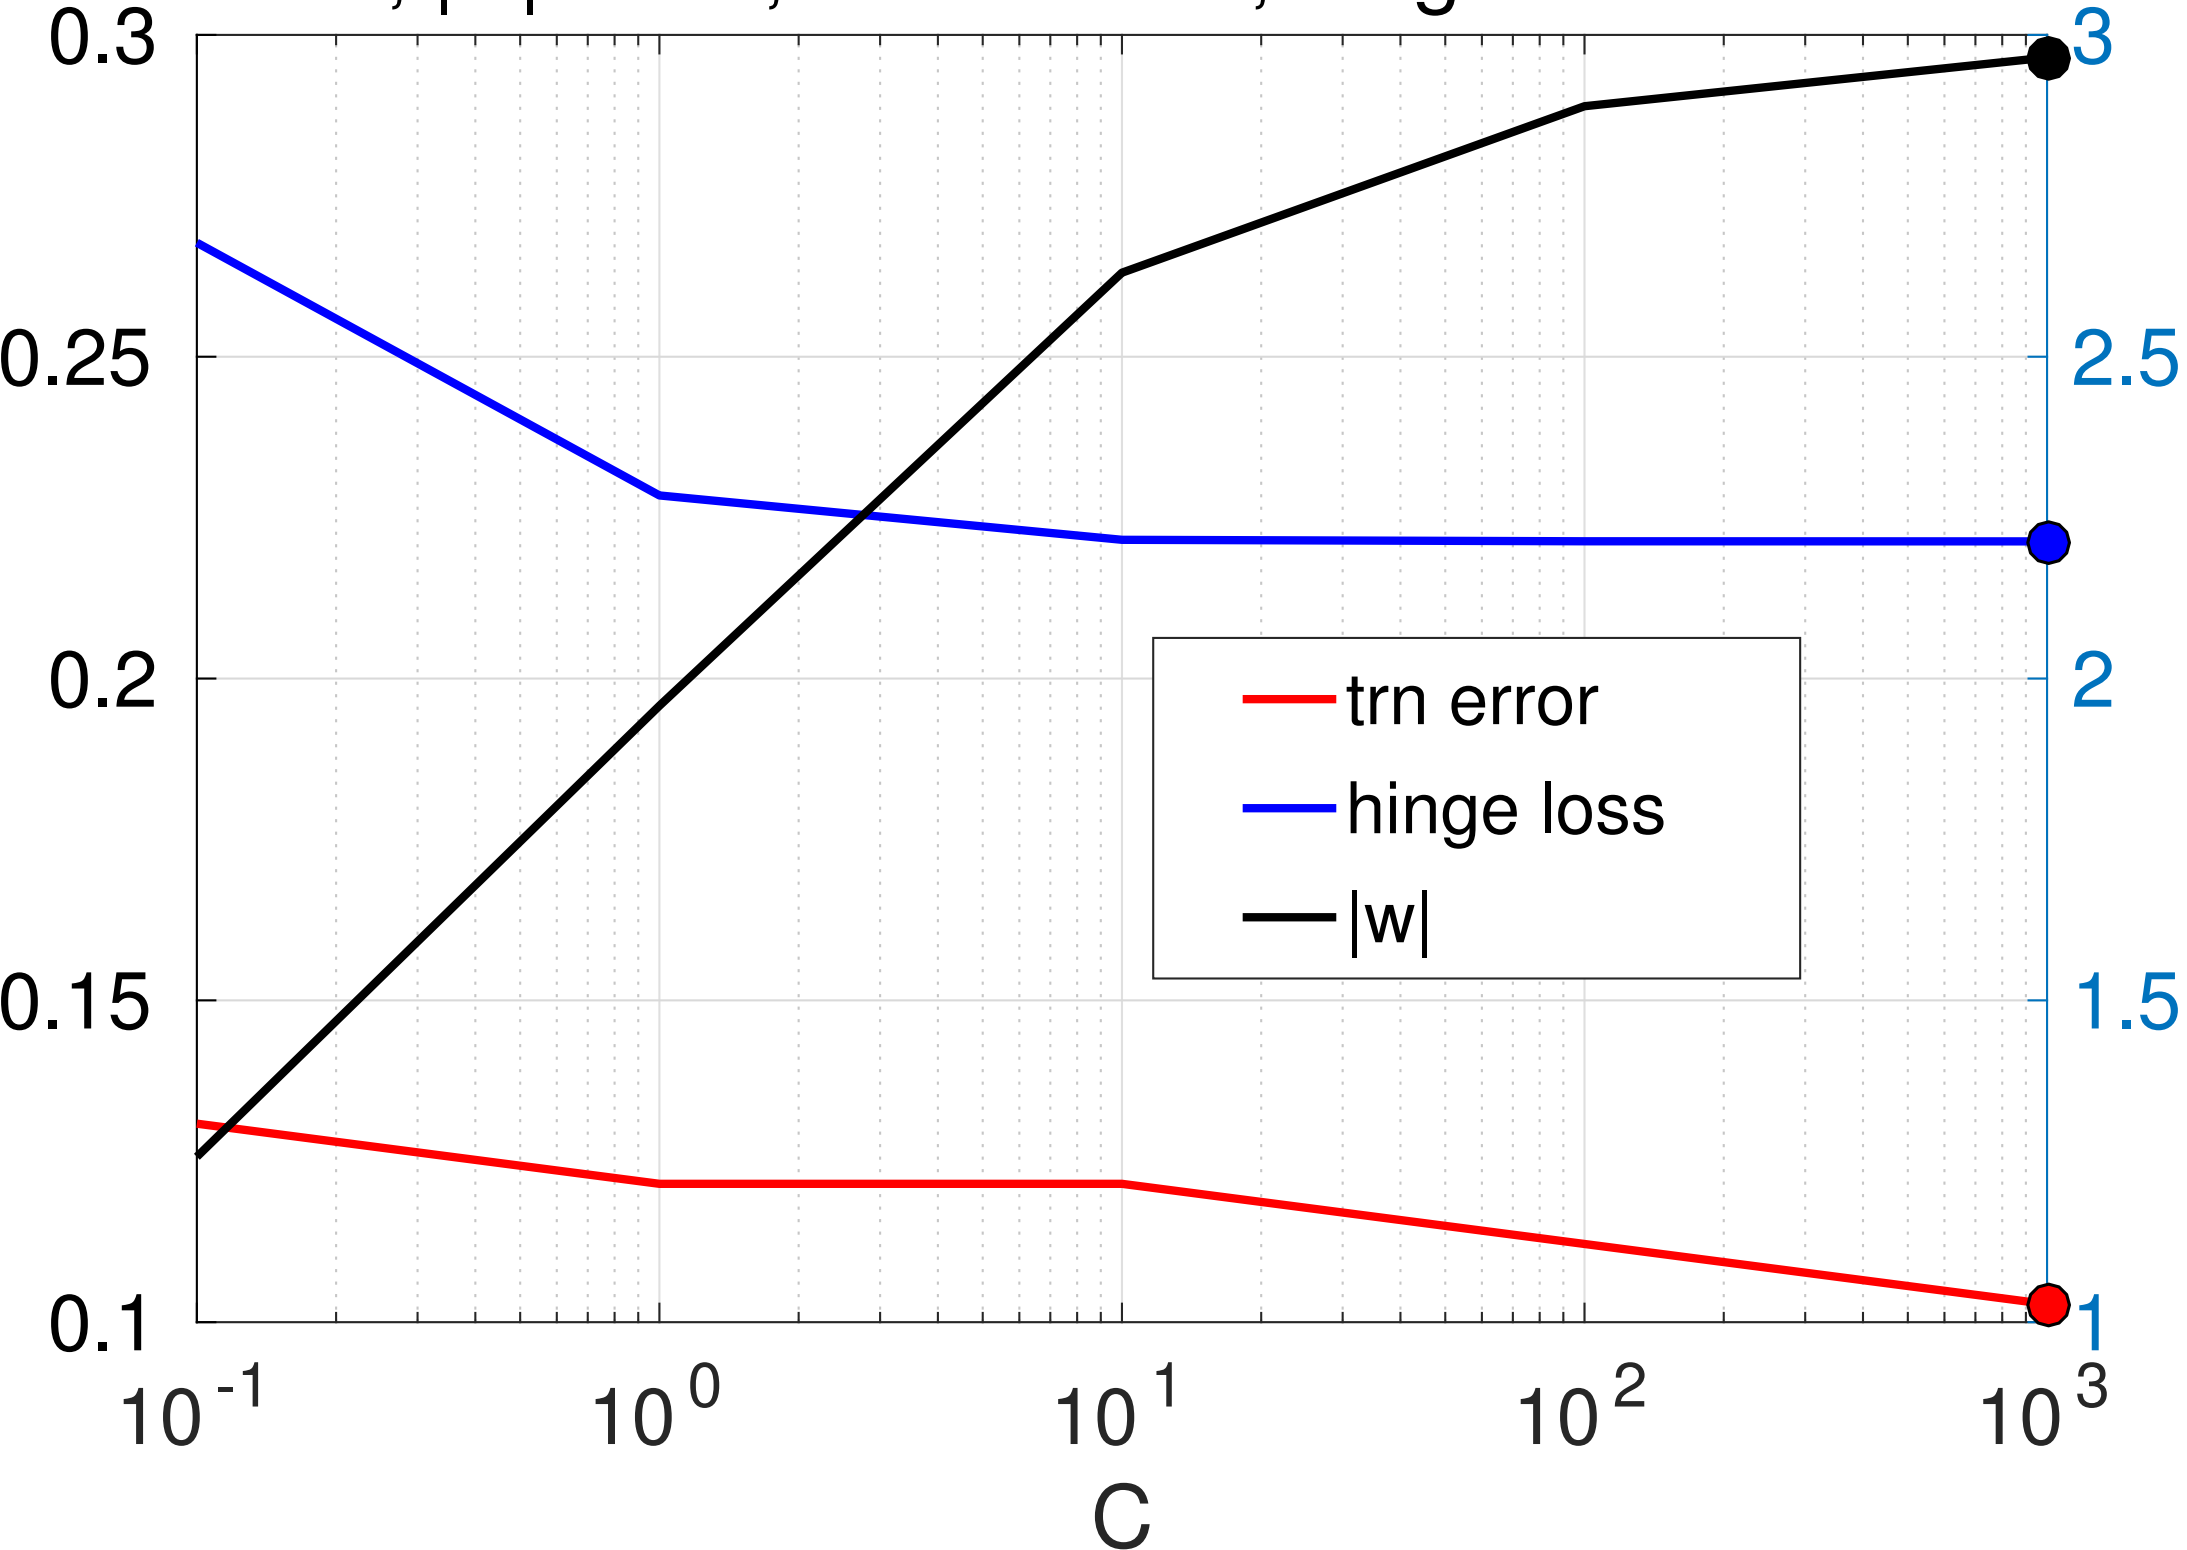
$C=100.0$, $|w|=2.89$, $\text{trnErr}=0.11$, $\text{hingeLoss}=0.22$



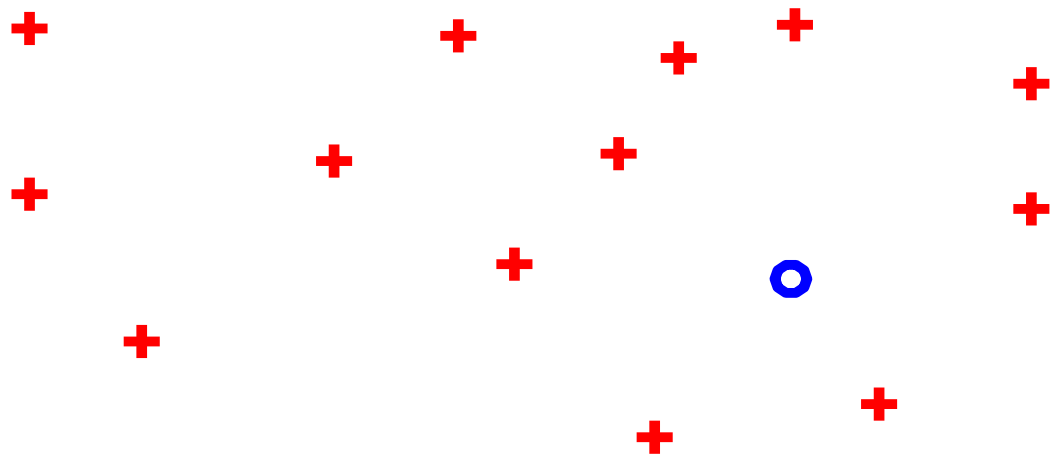
$$1/|w|=0.34$$



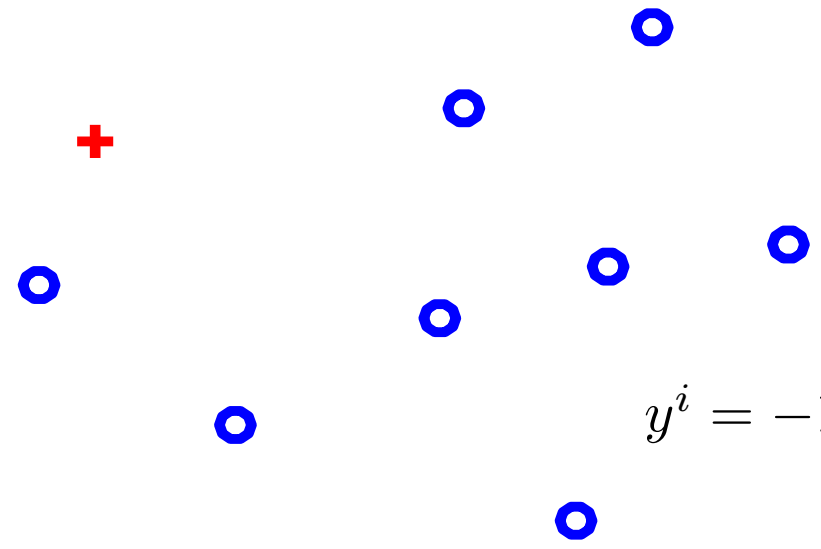
$C=1000.0$, $|w|=2.96$, $\text{trnErr}=0.10$, $\text{hingeLoss}=0.22$

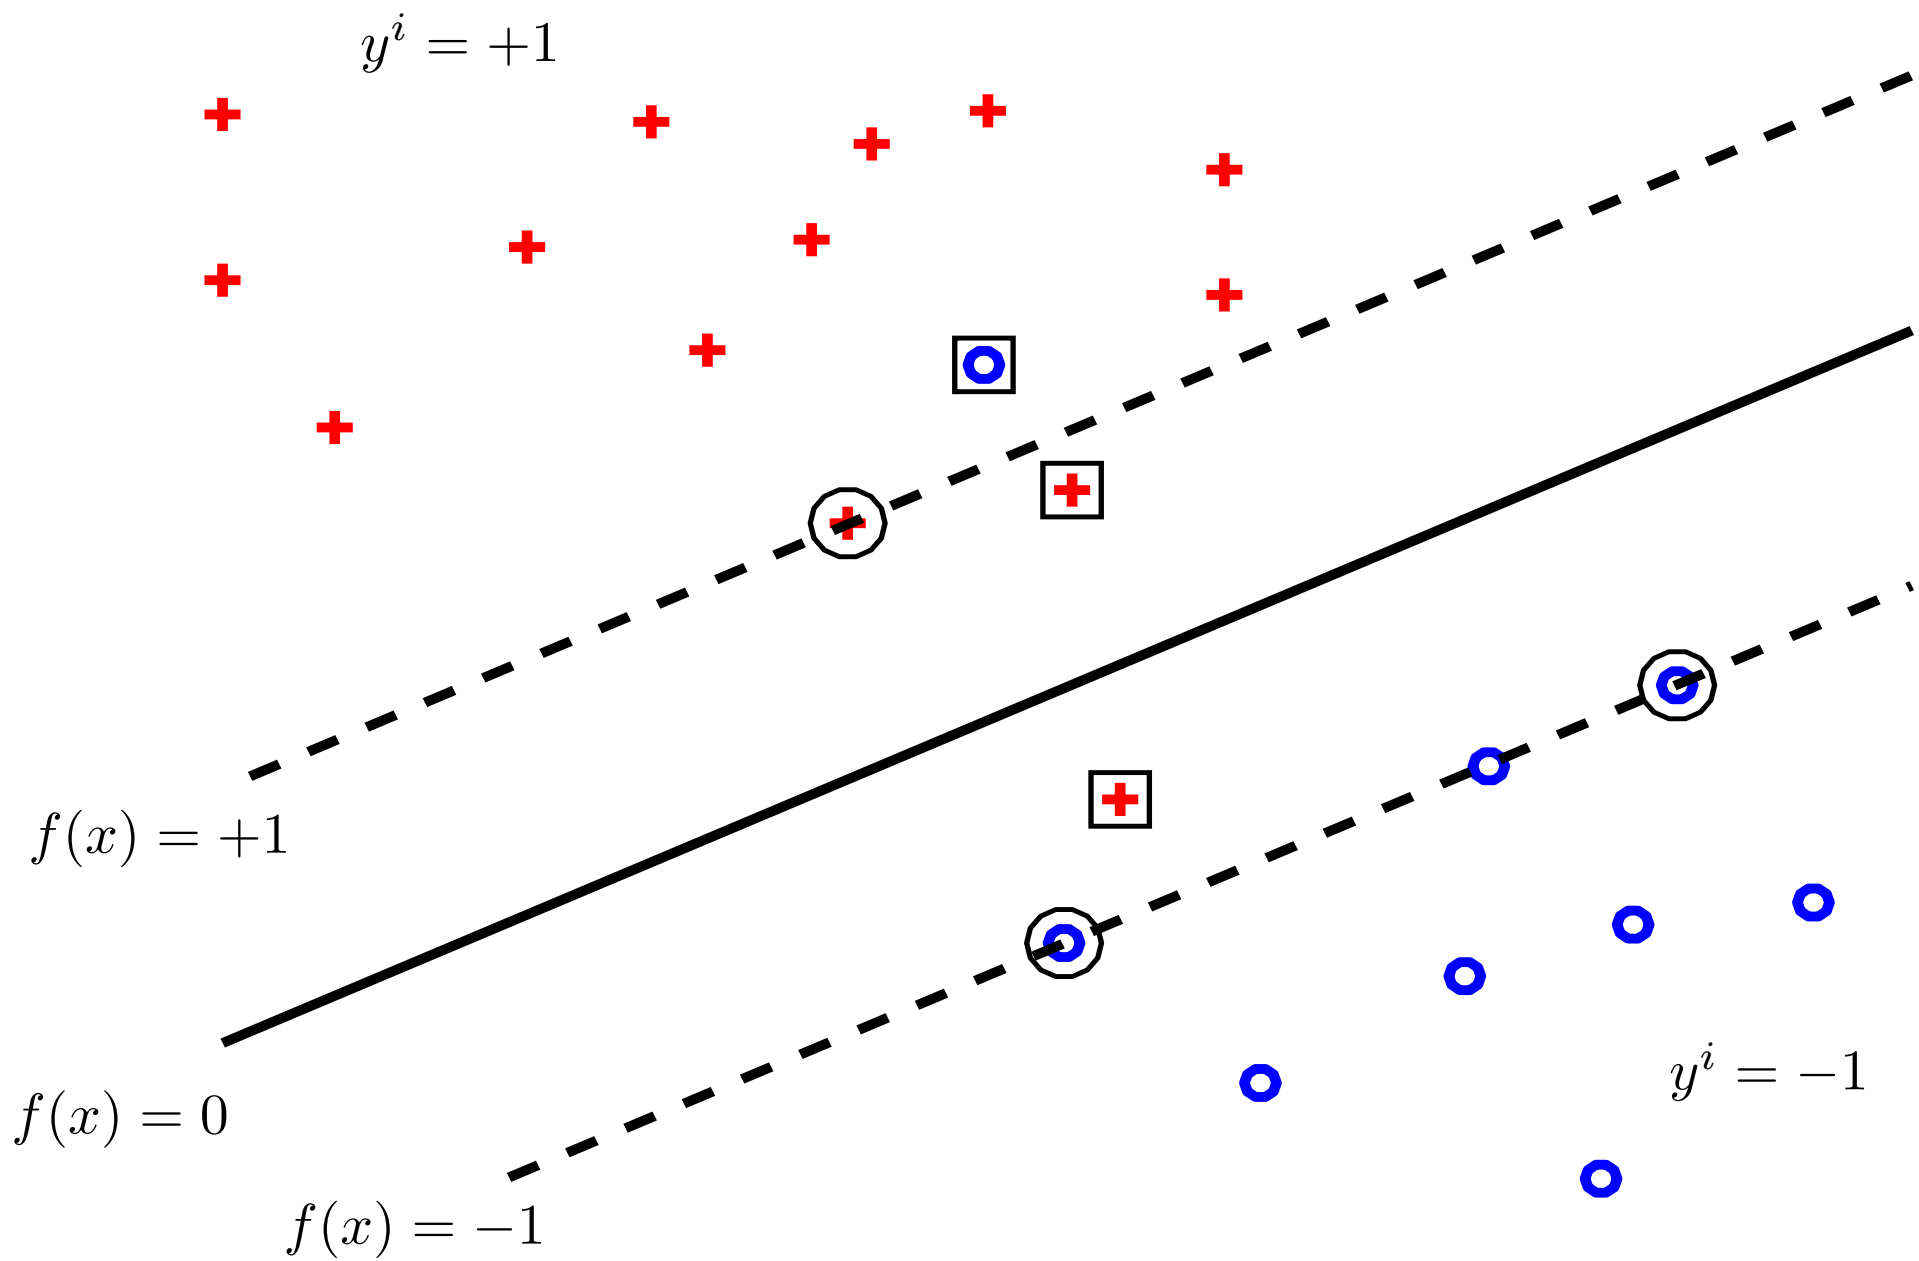


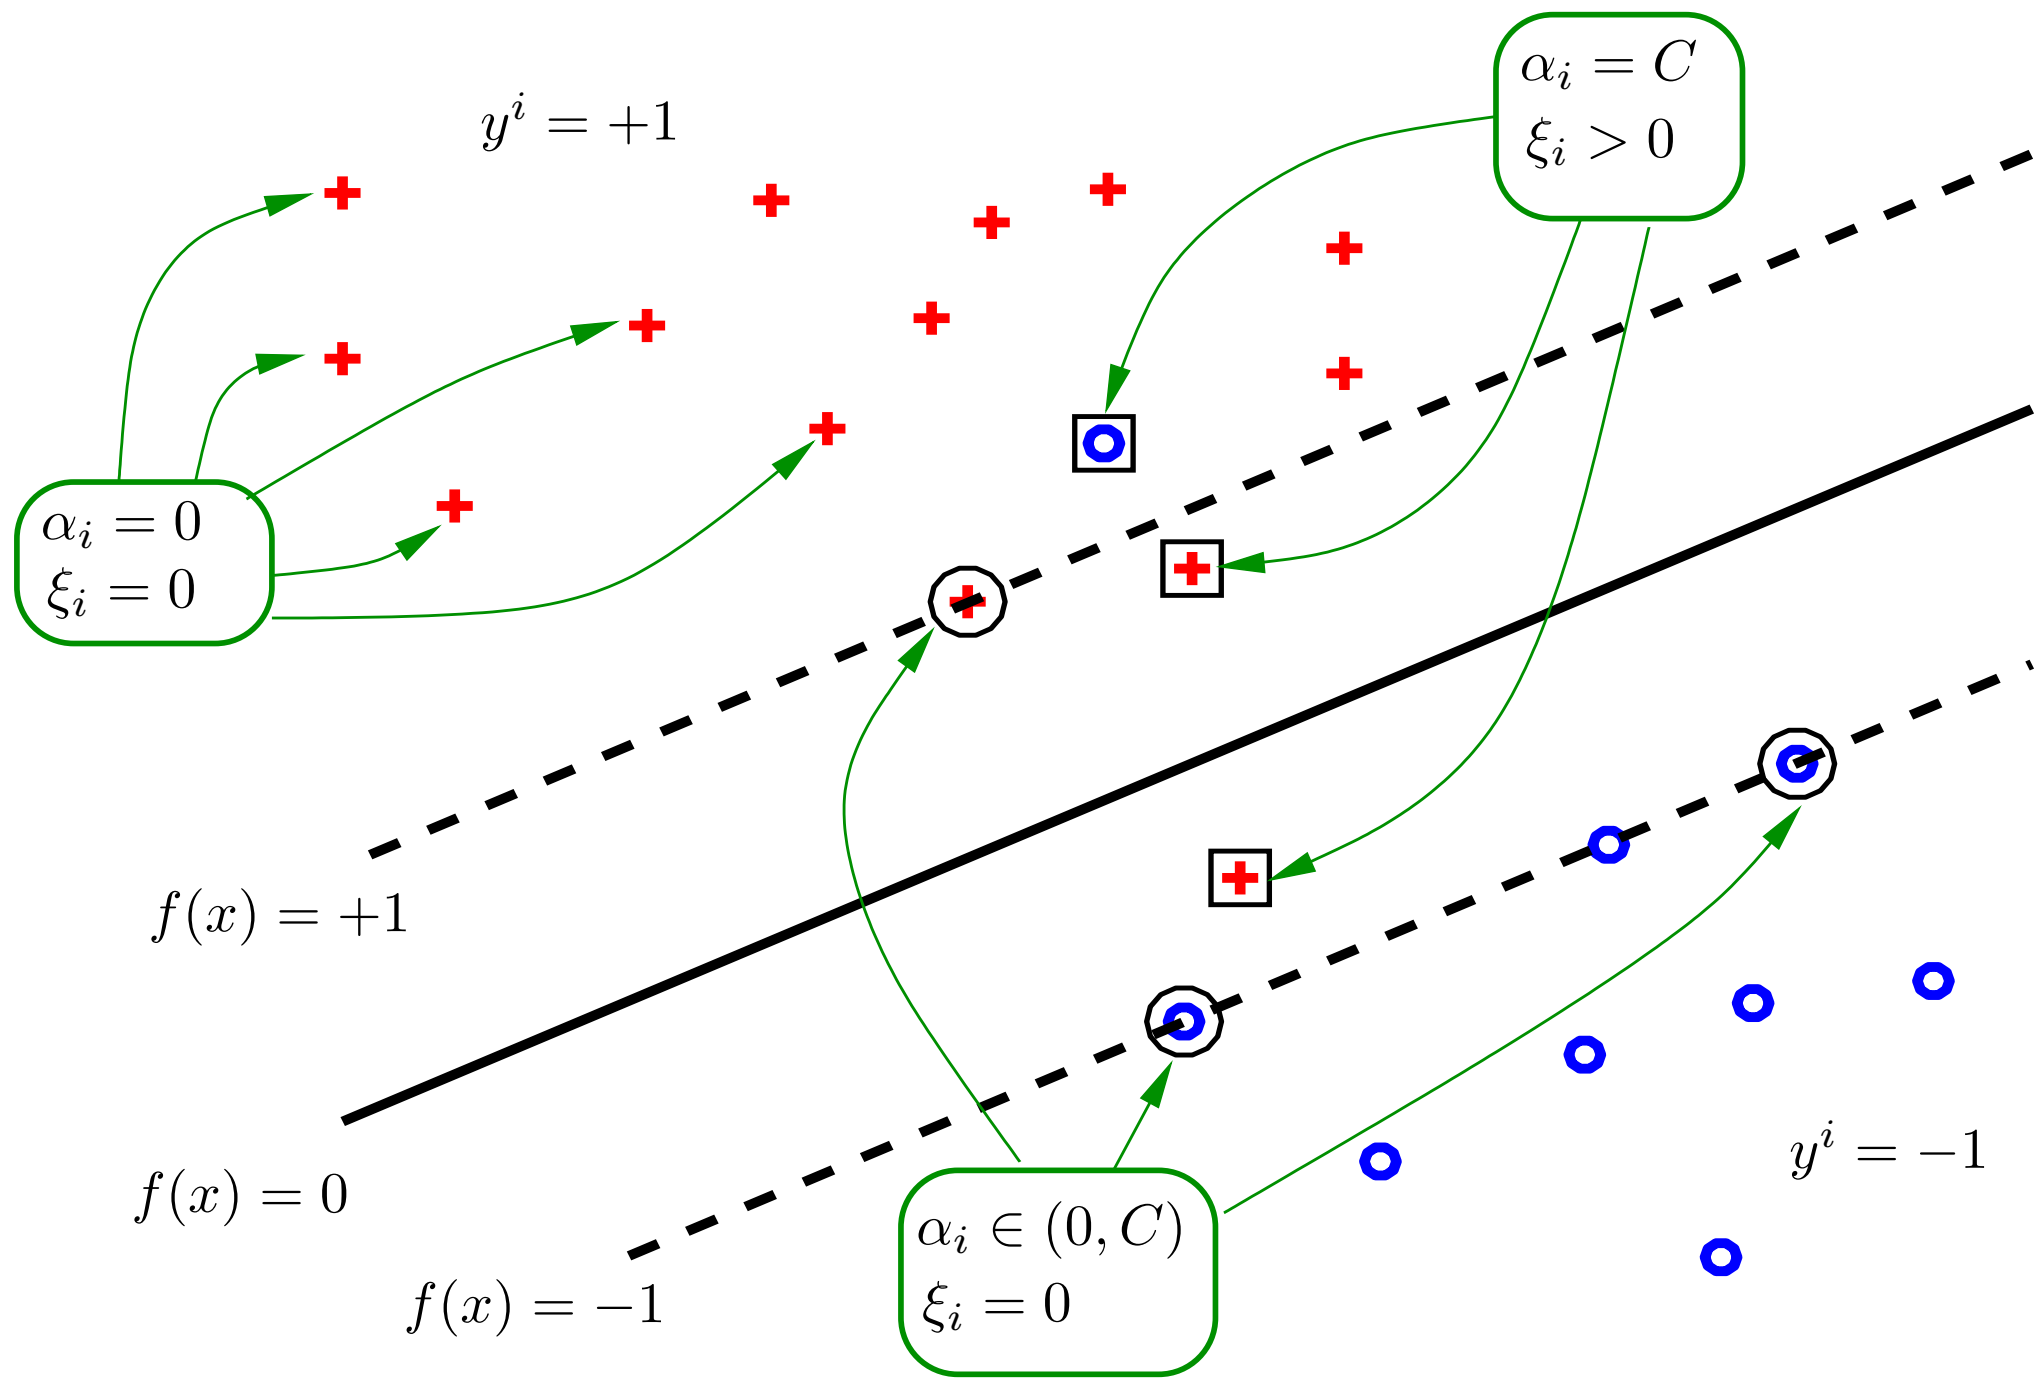
$y^i = +1$

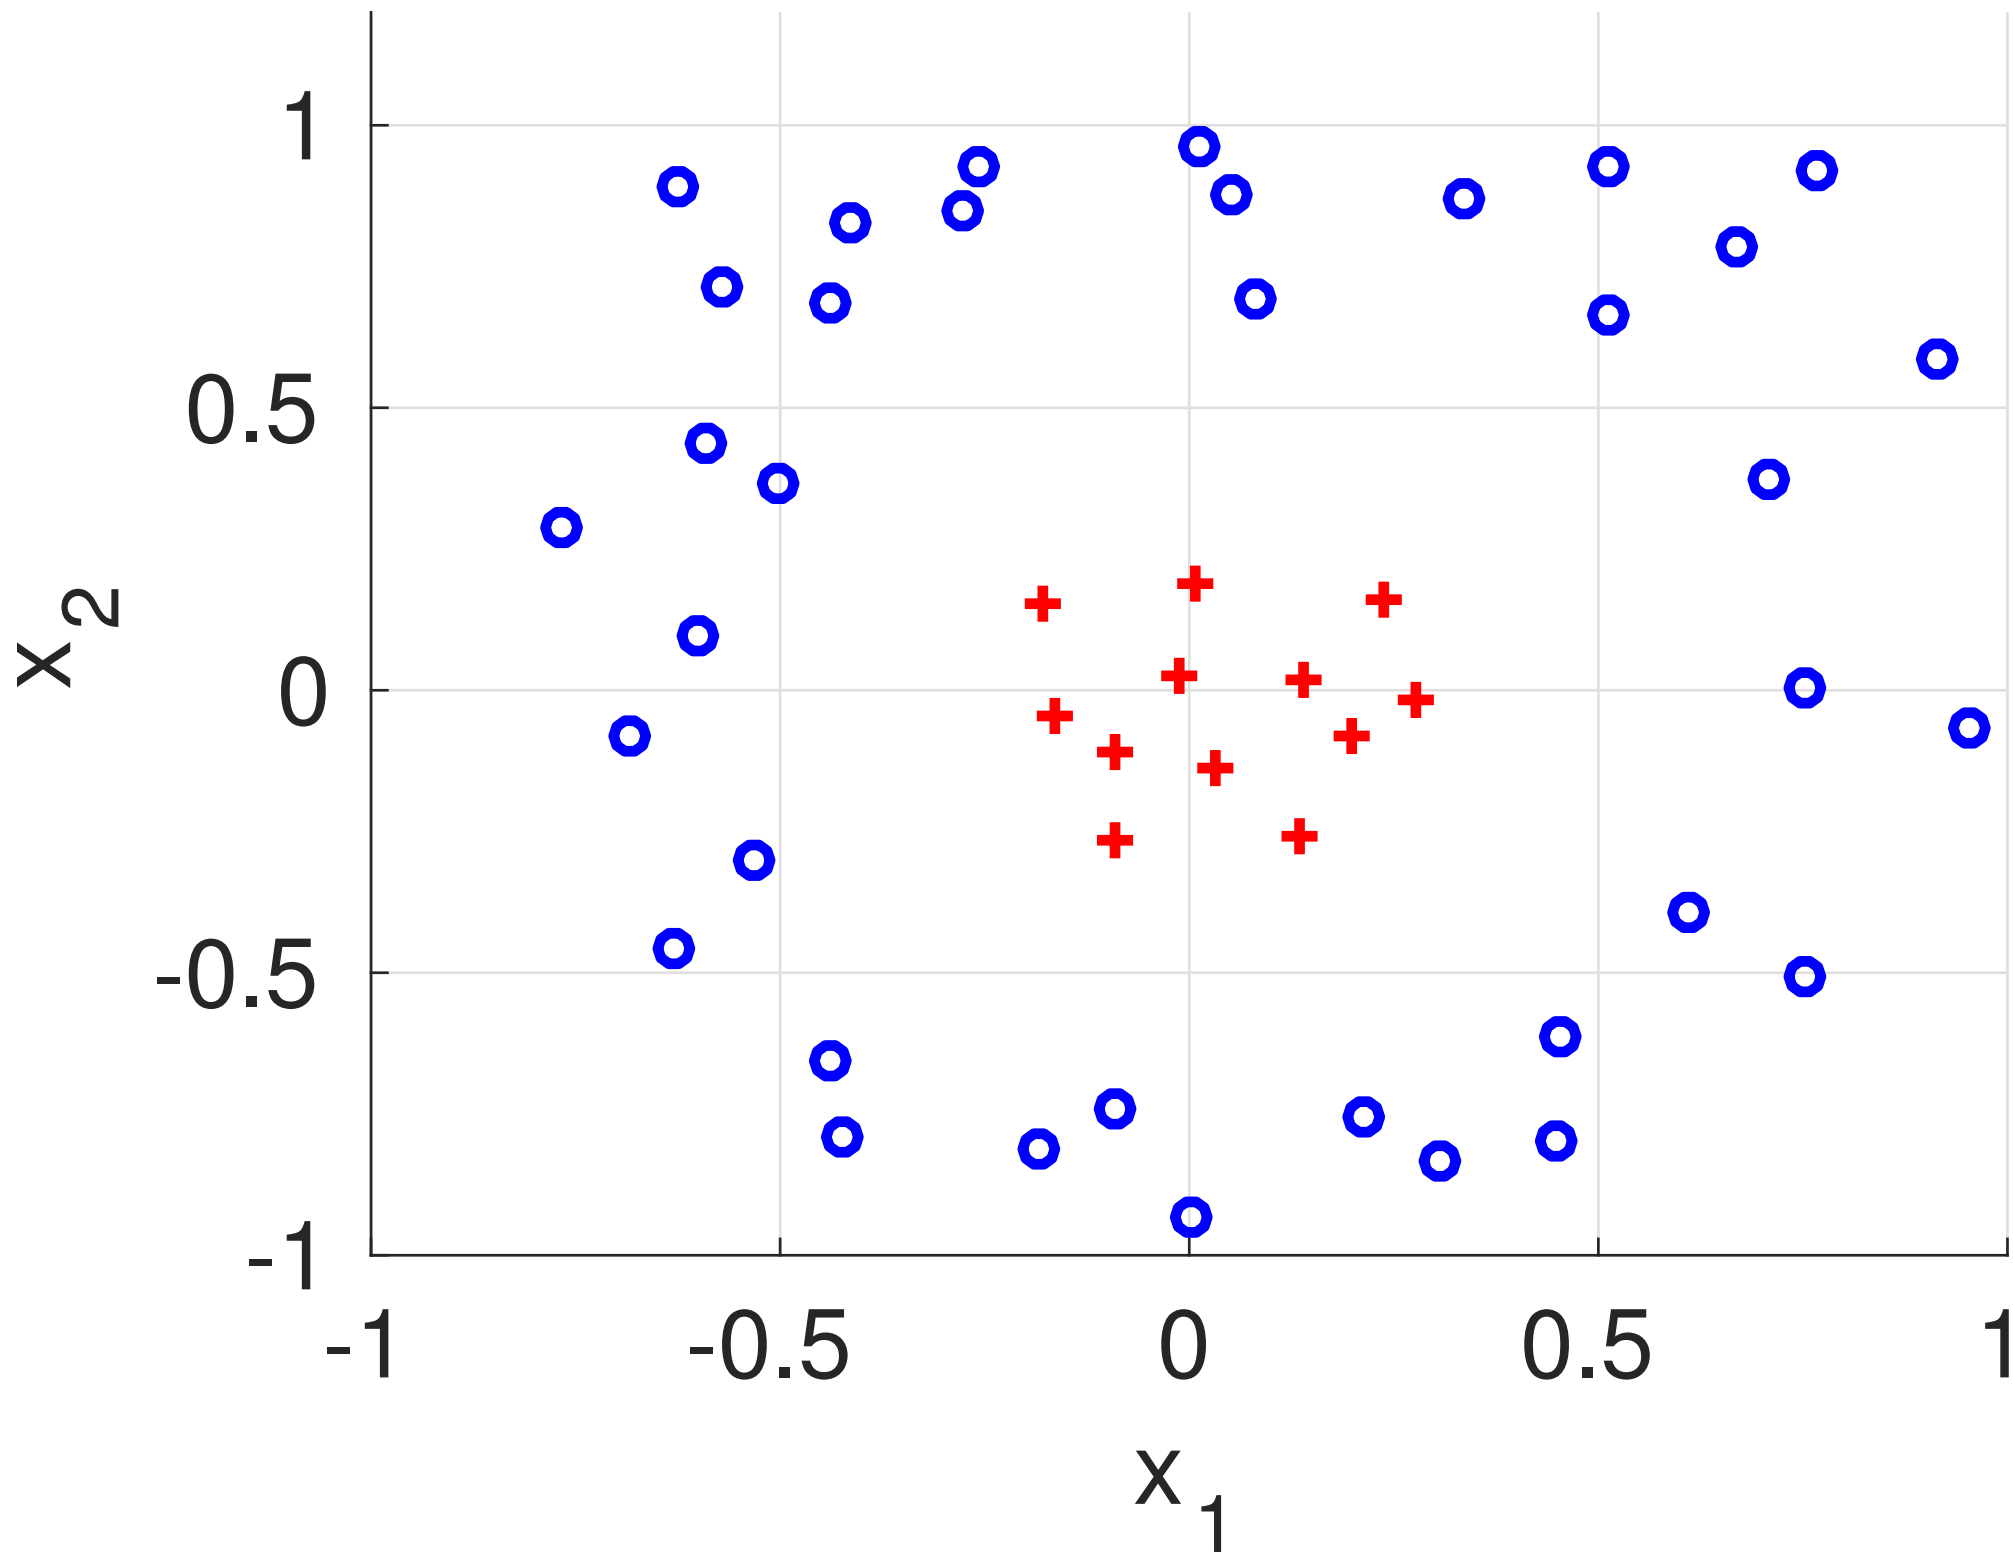


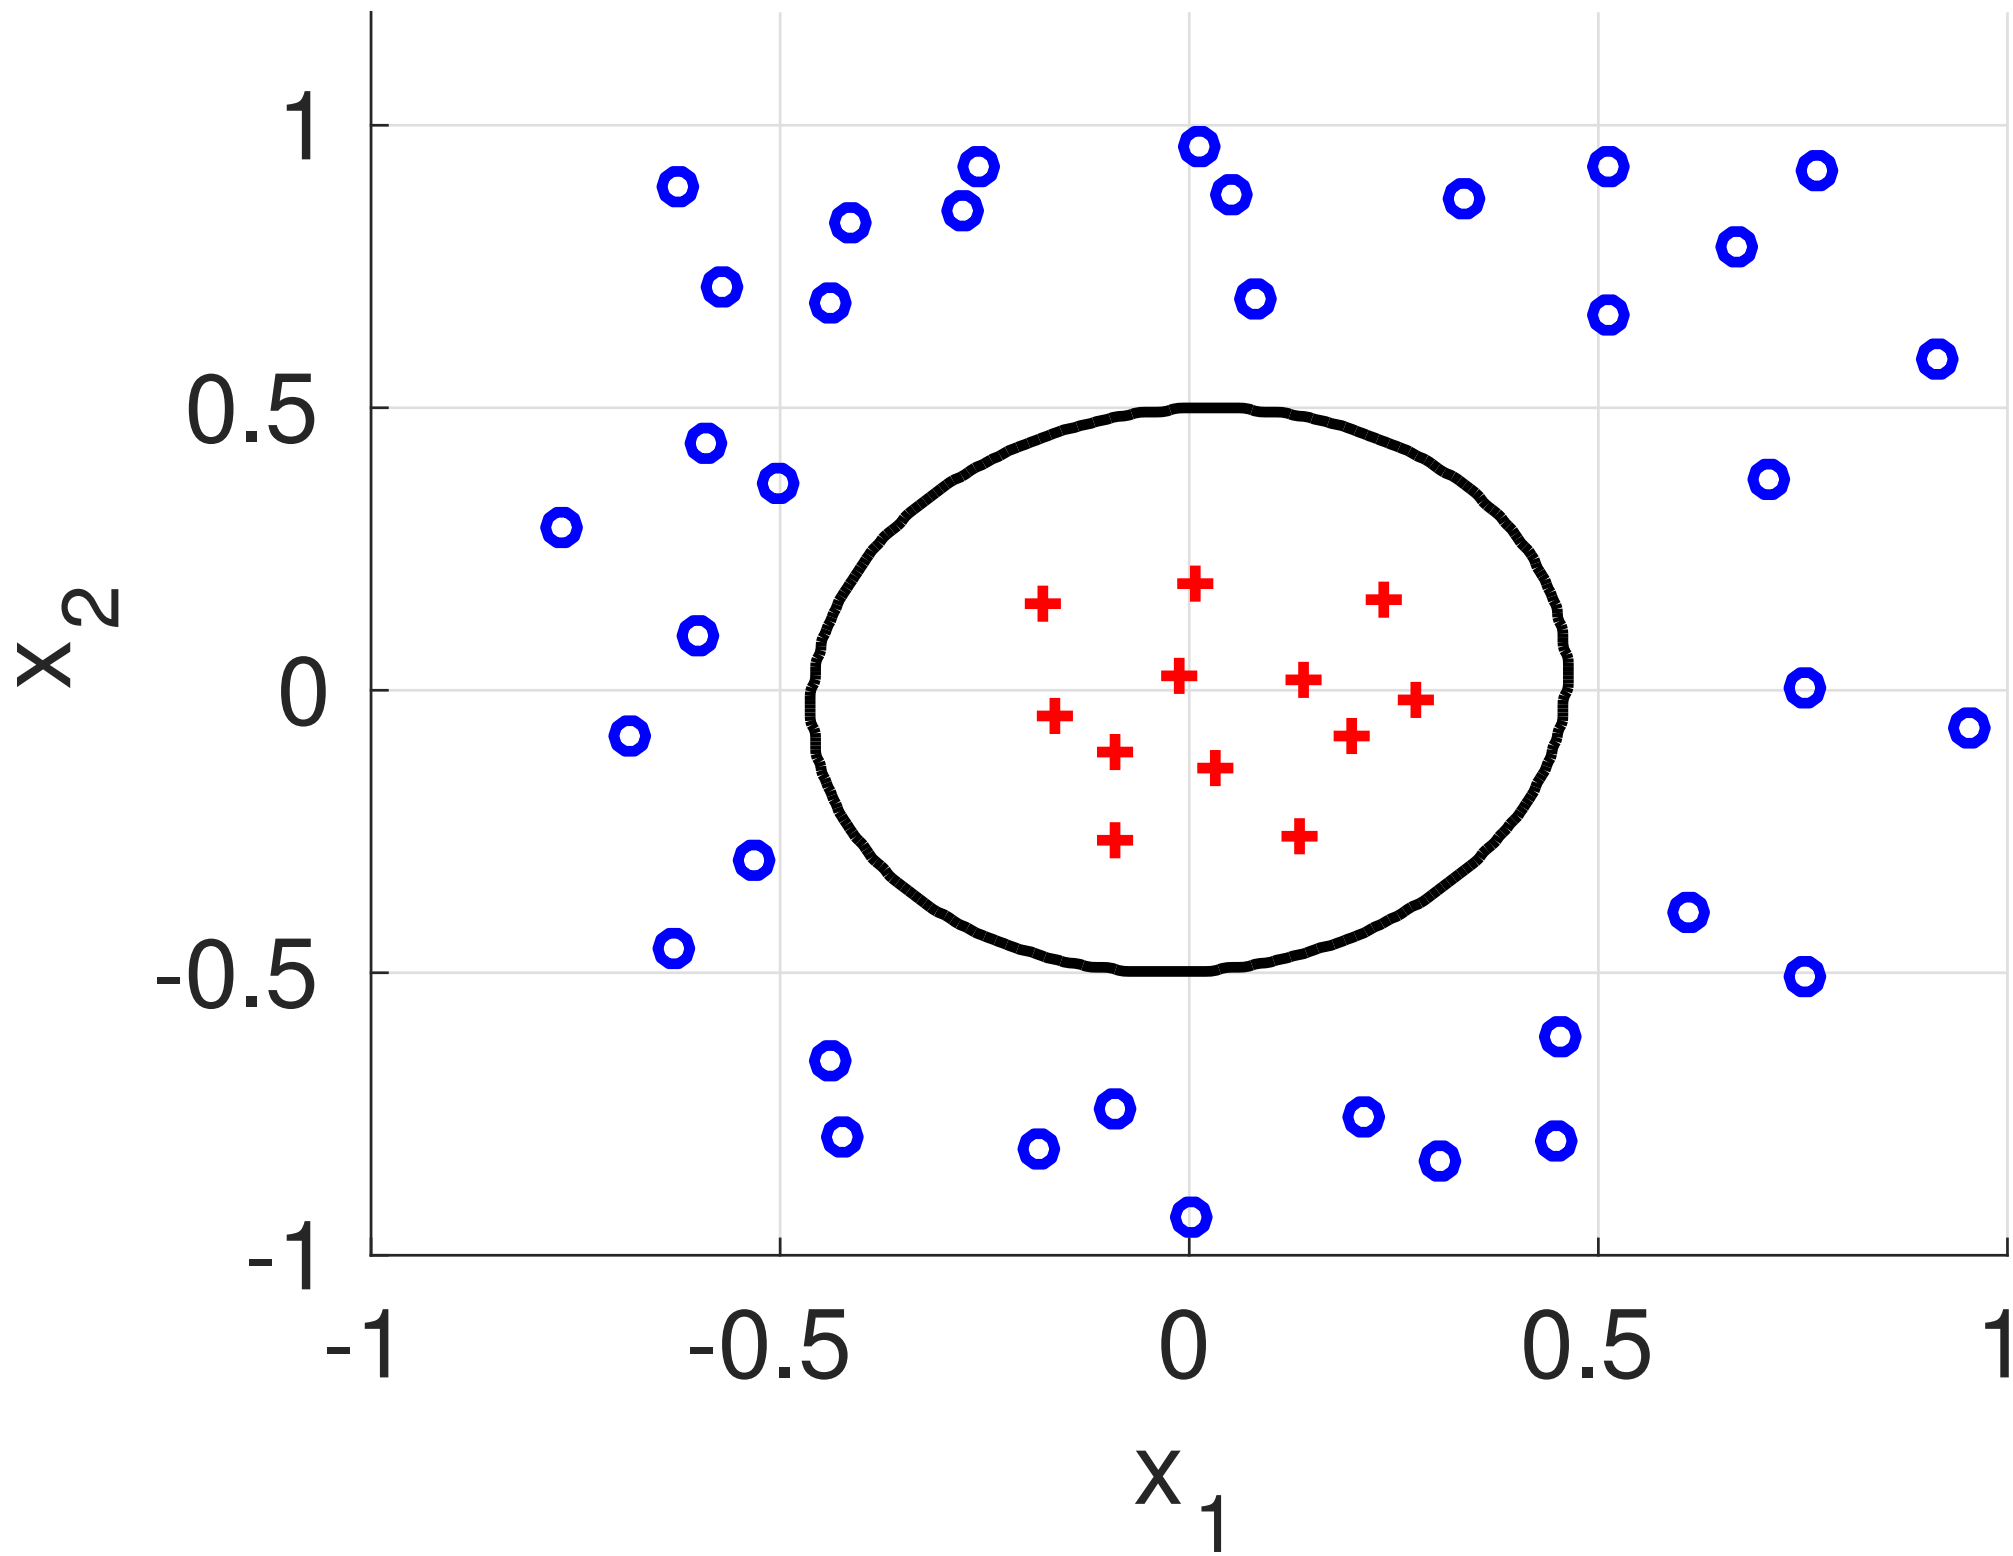
$y^i = -1$

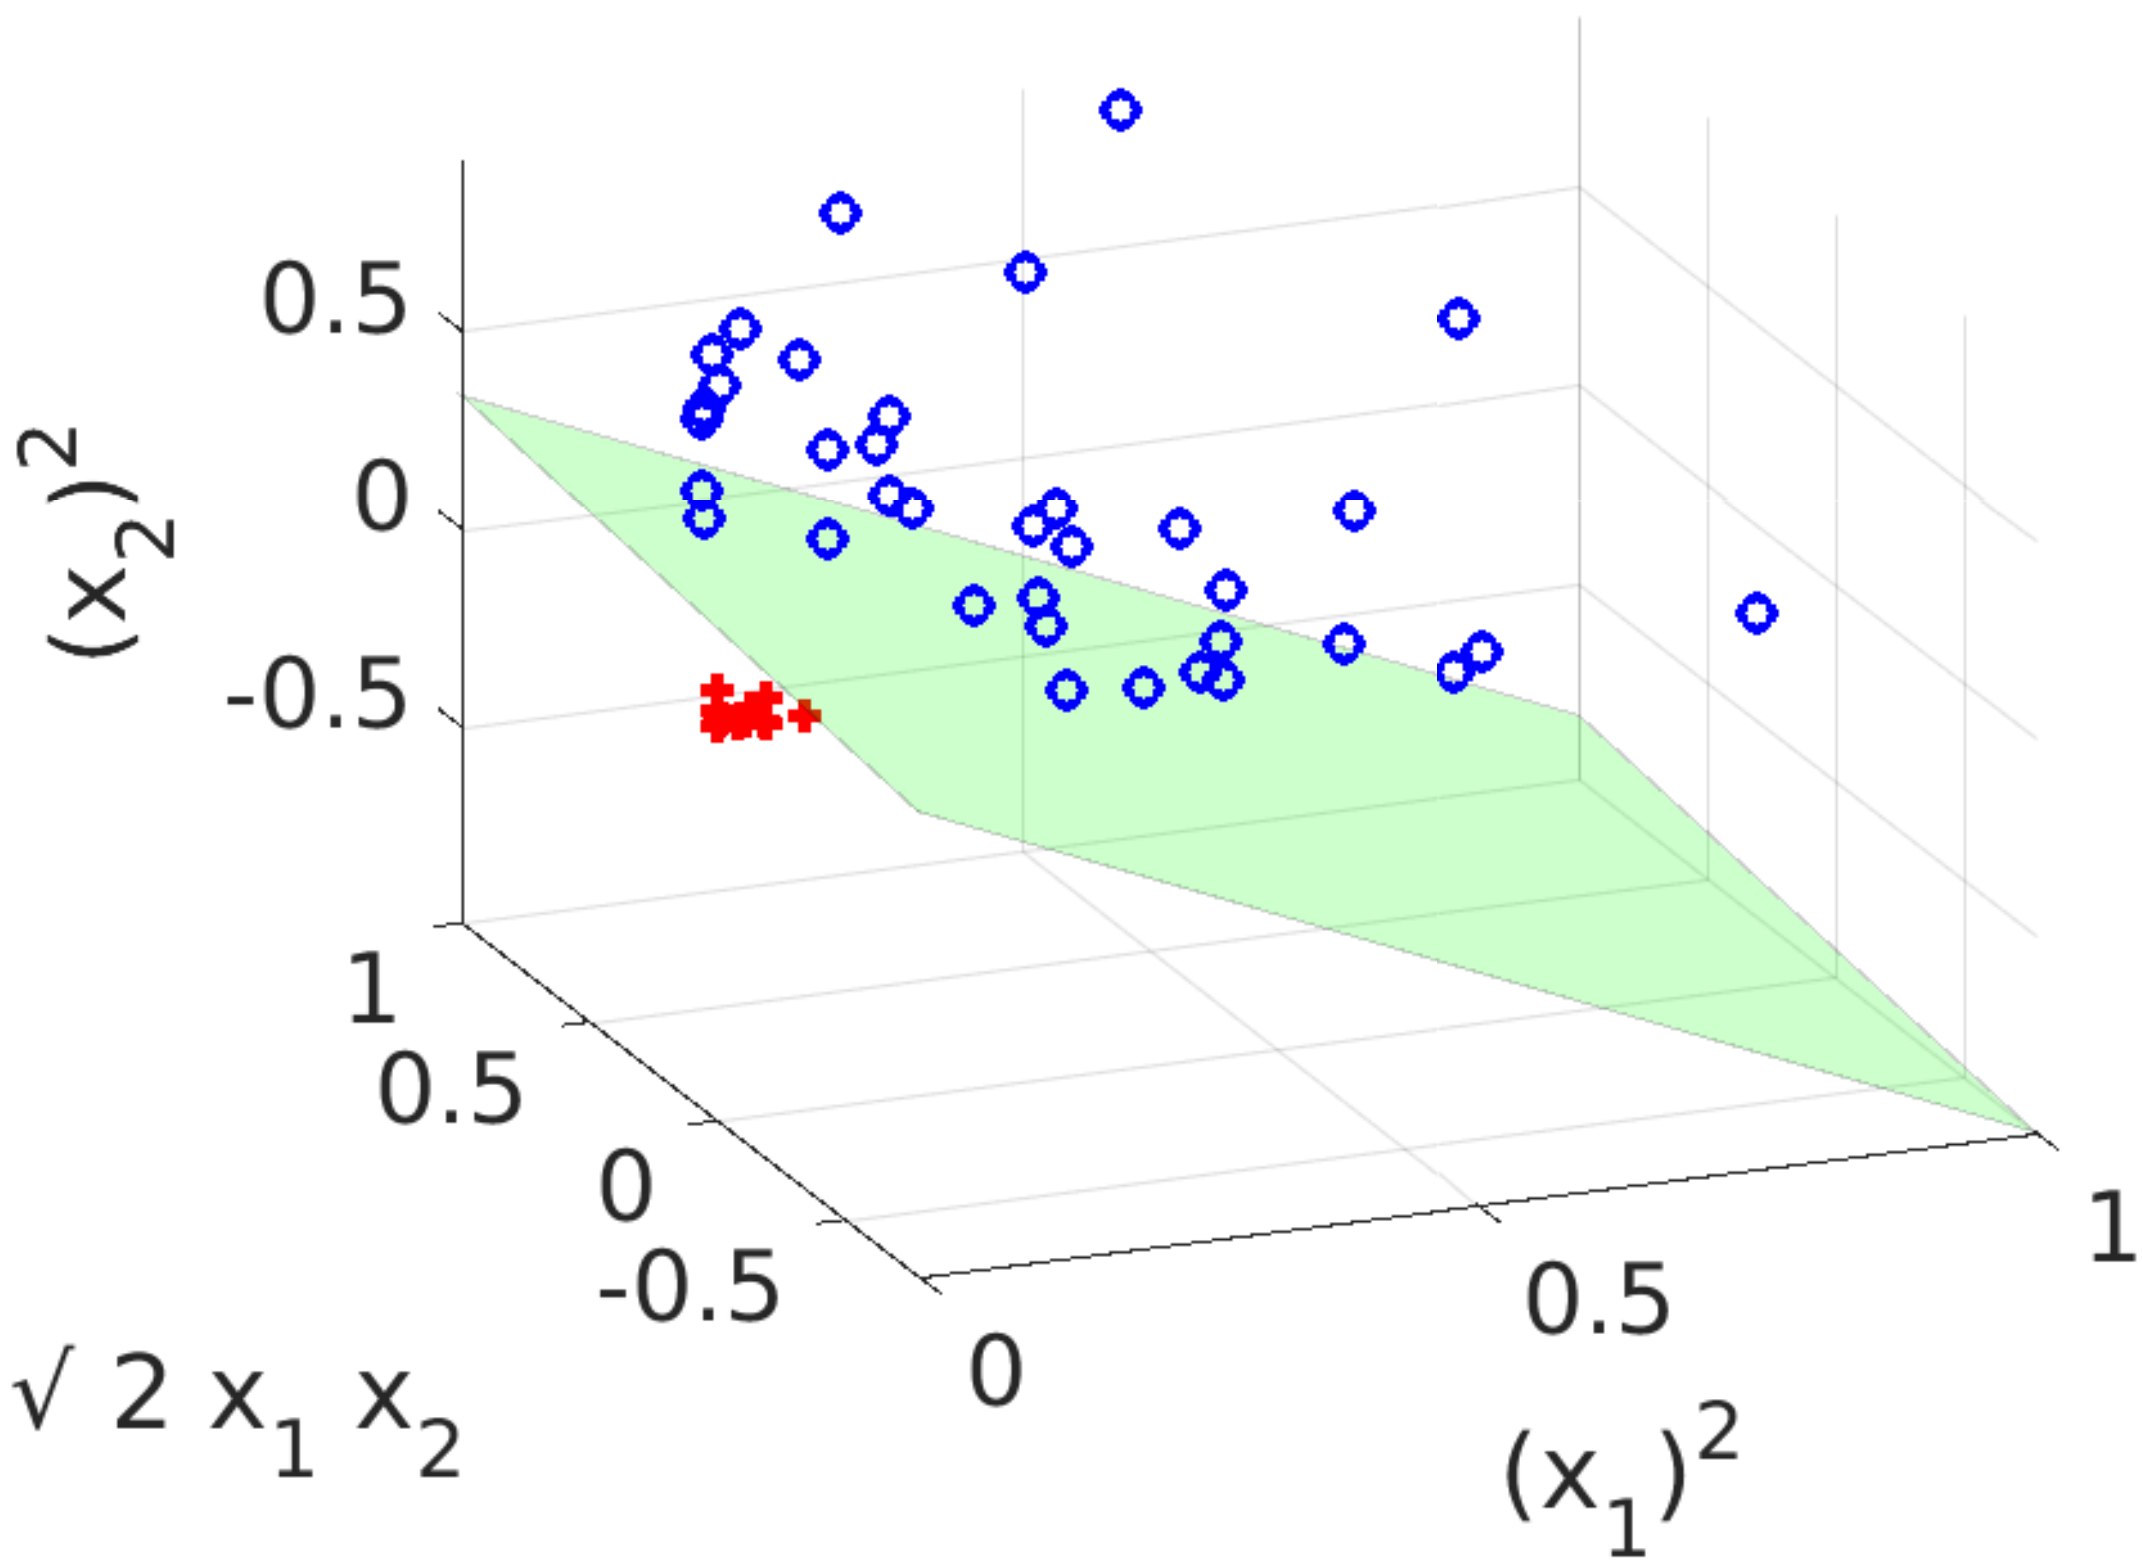


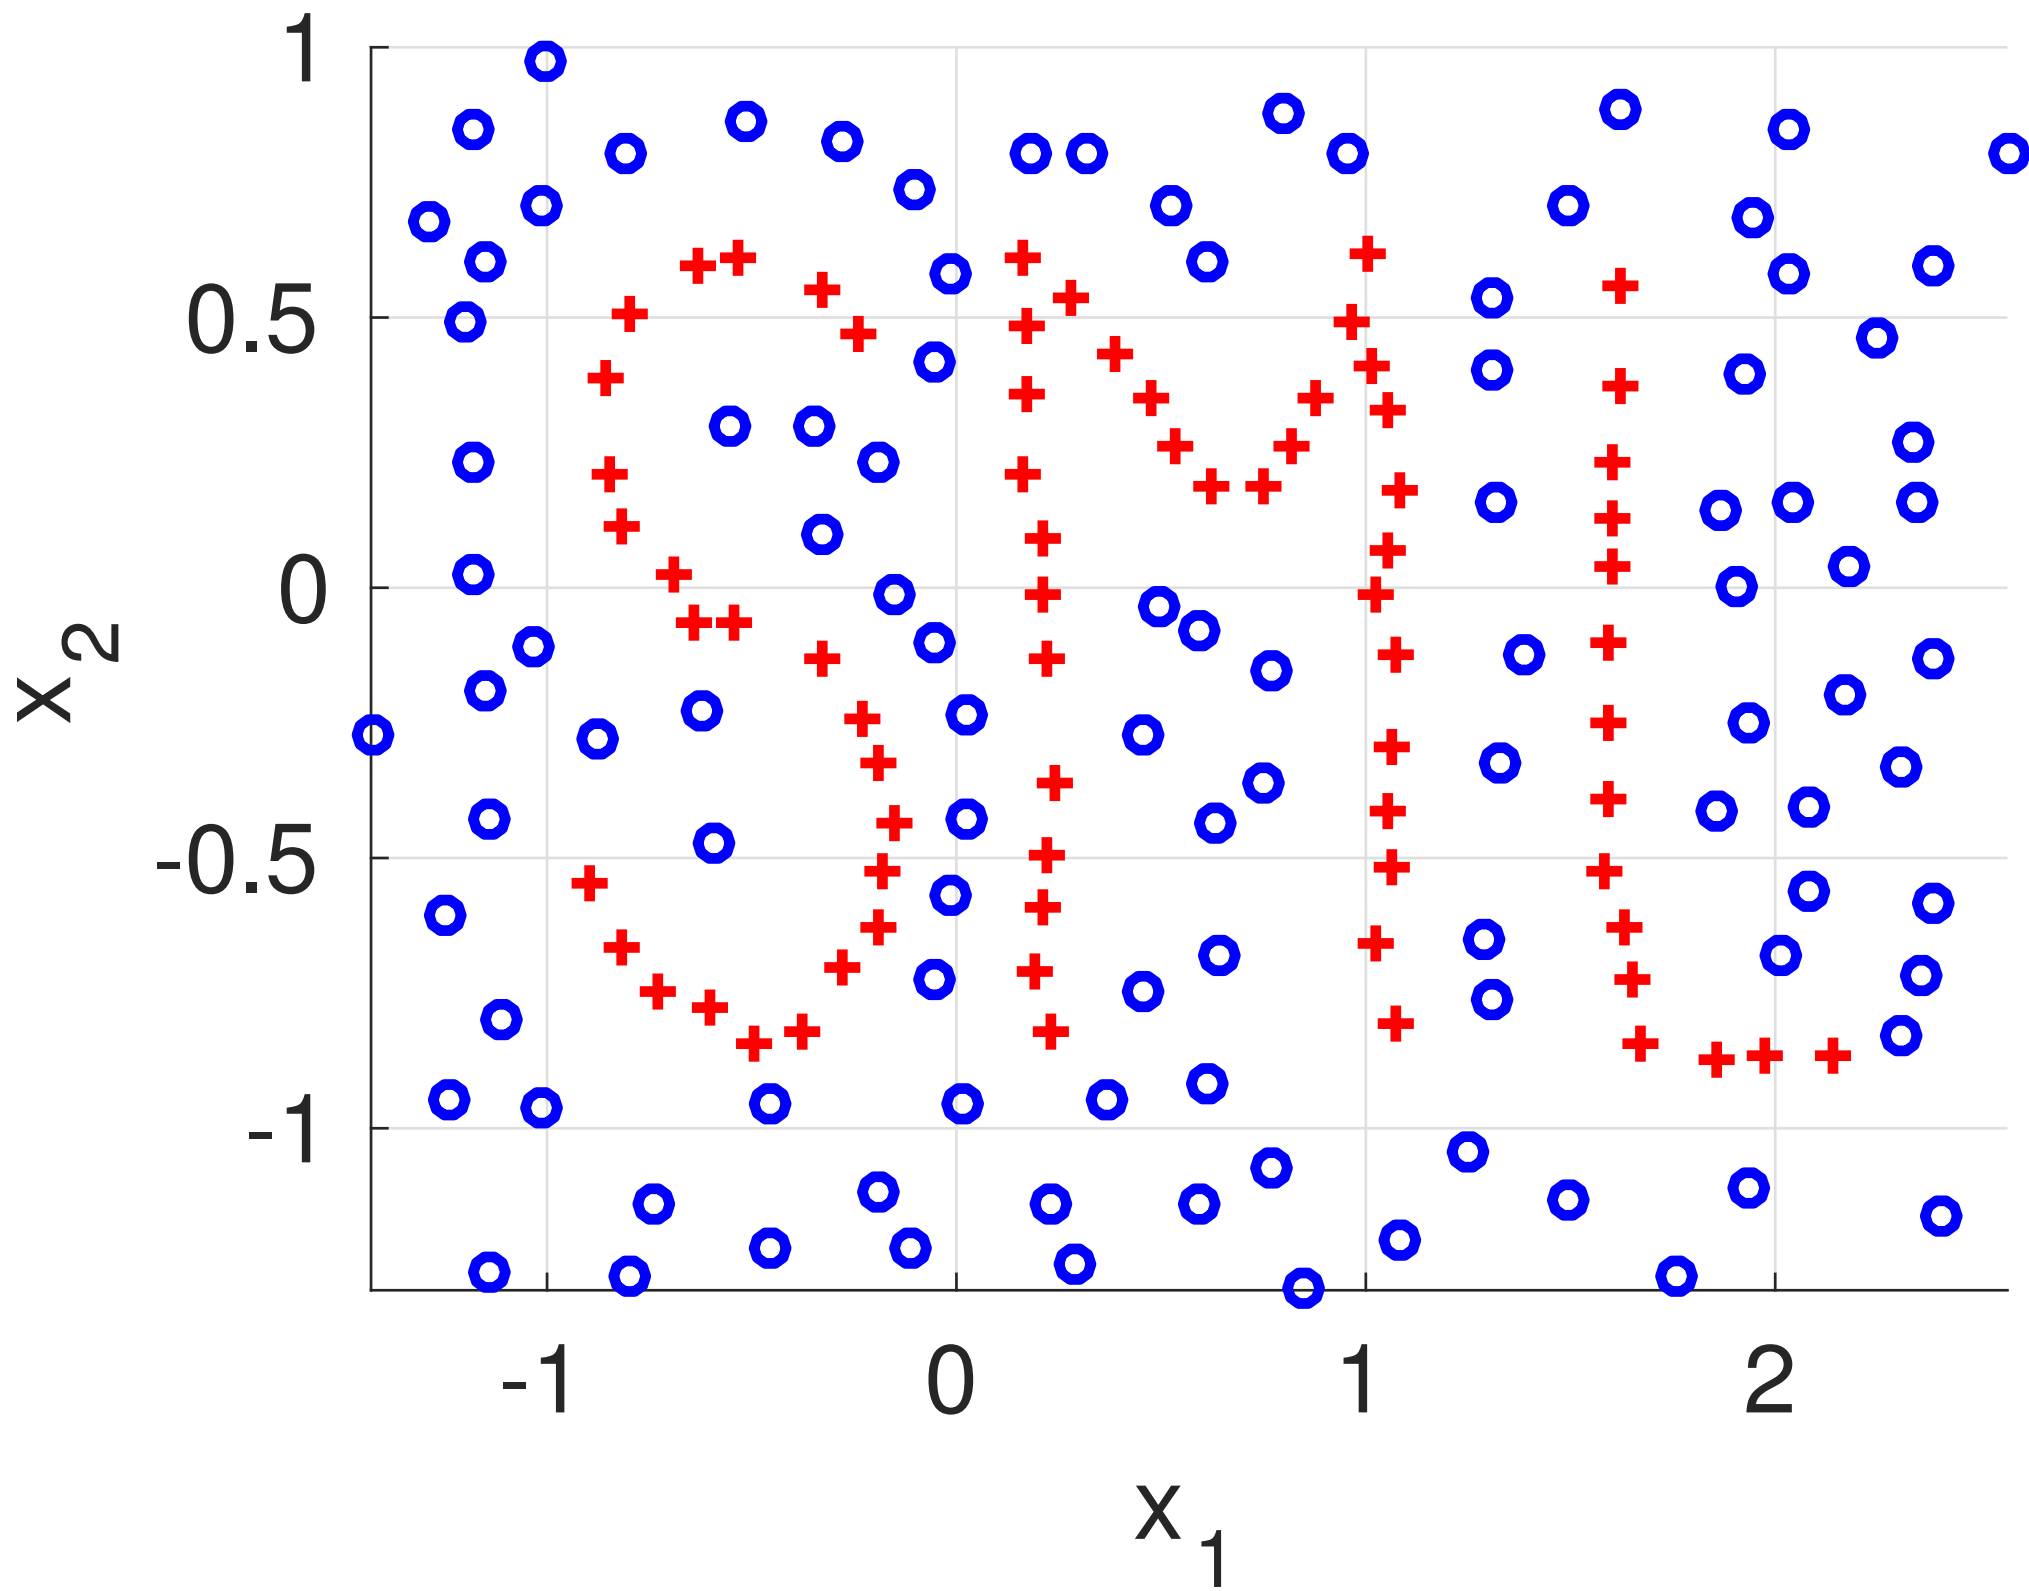


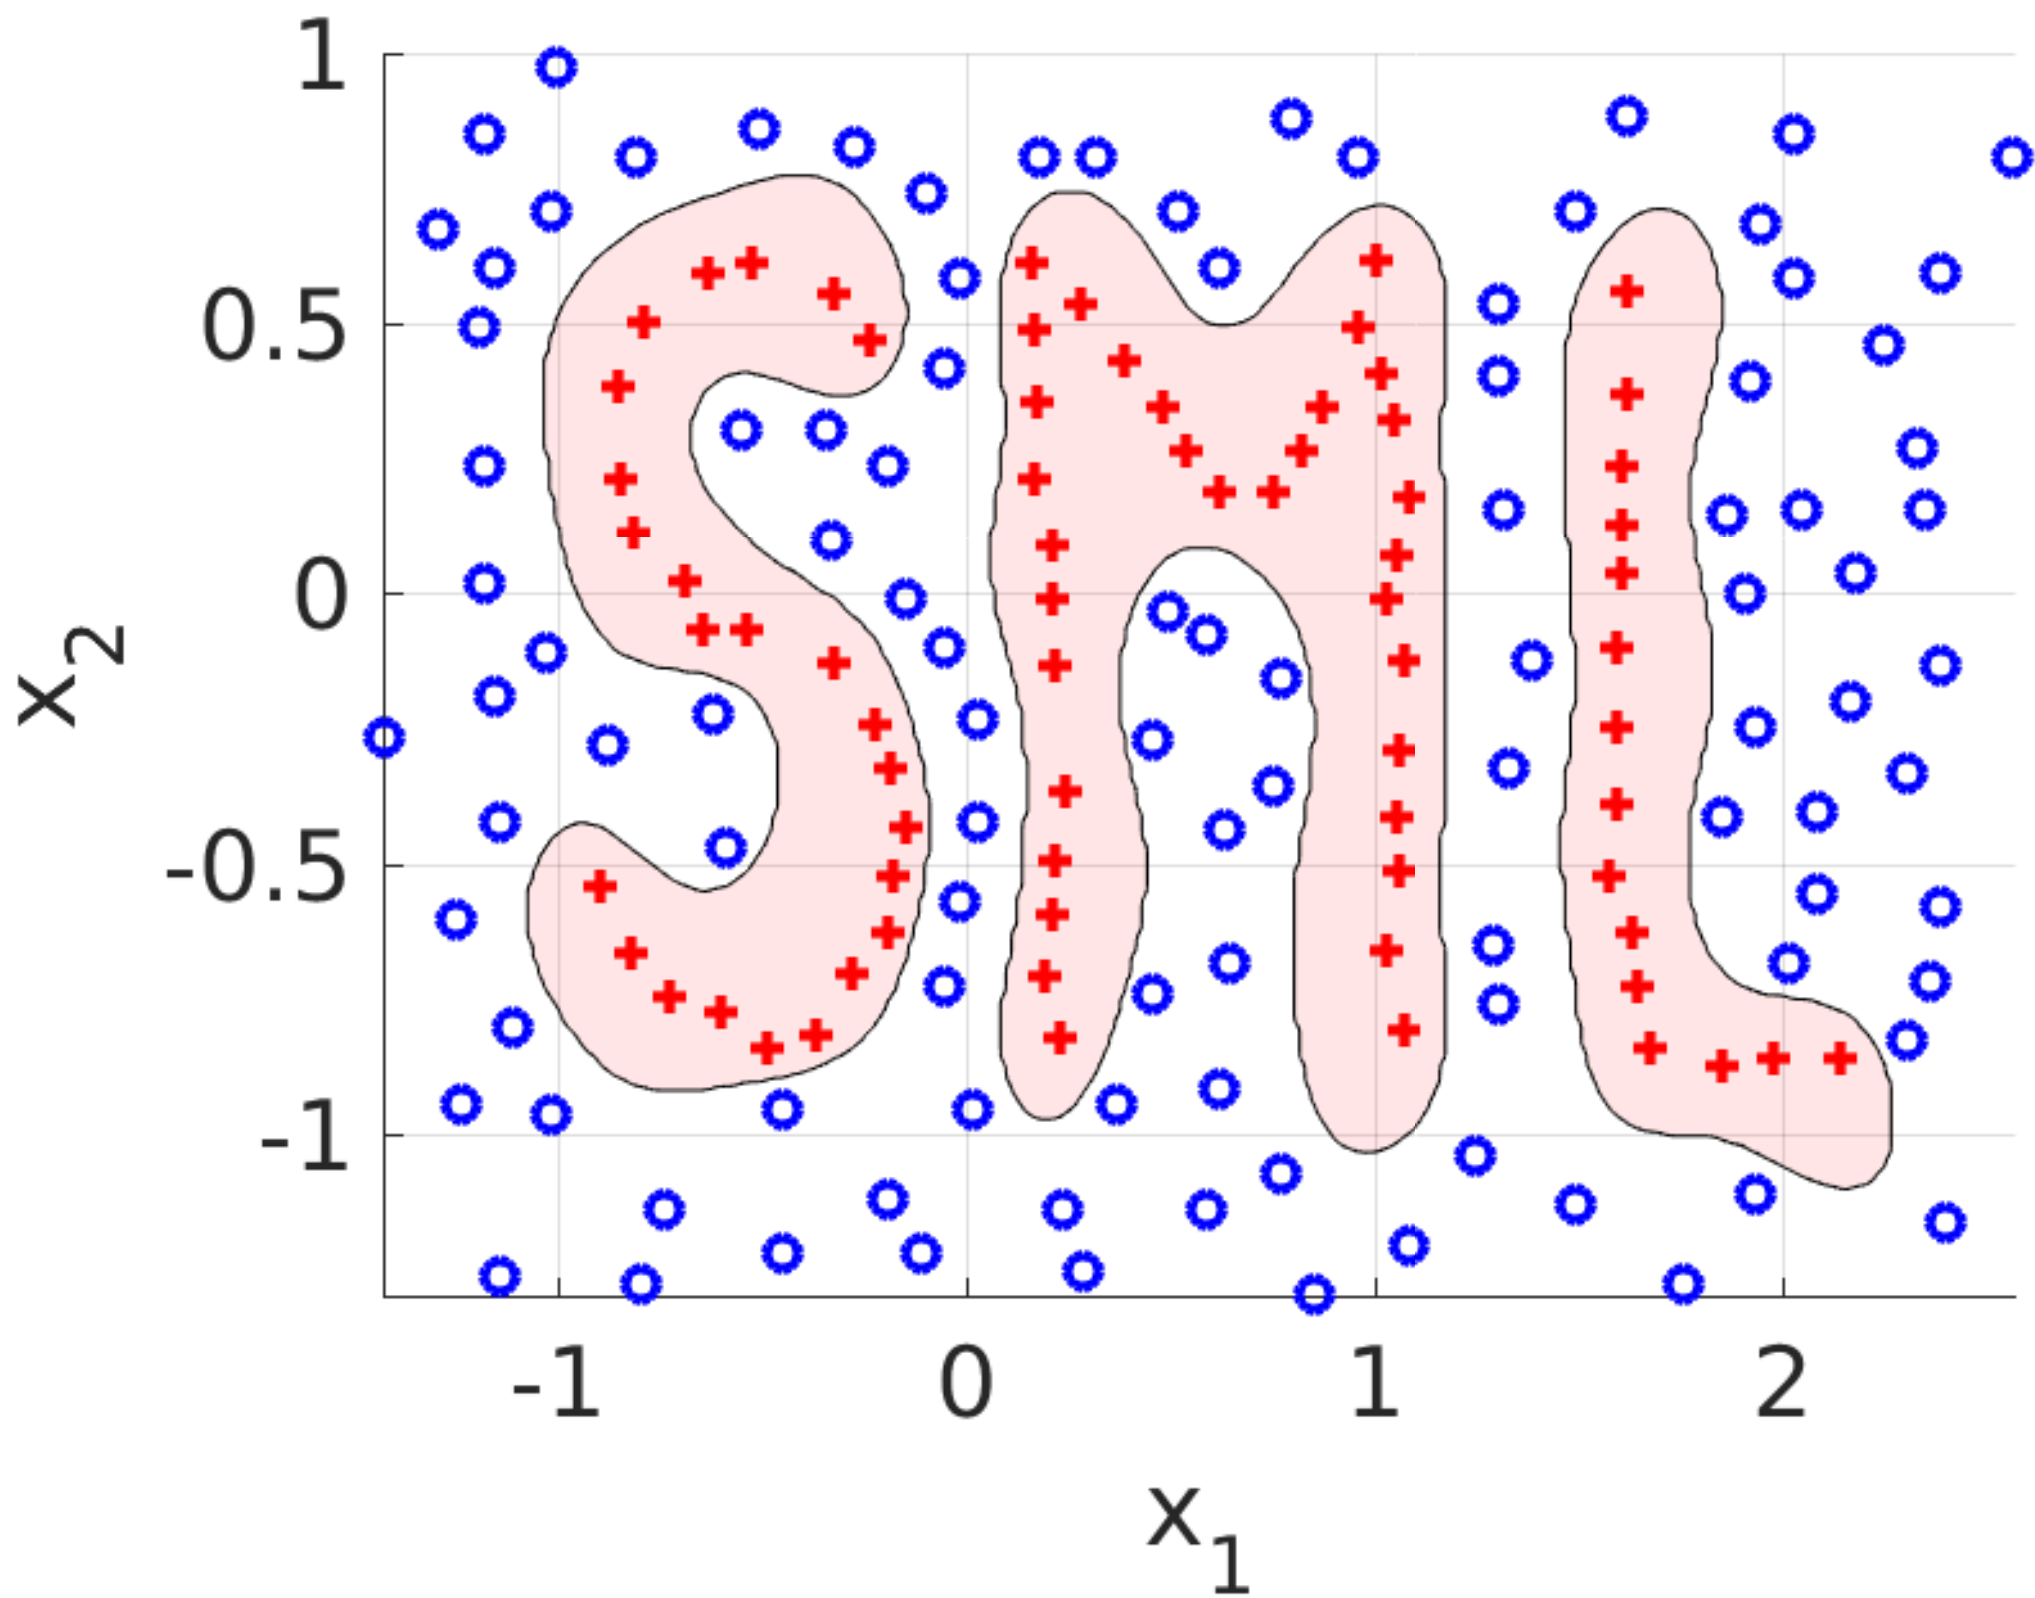


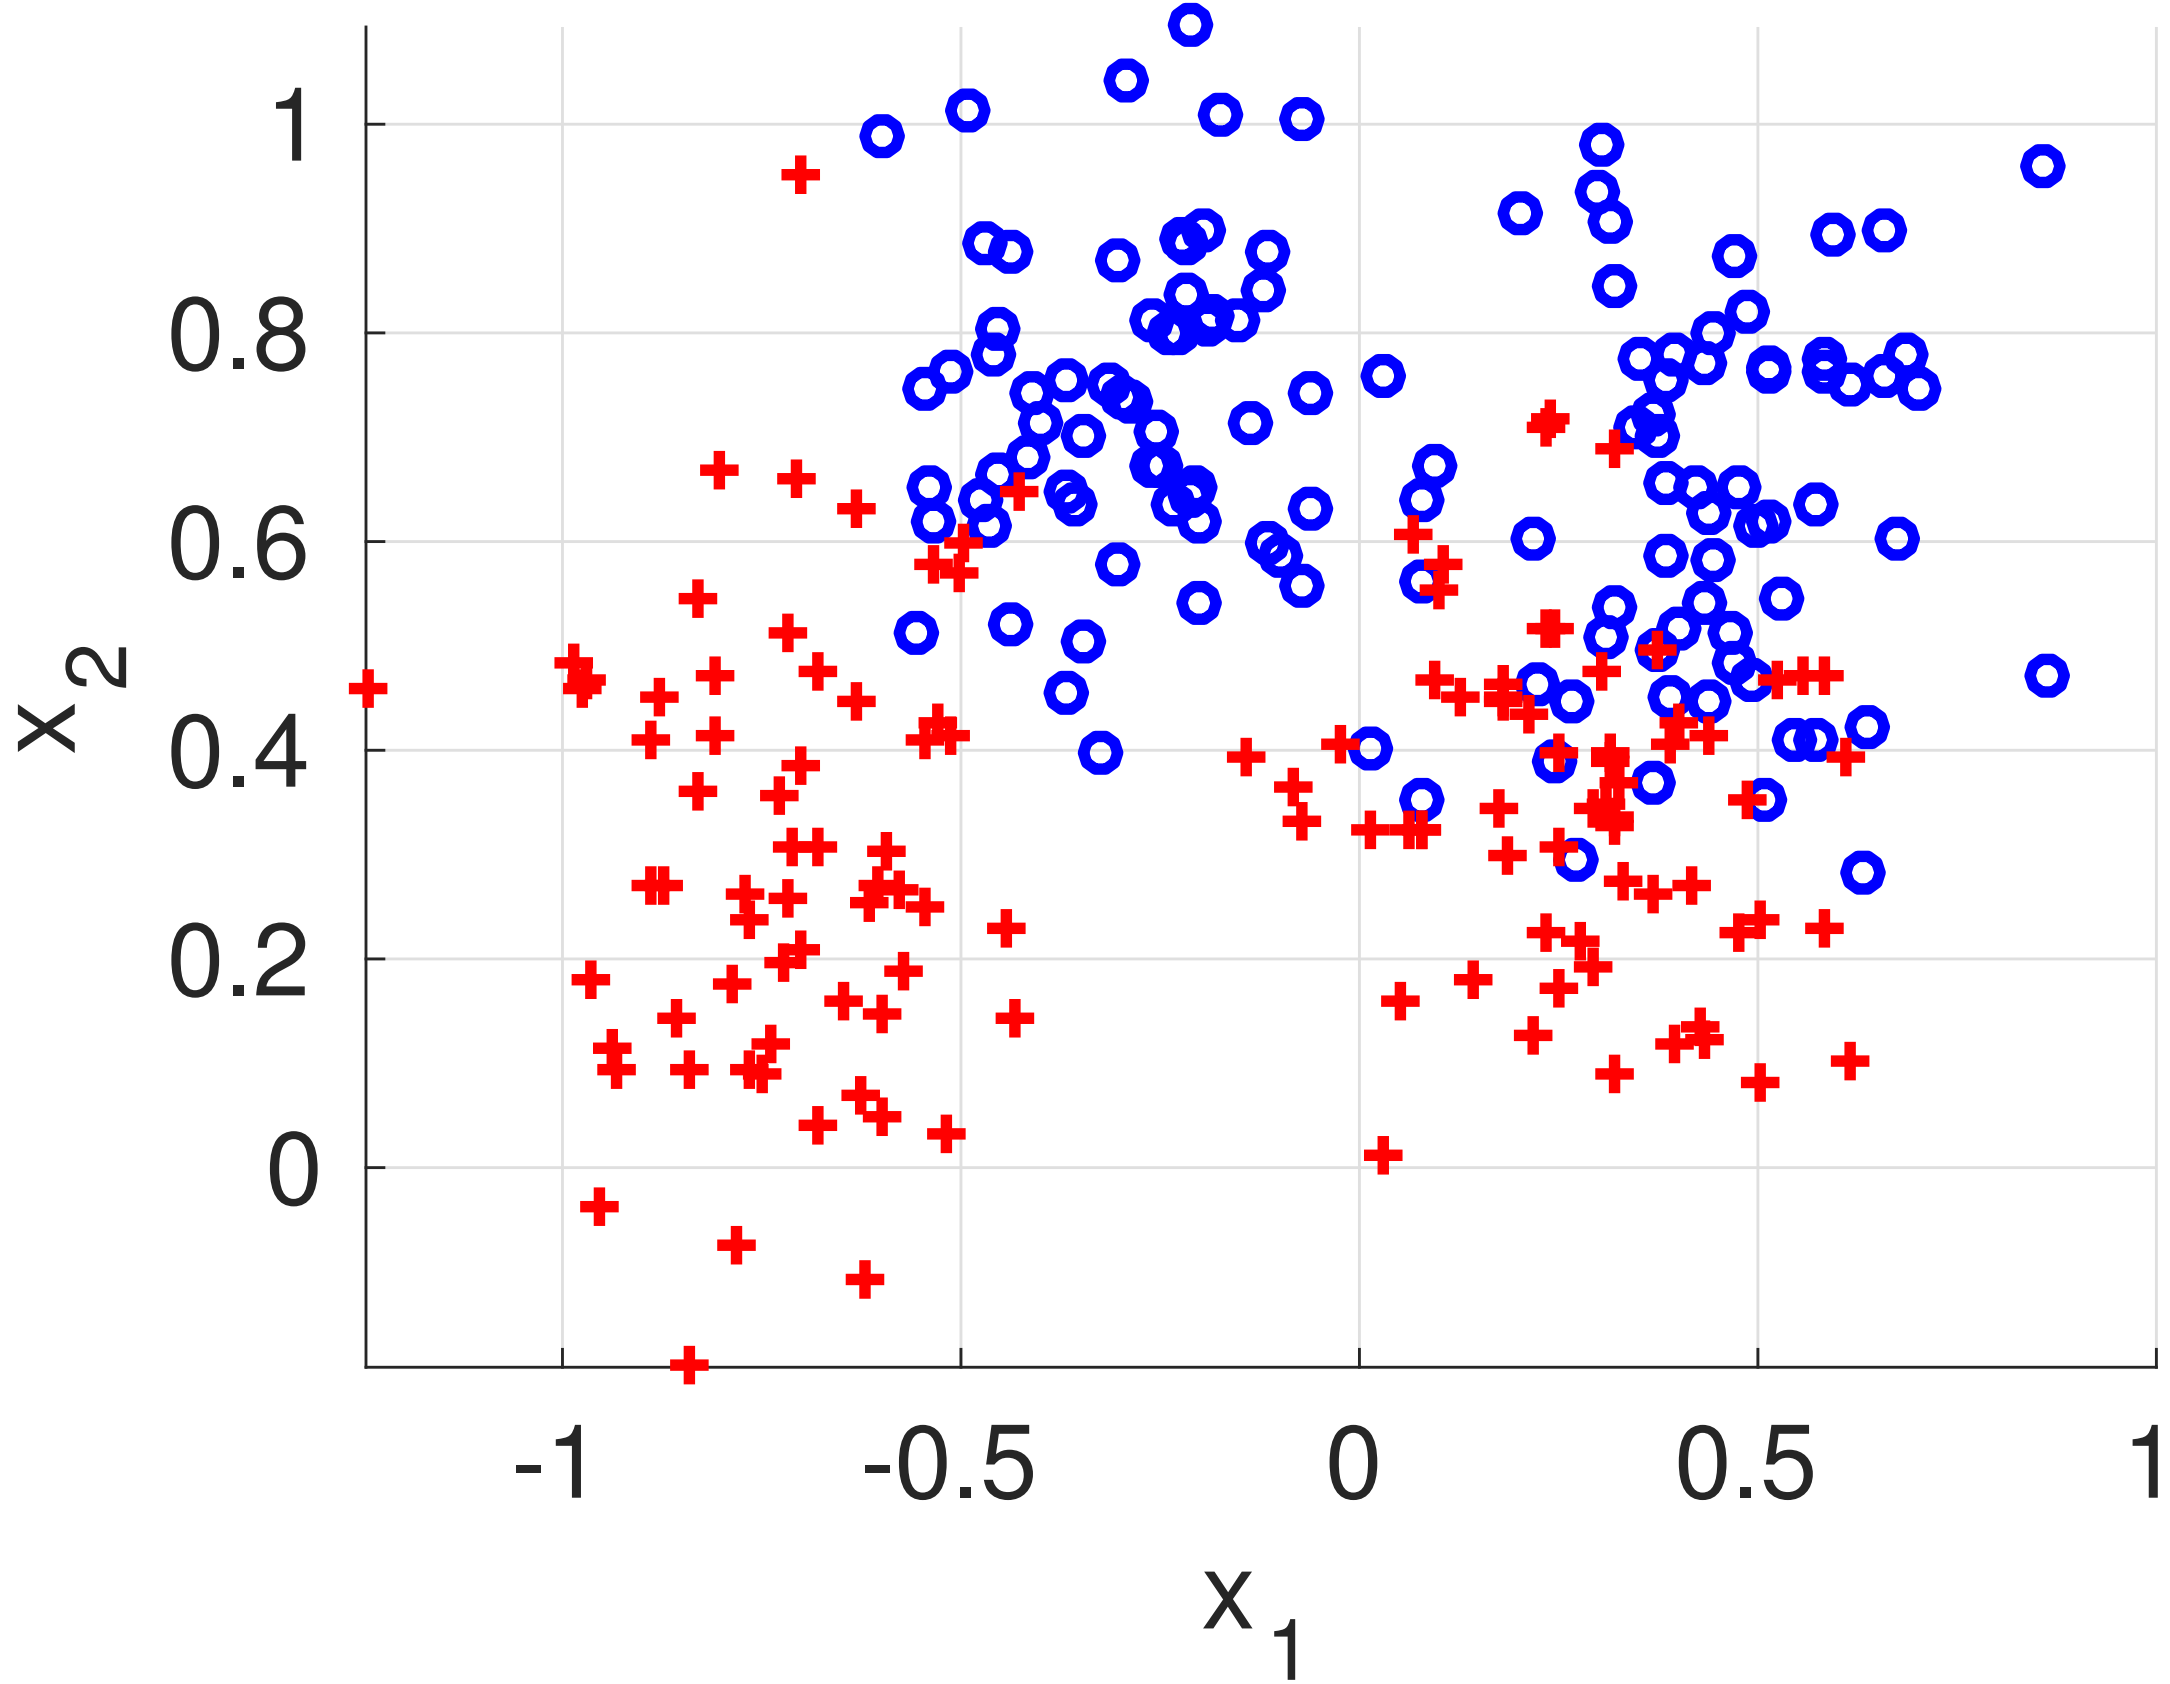


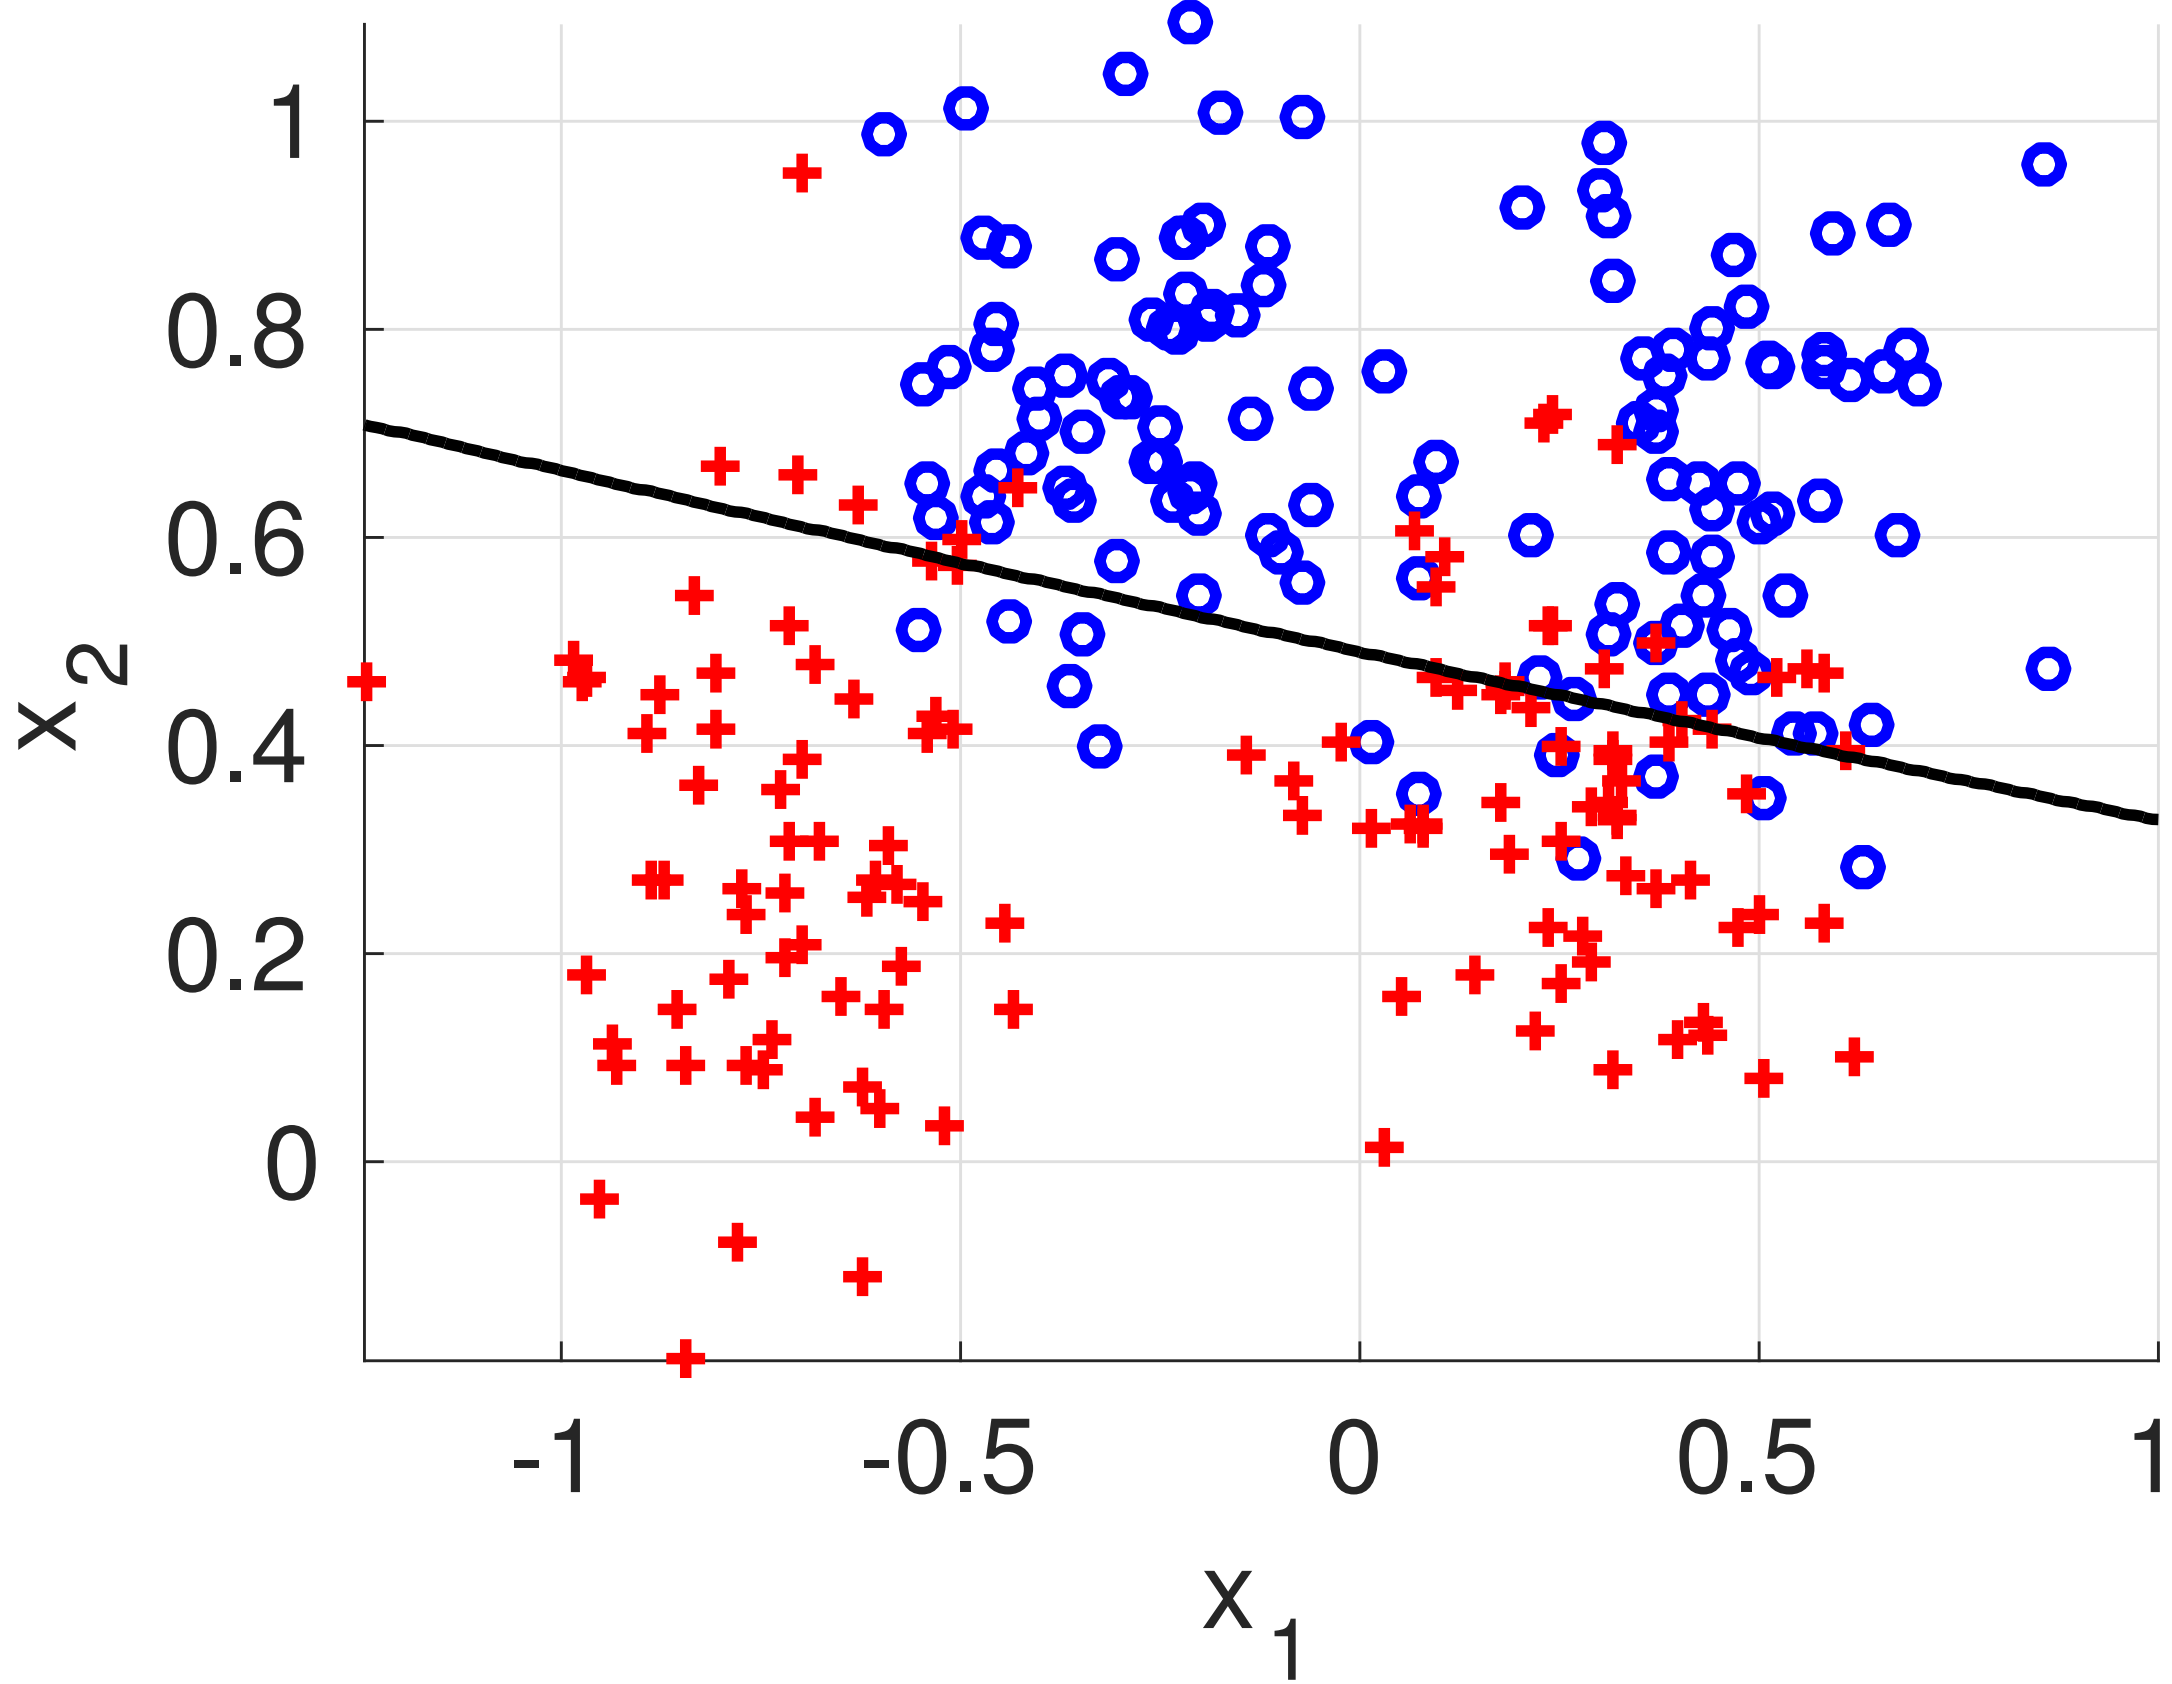




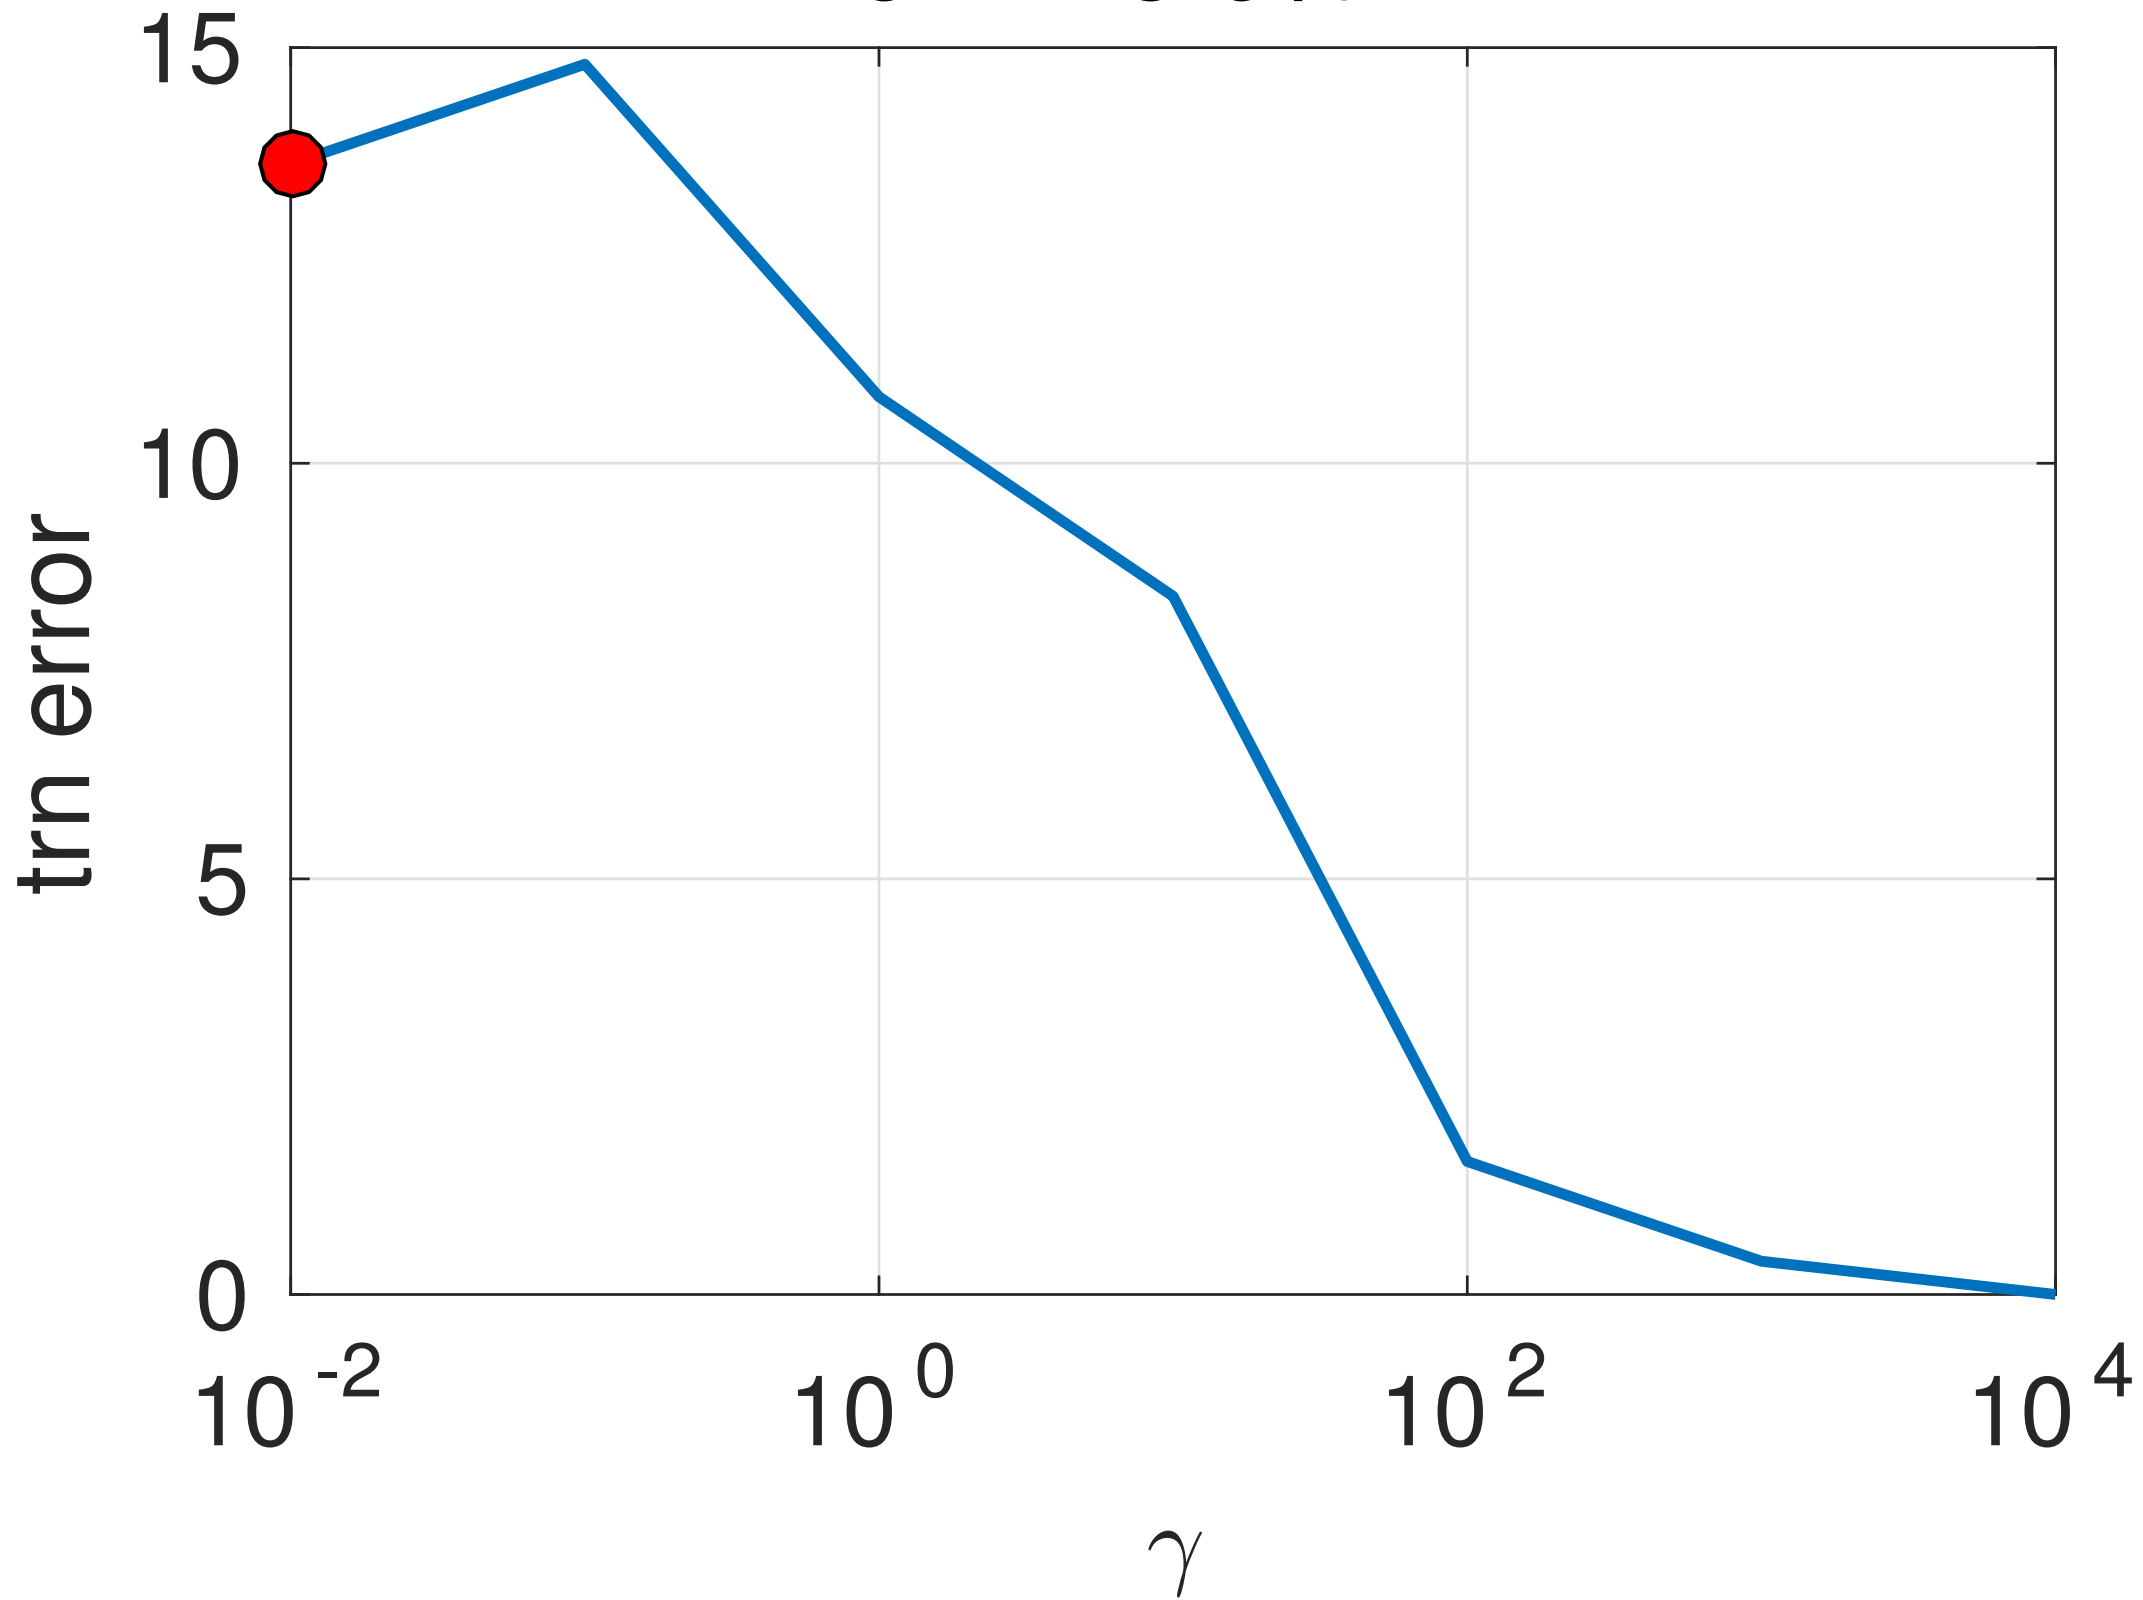


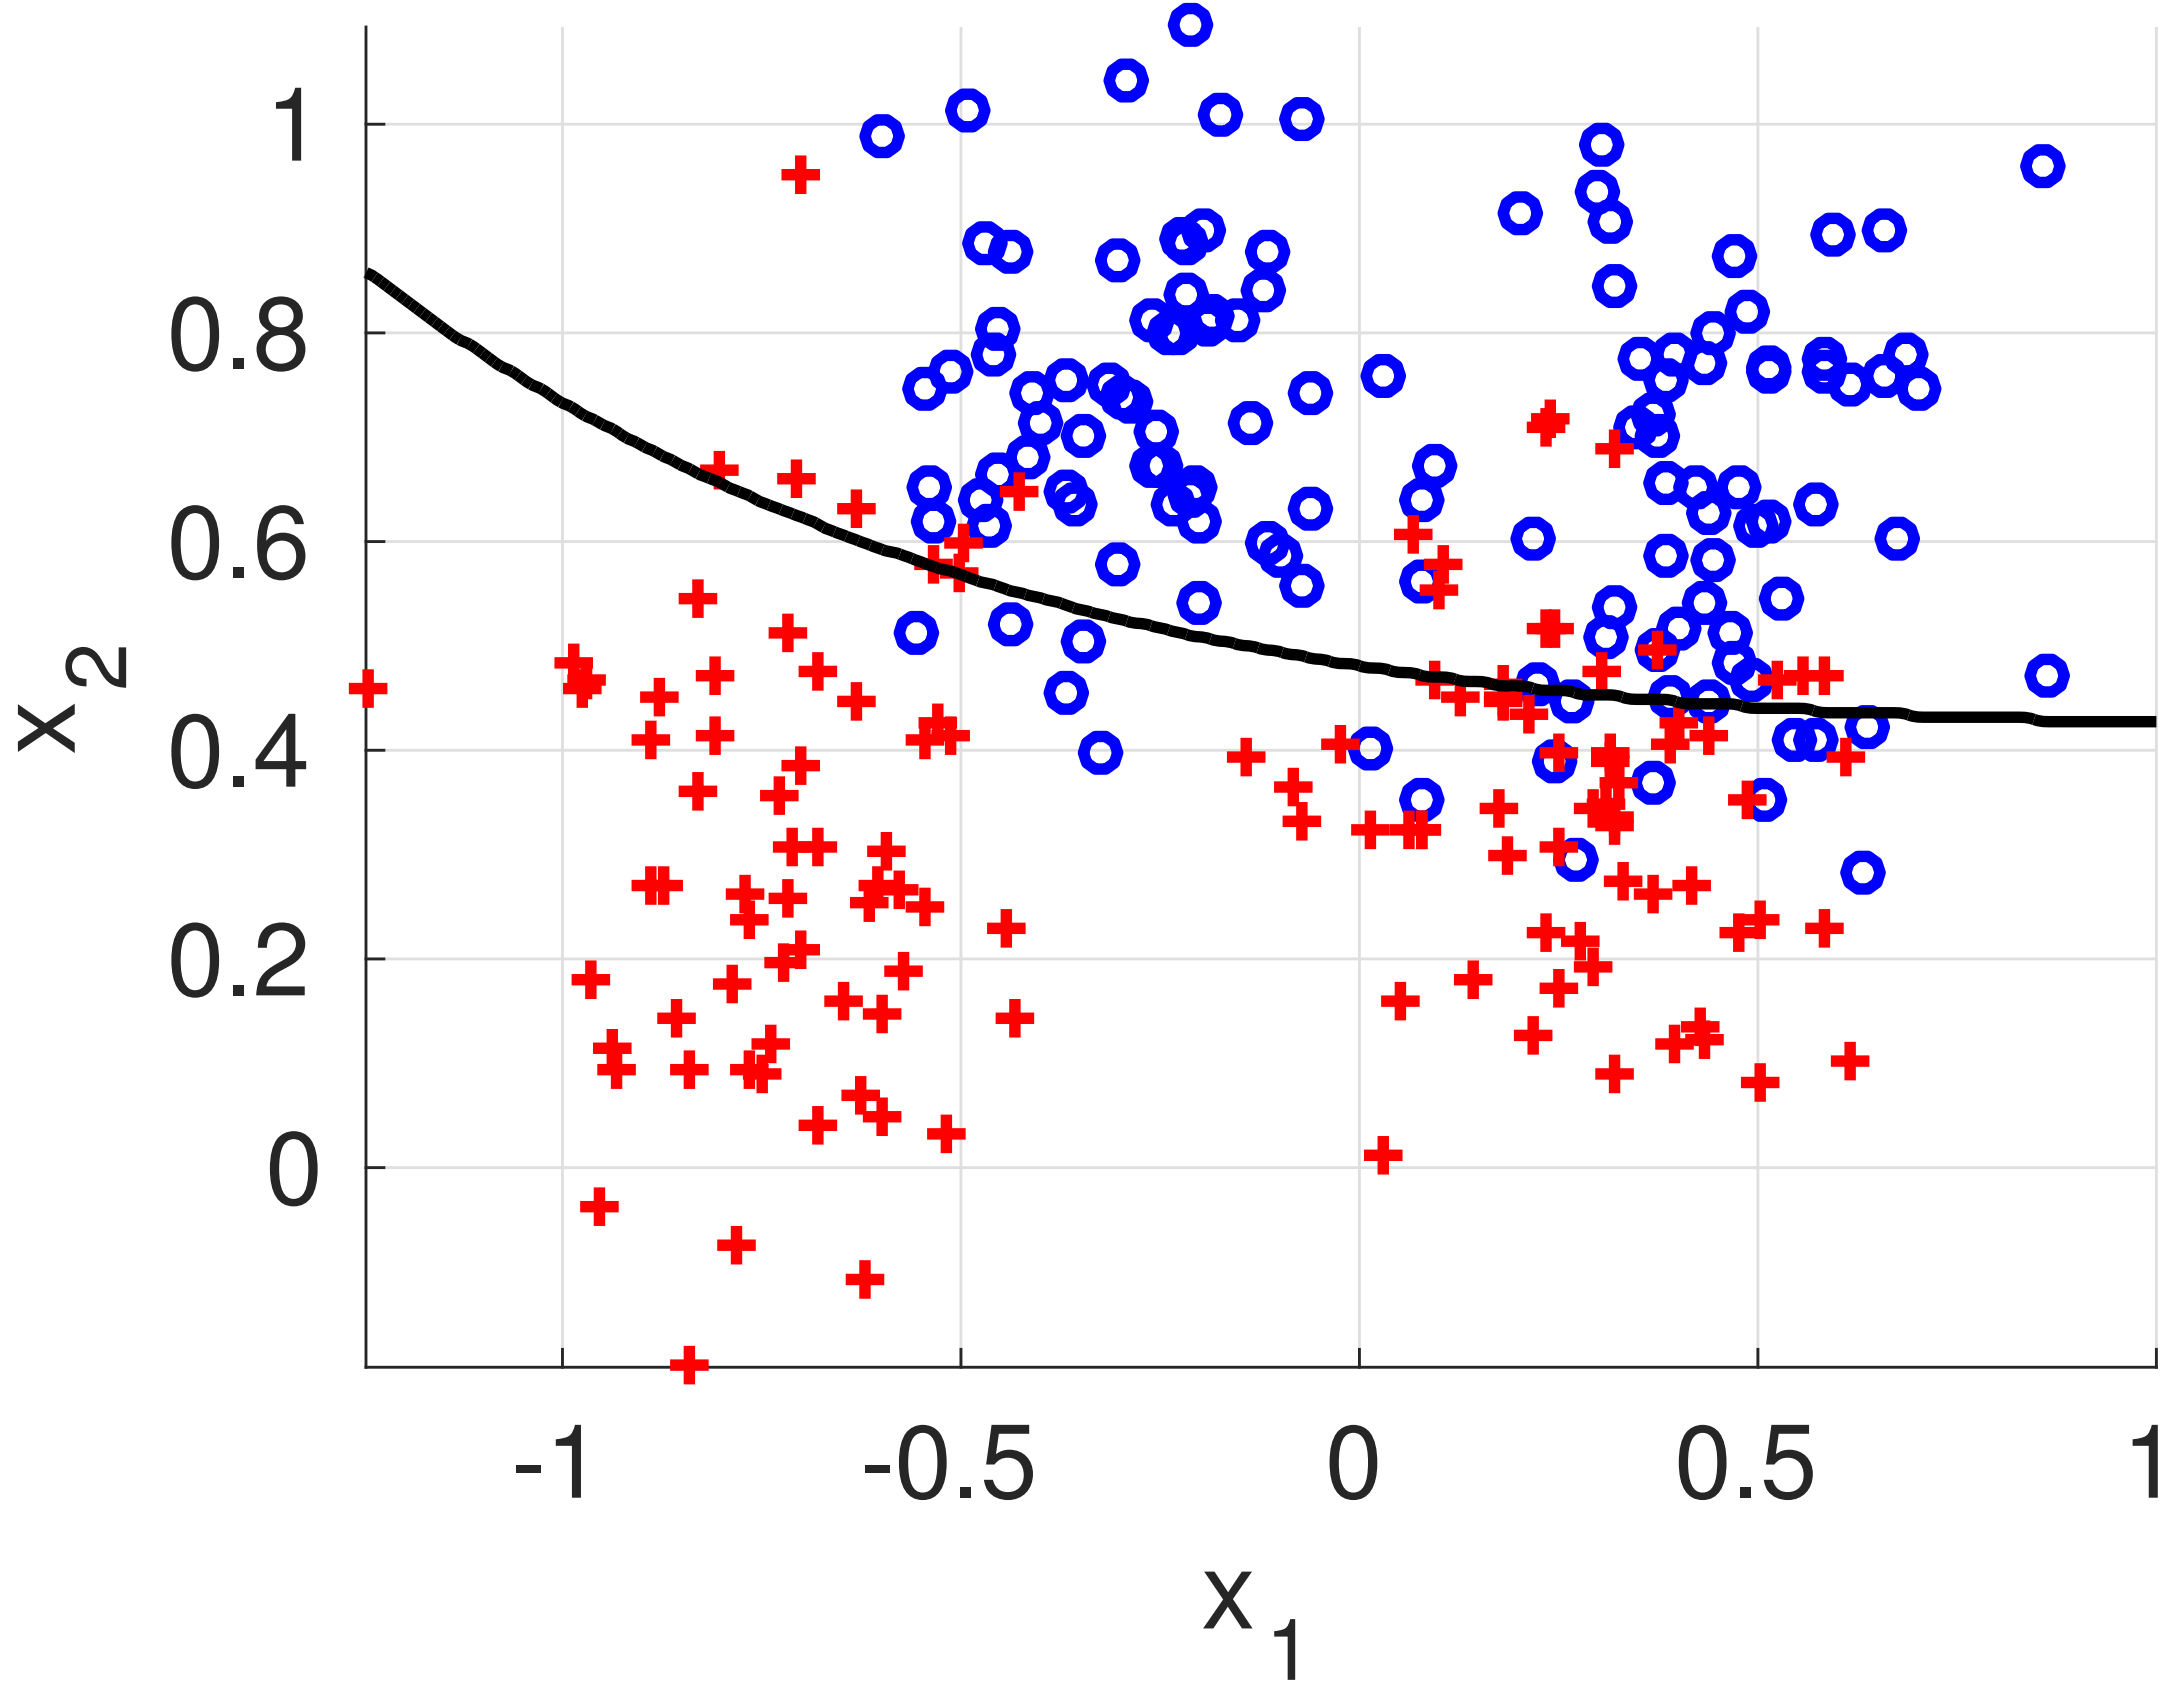




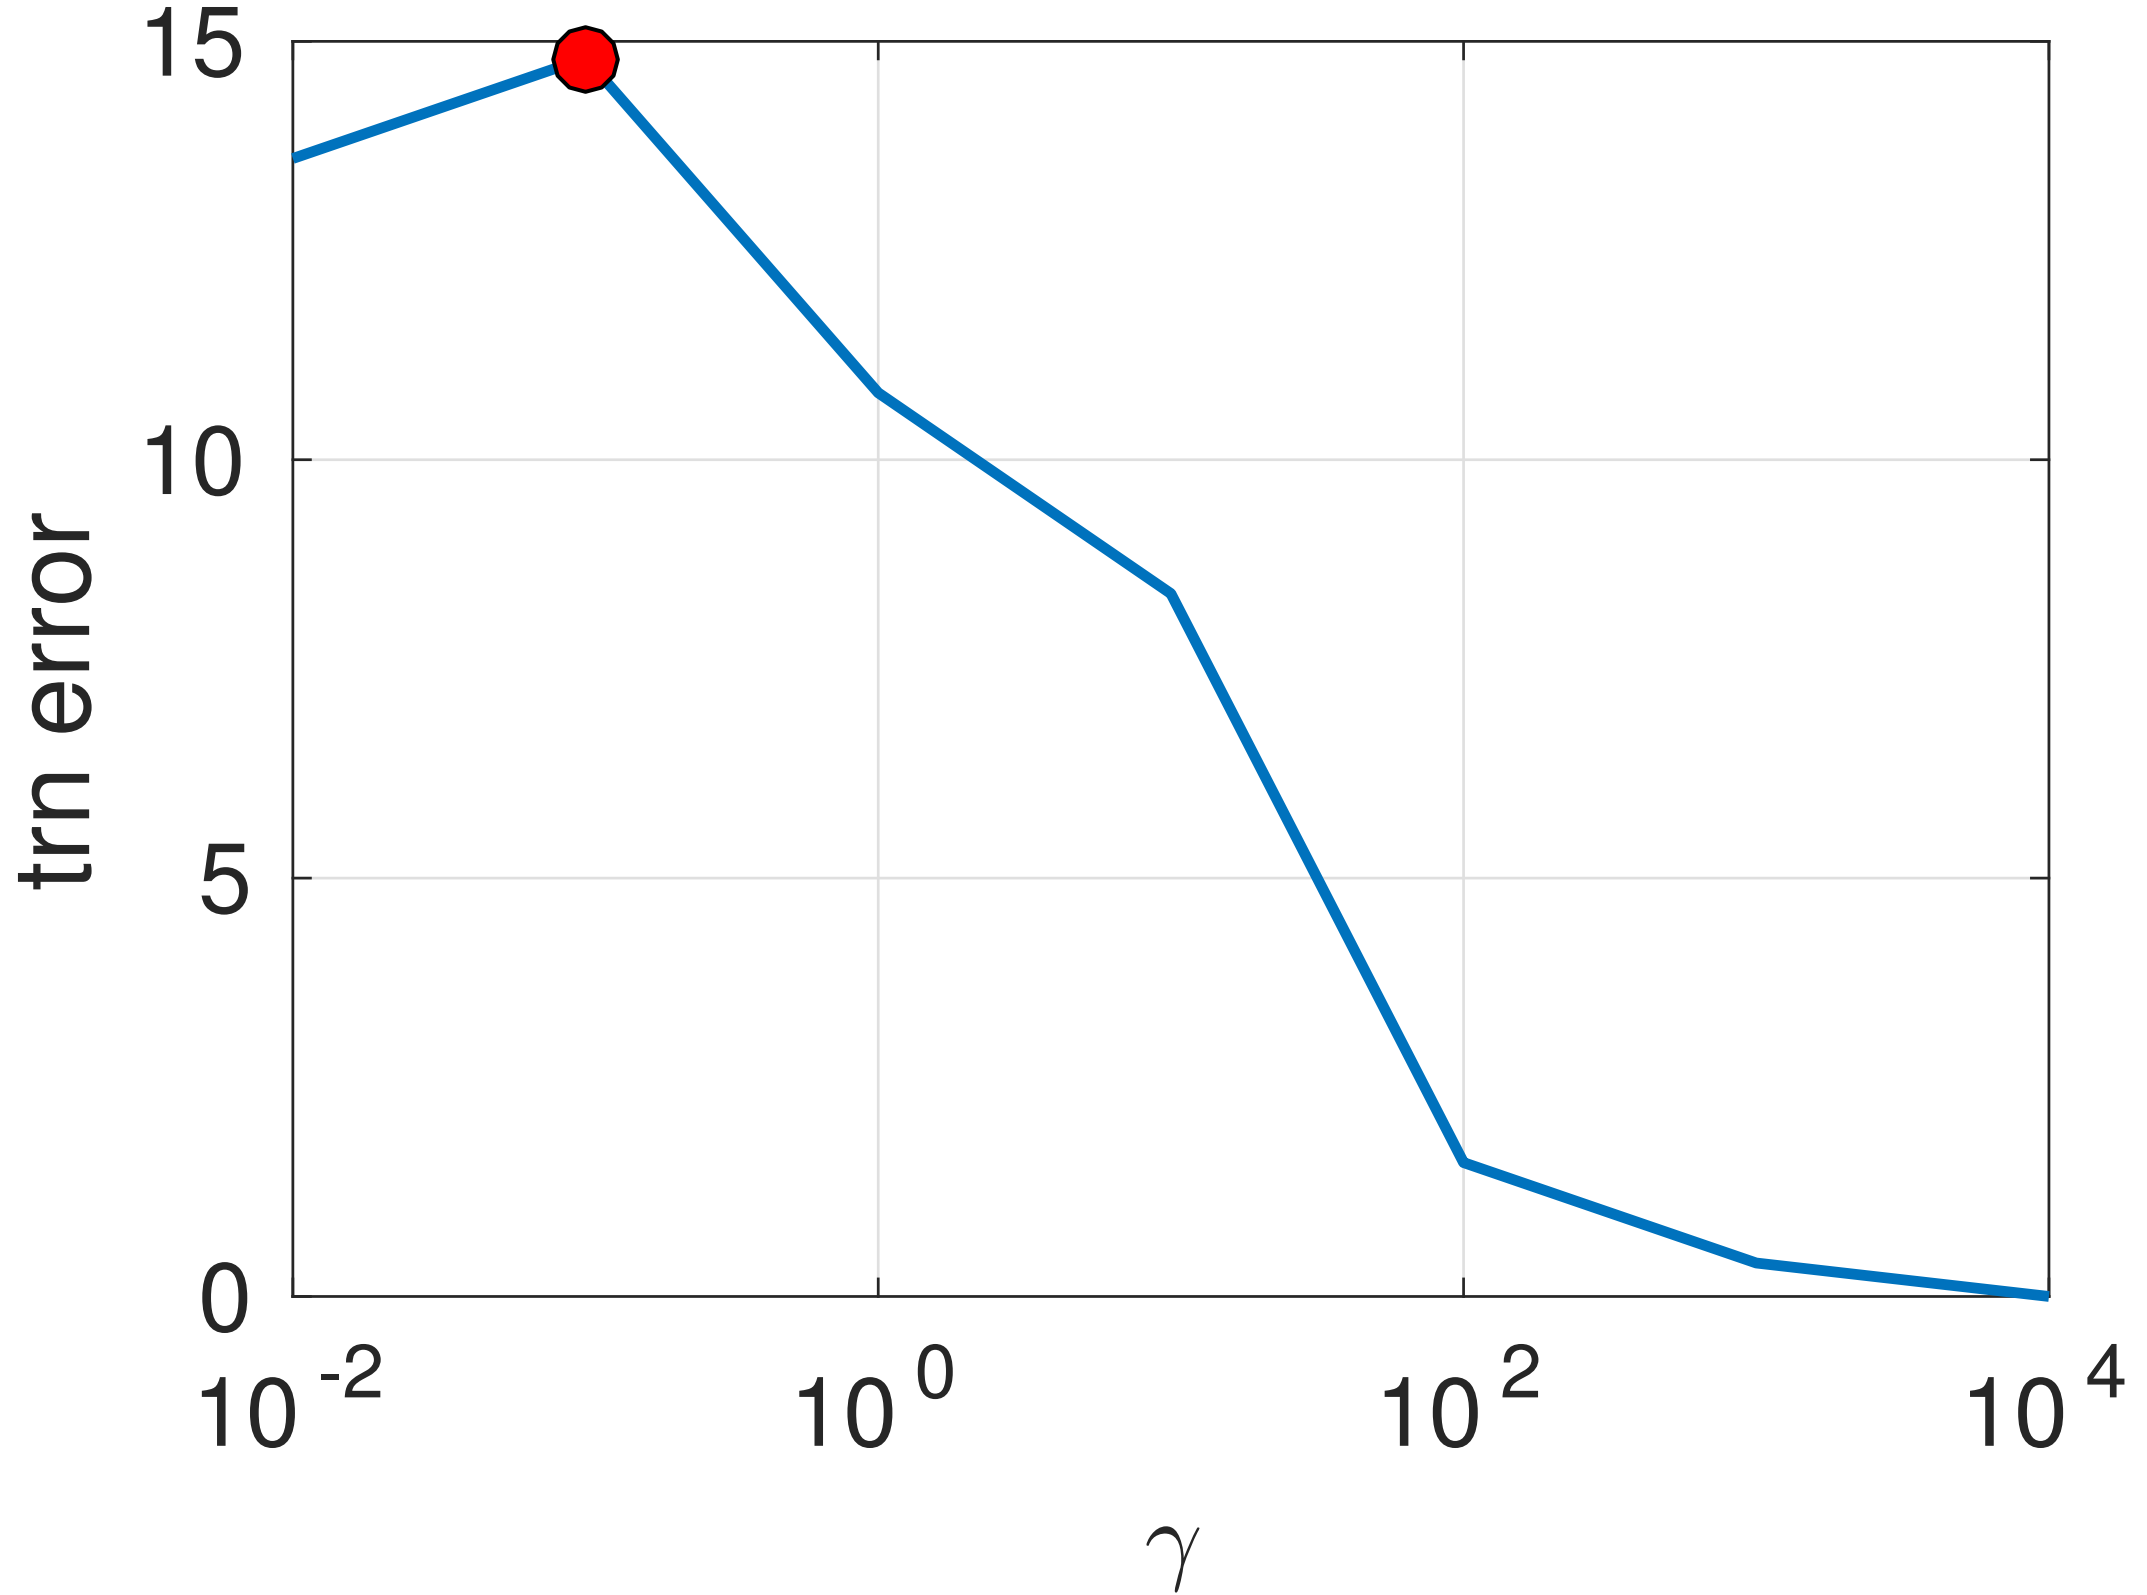


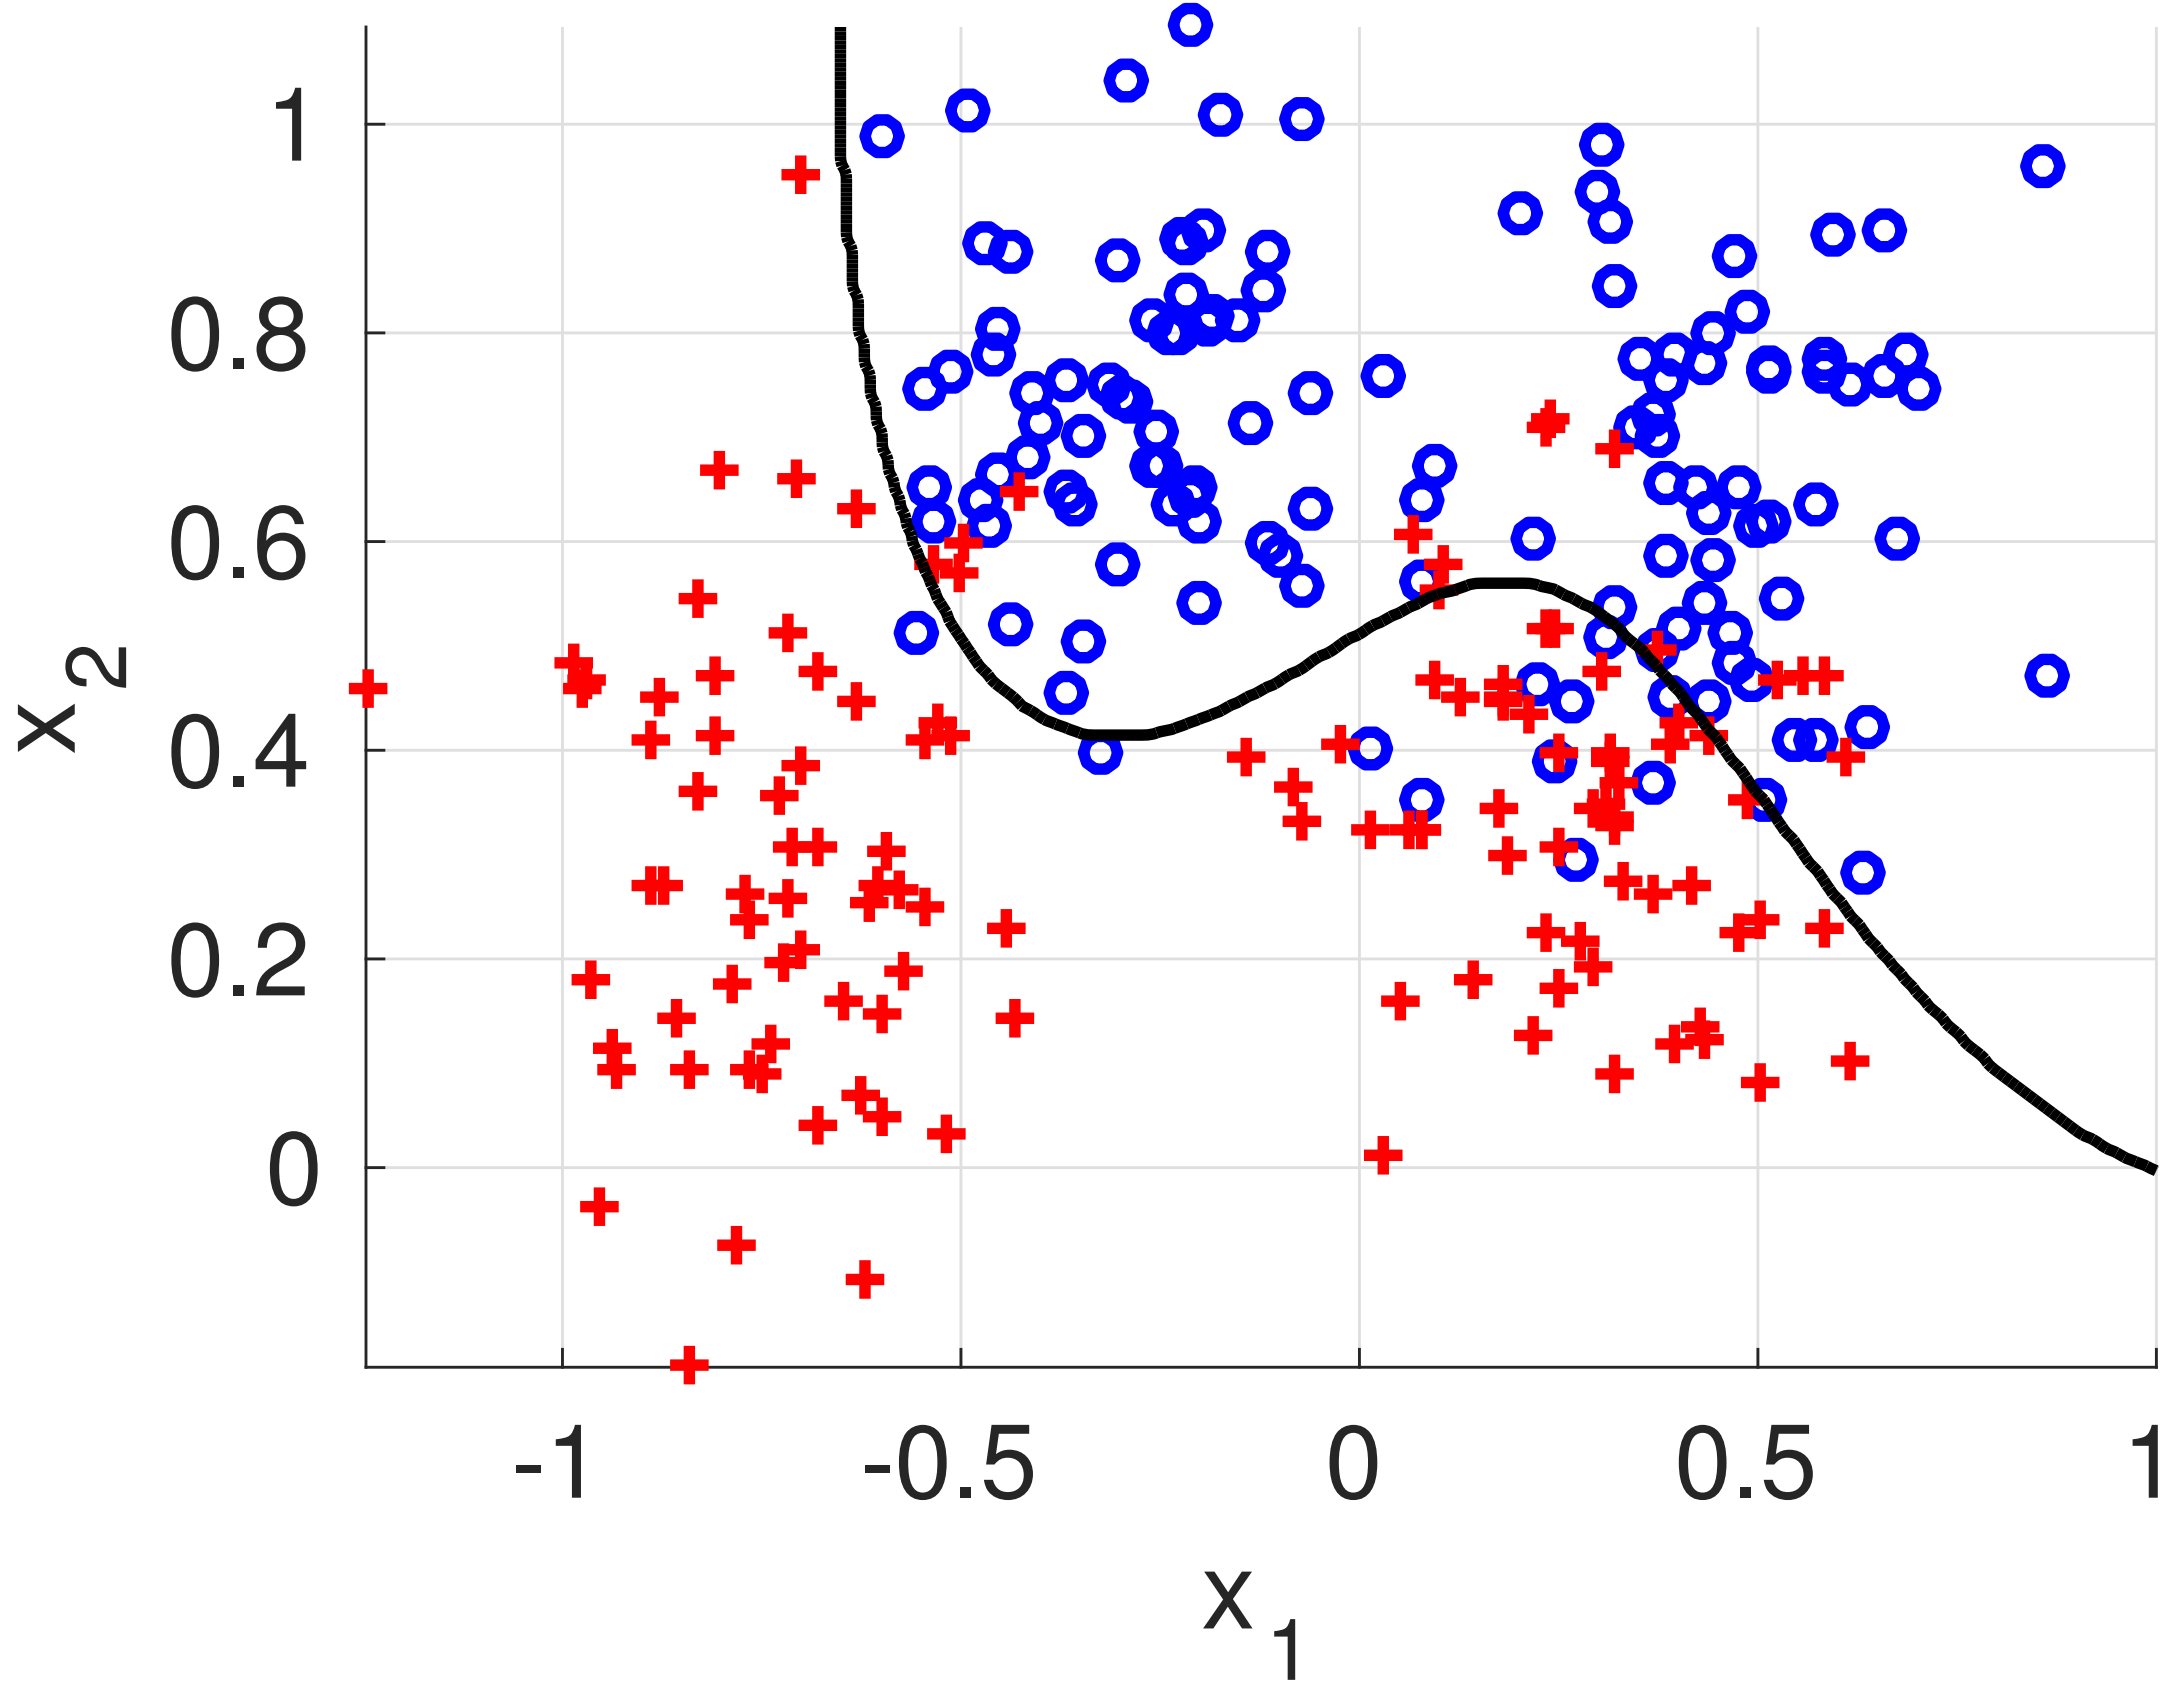
err=13.6%



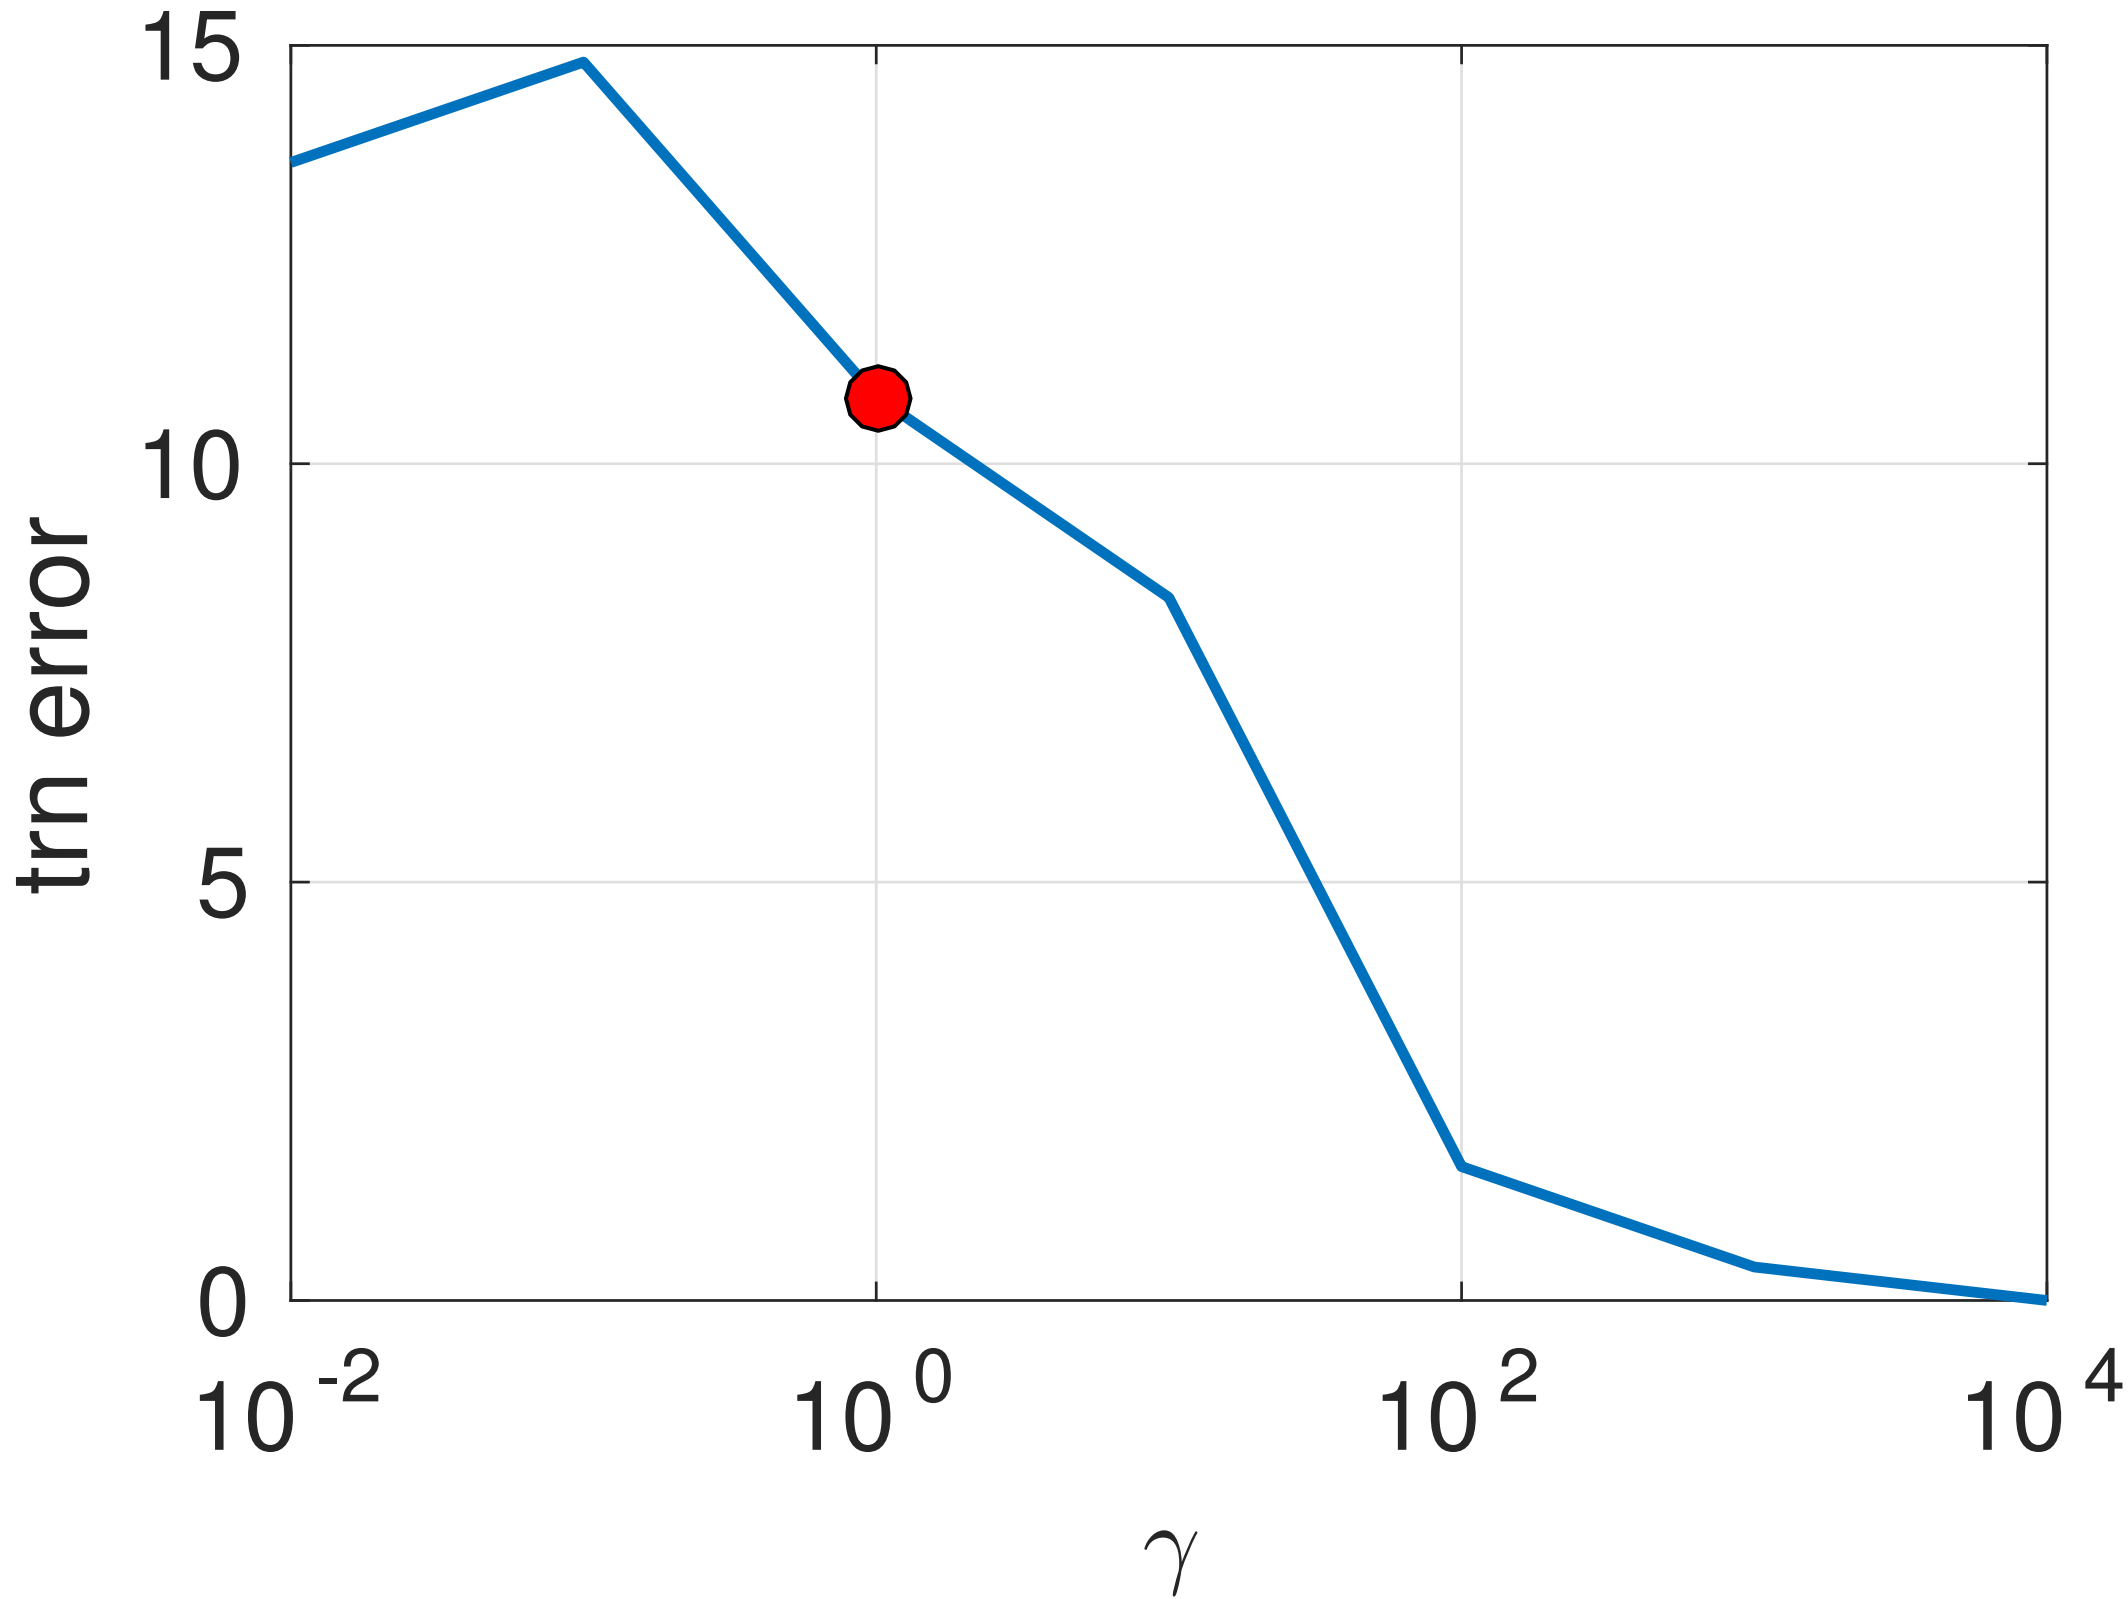


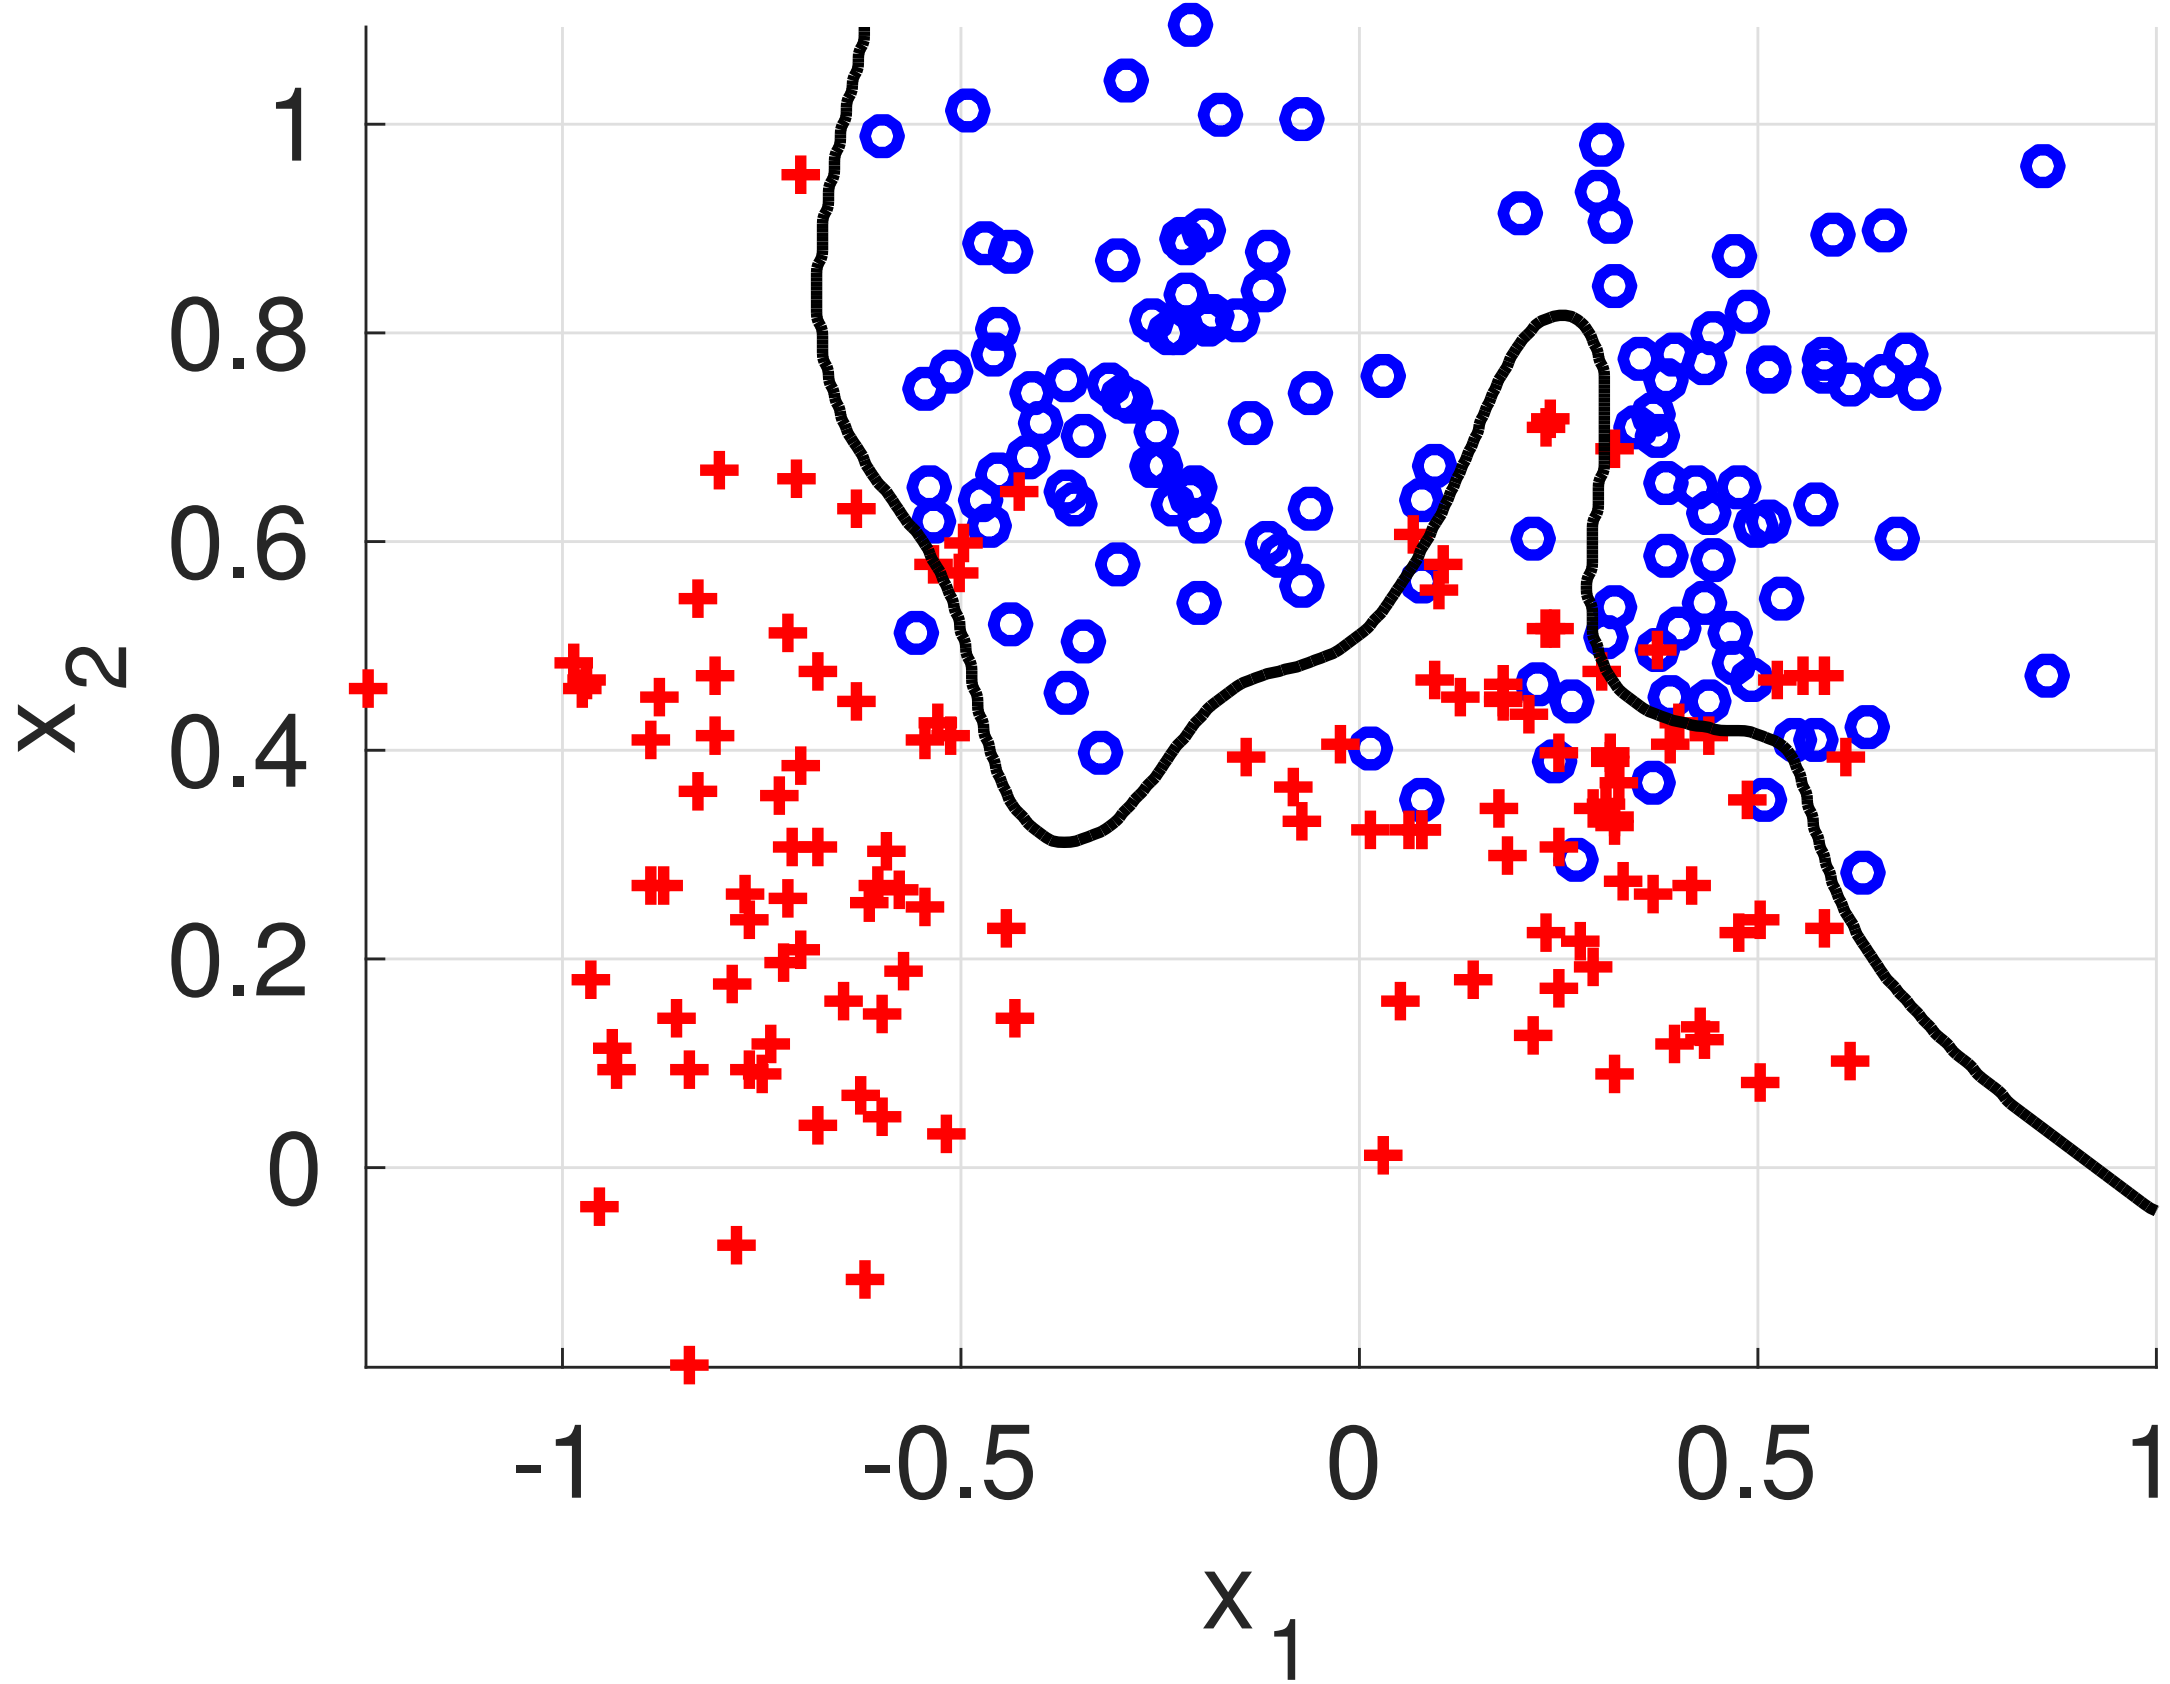
err=14.8%



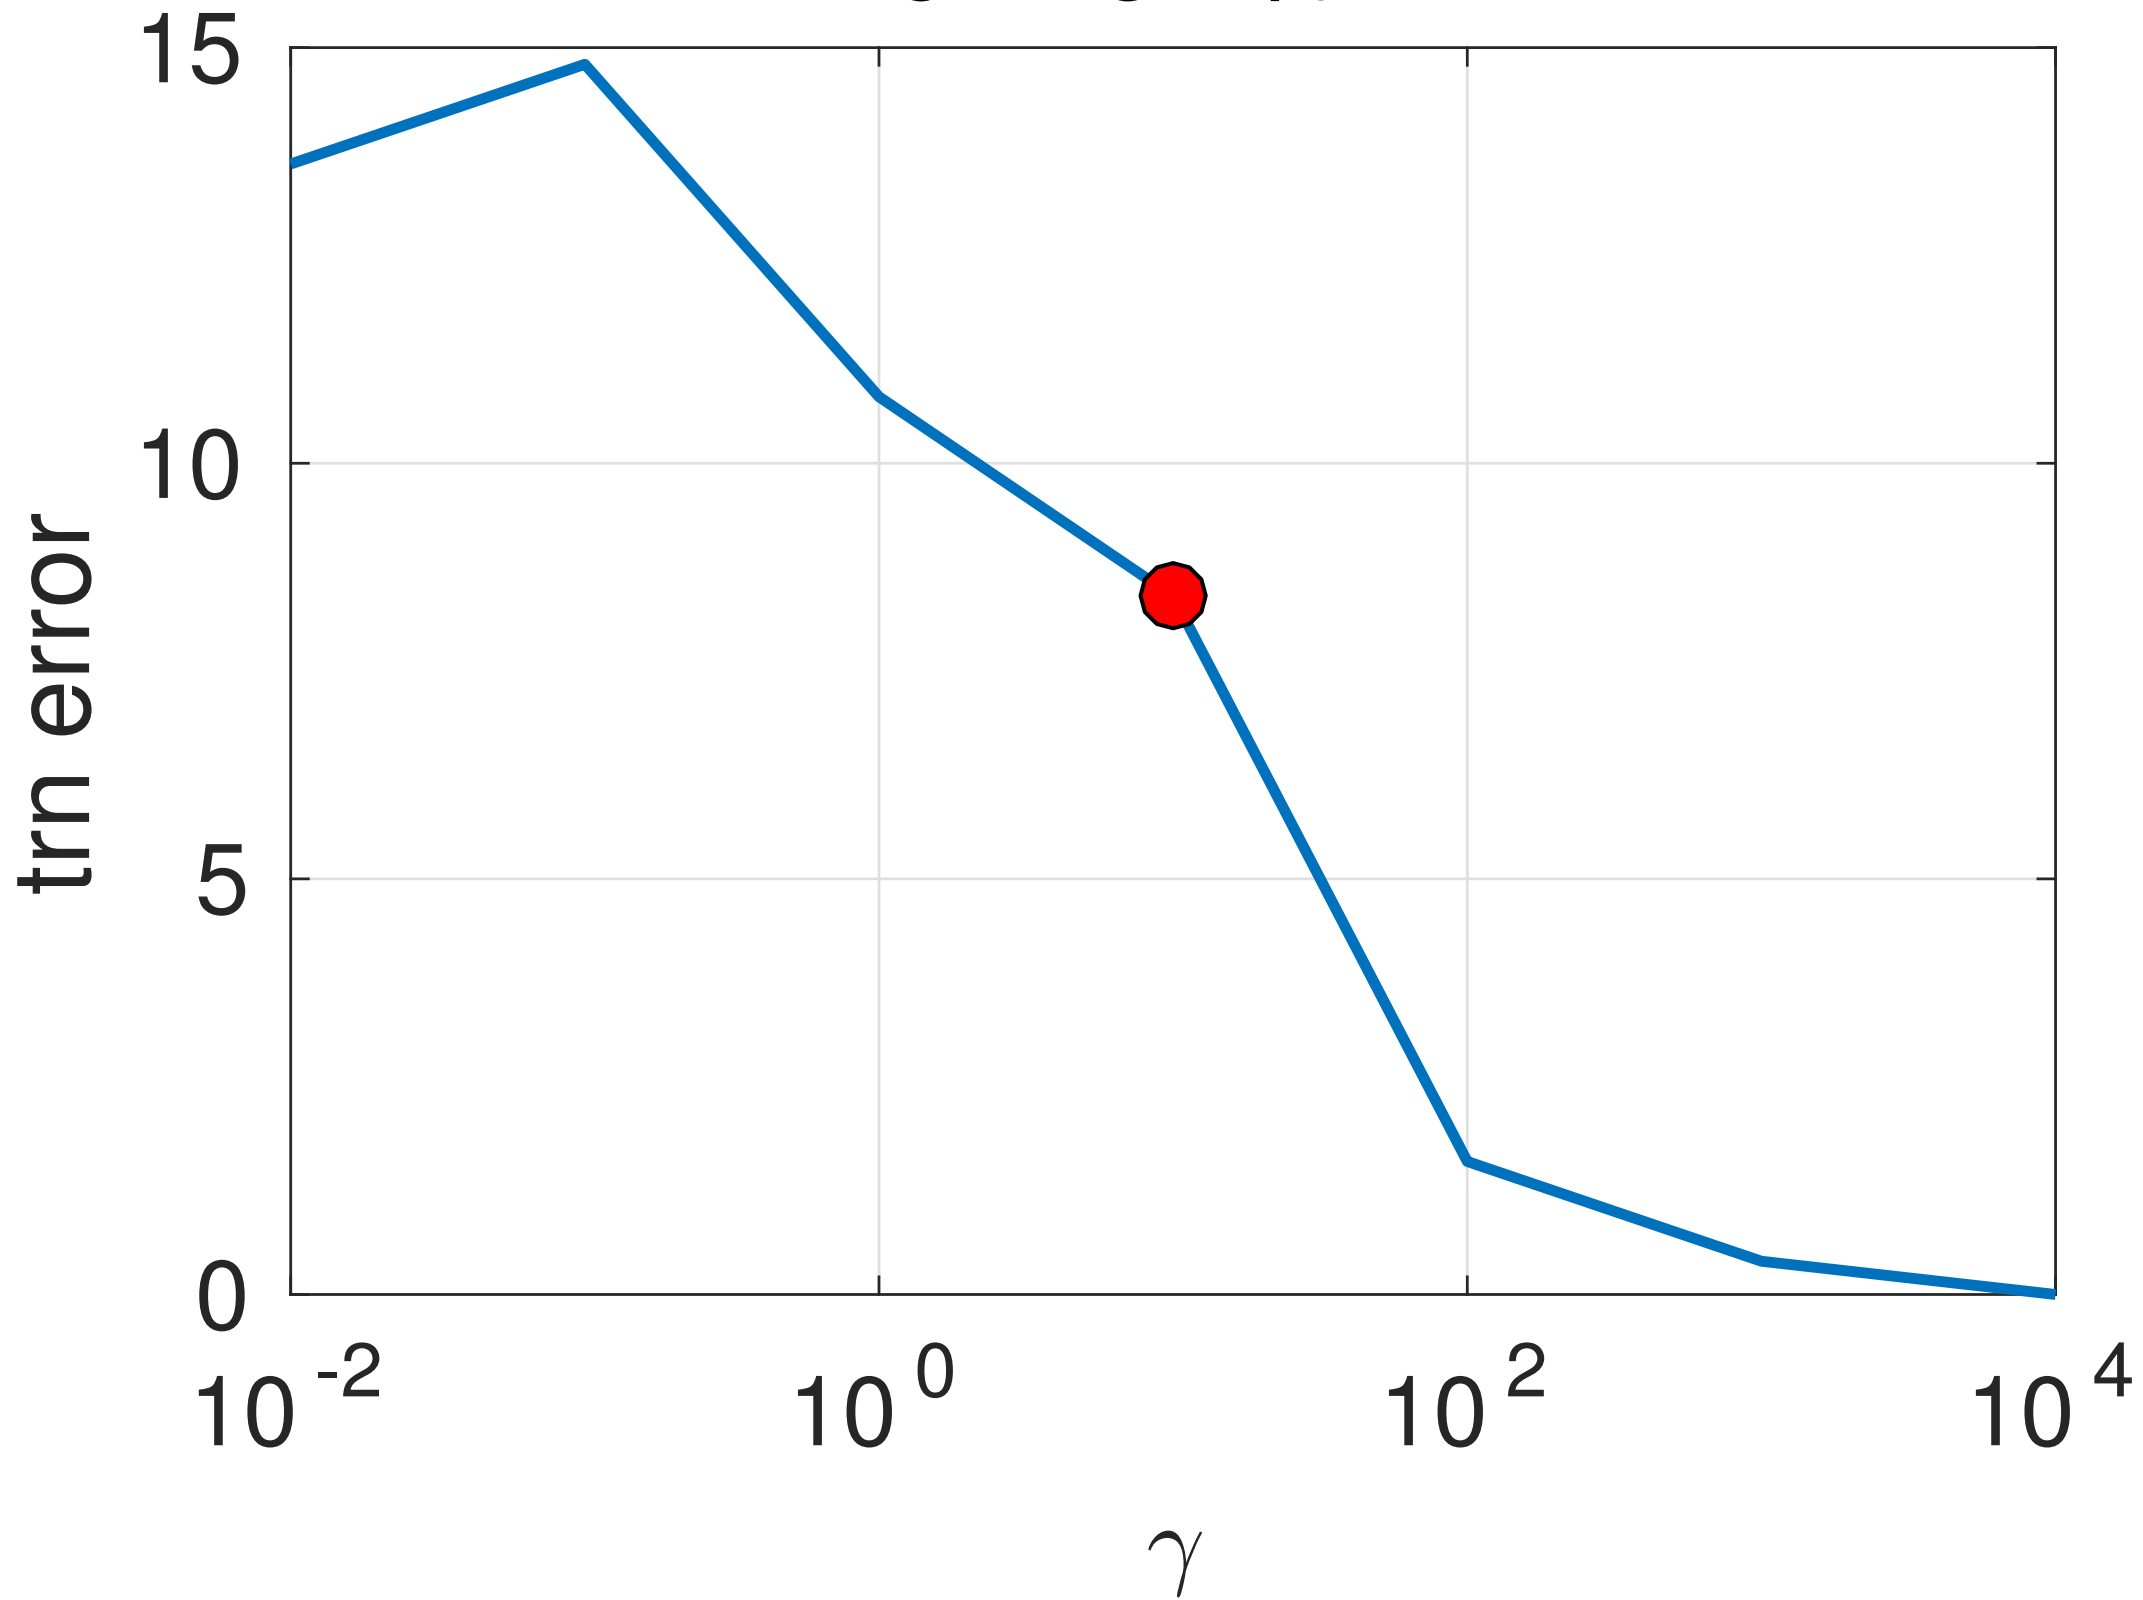


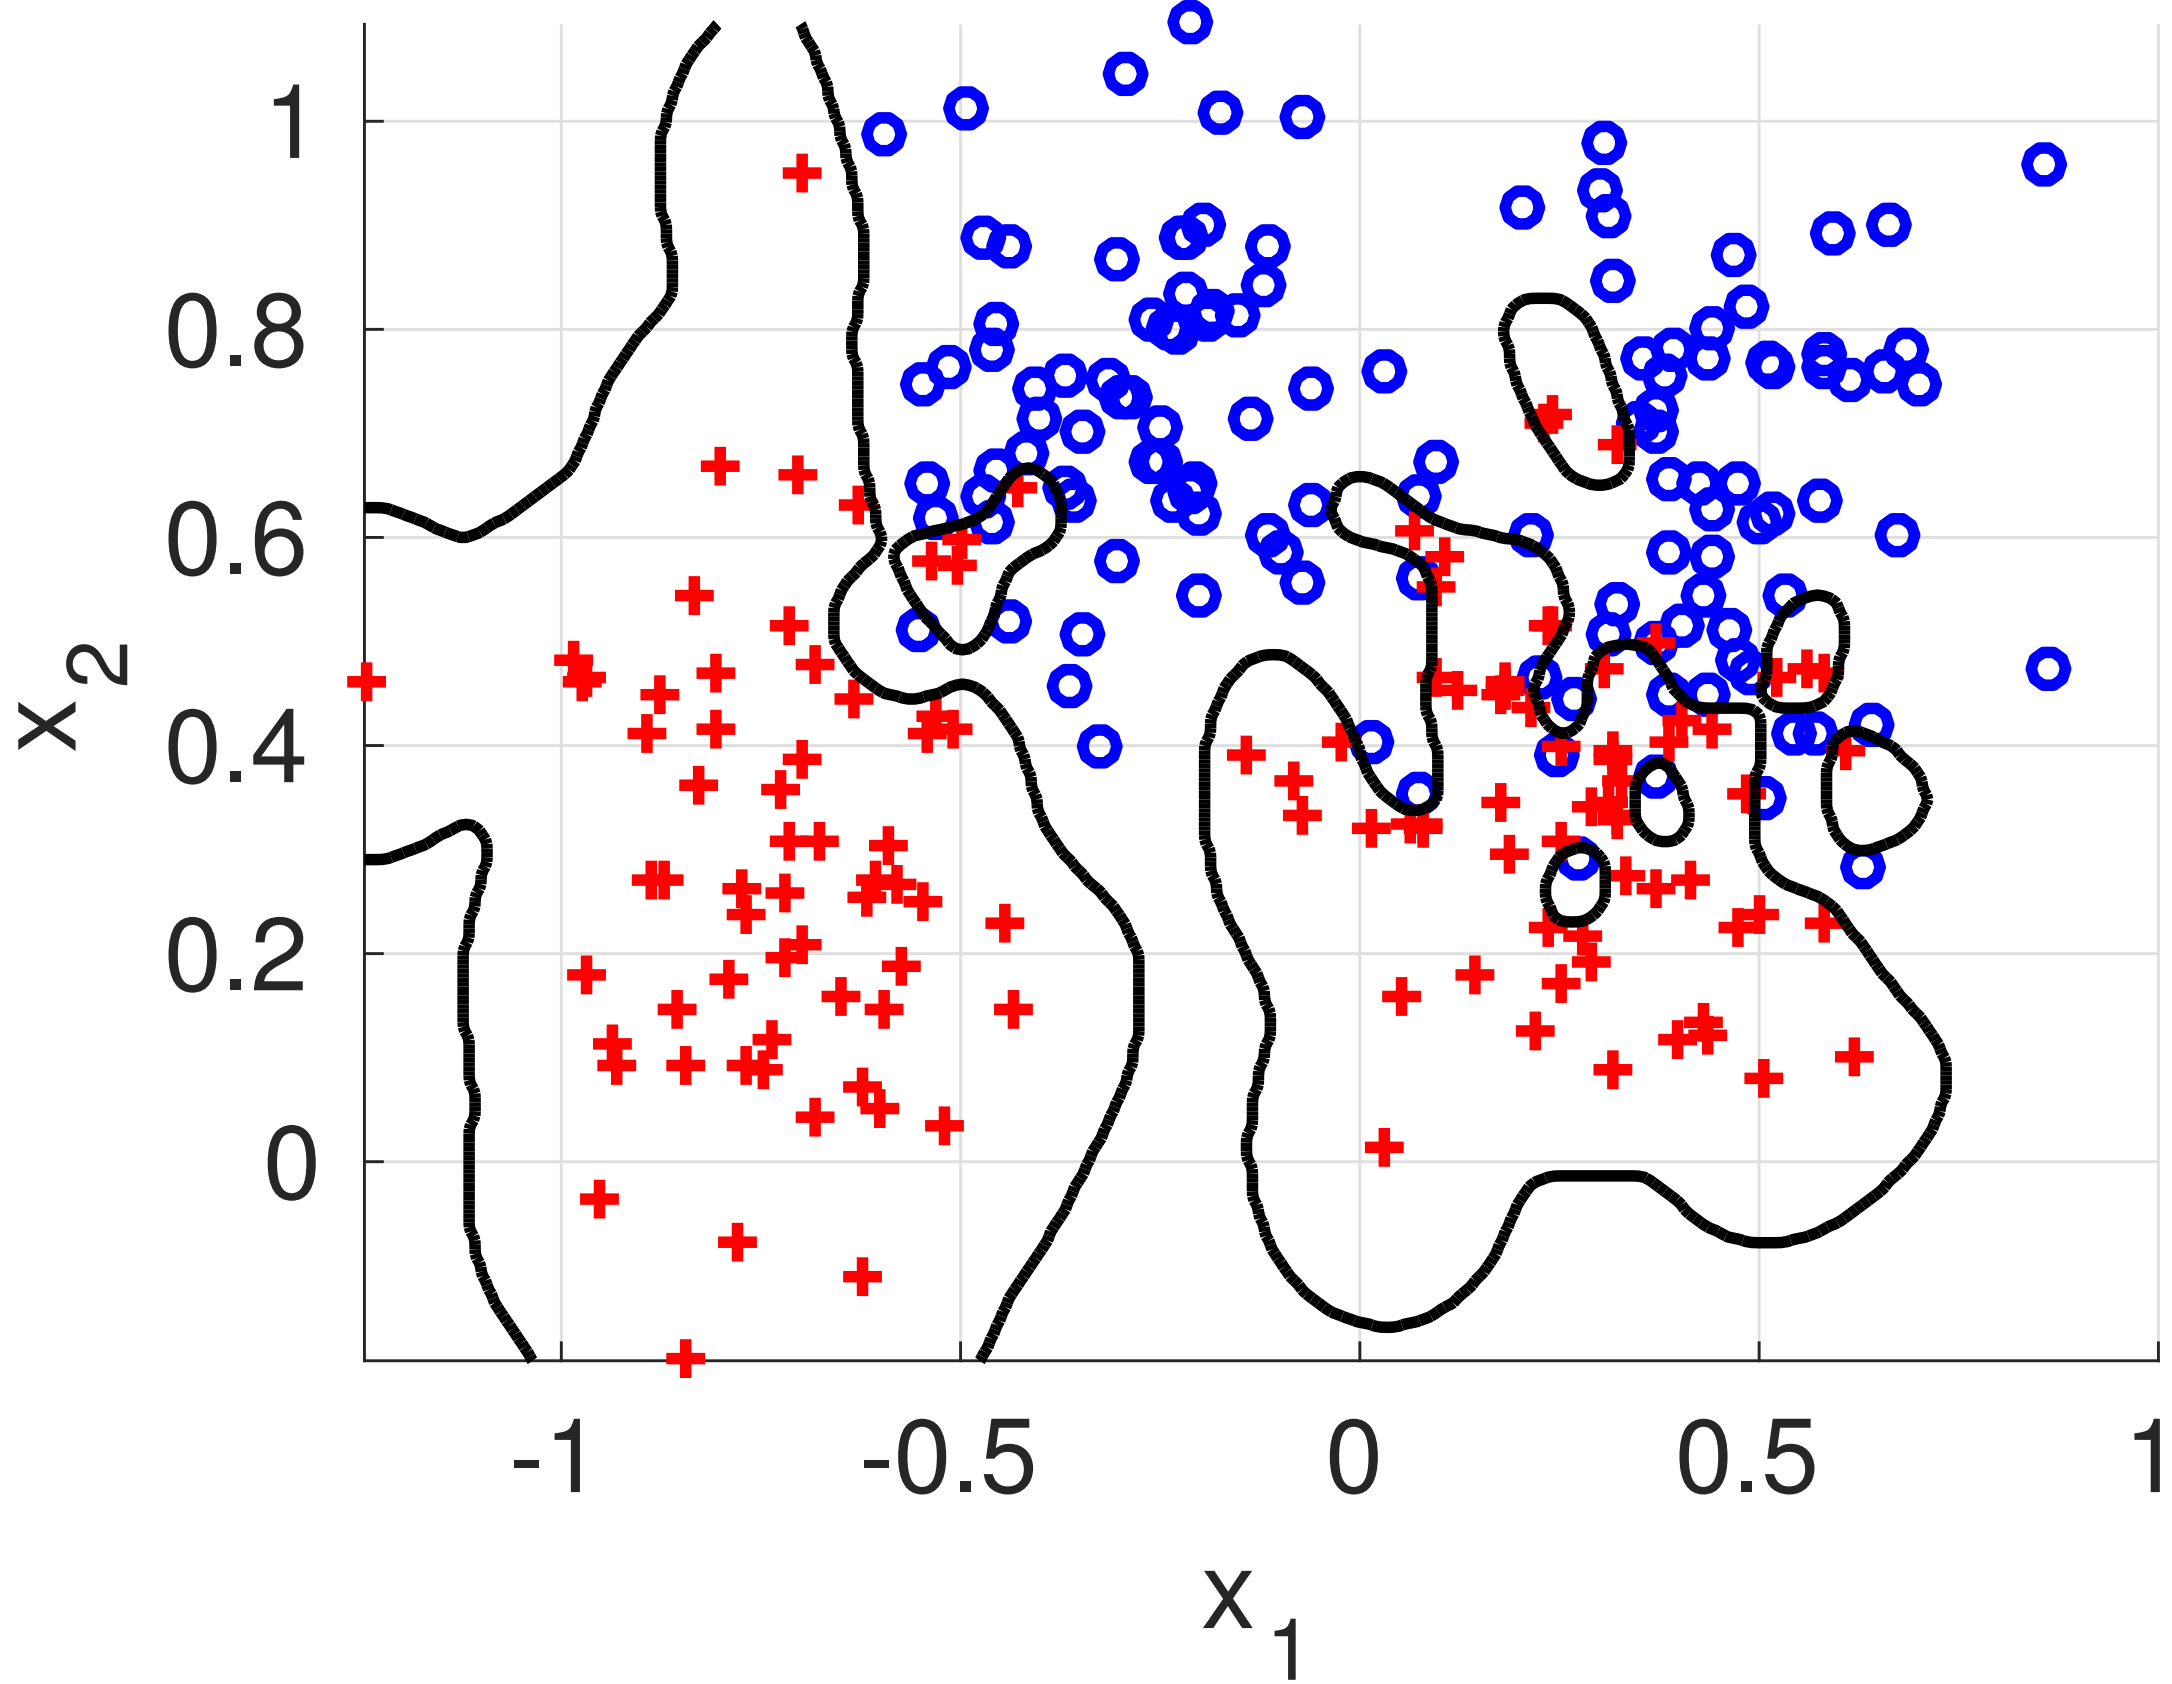
err=10.8%



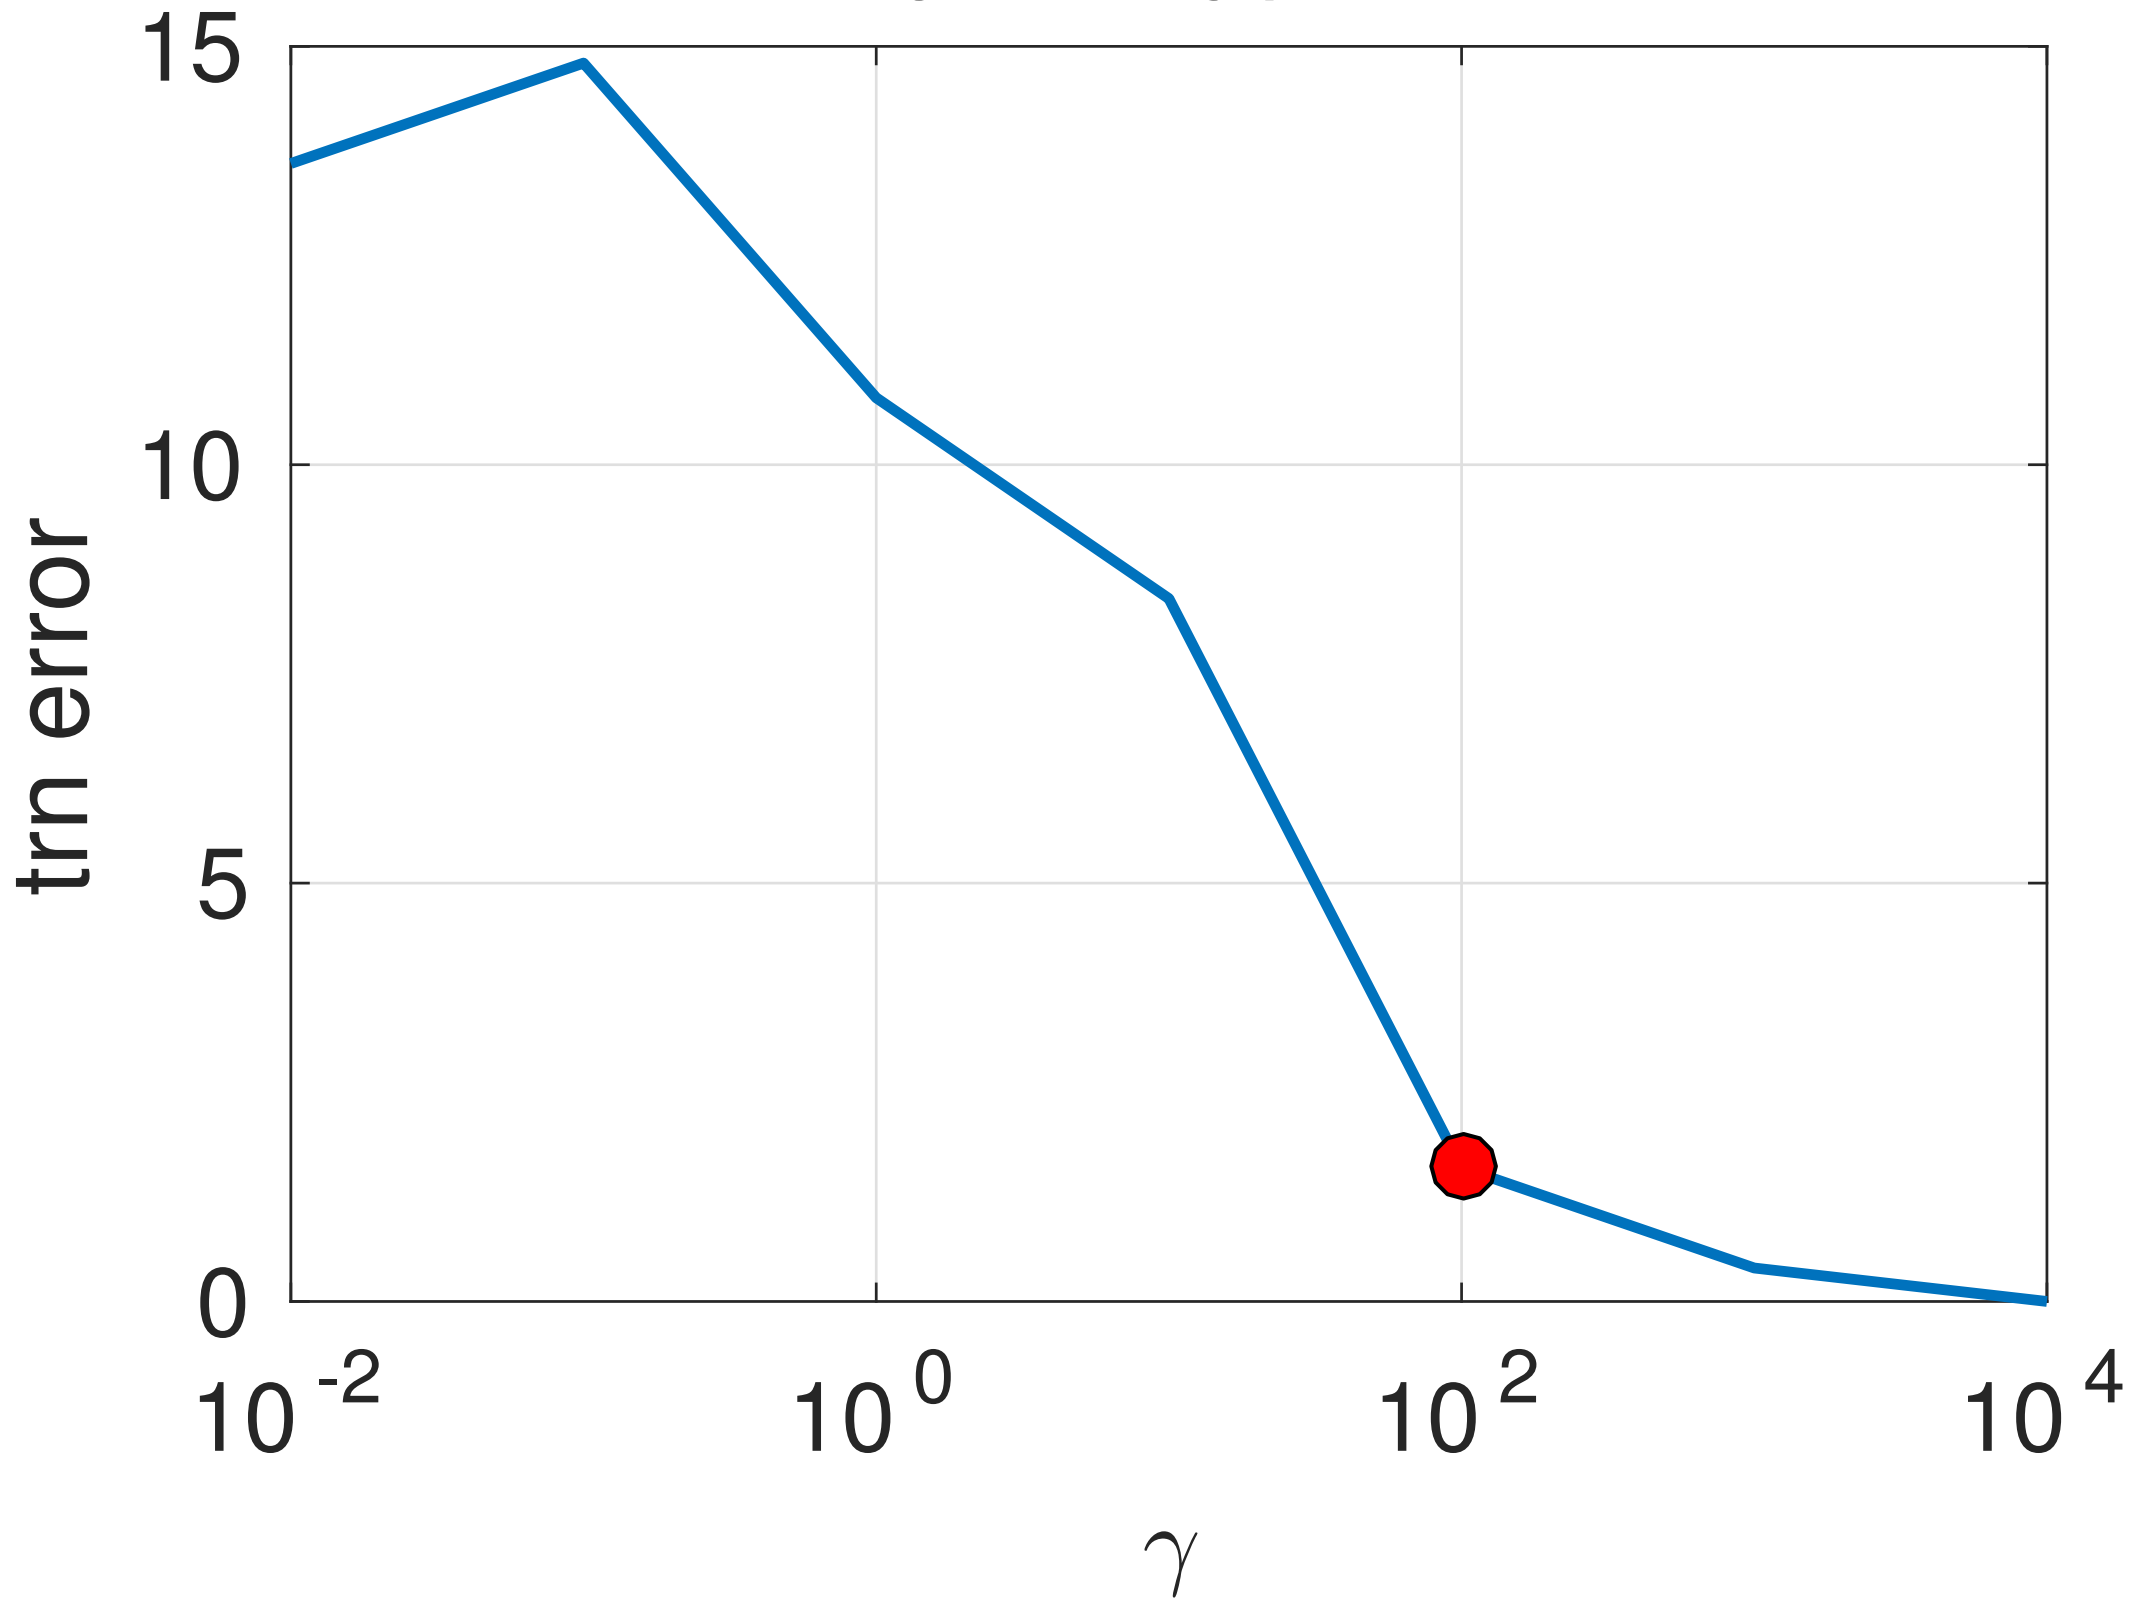


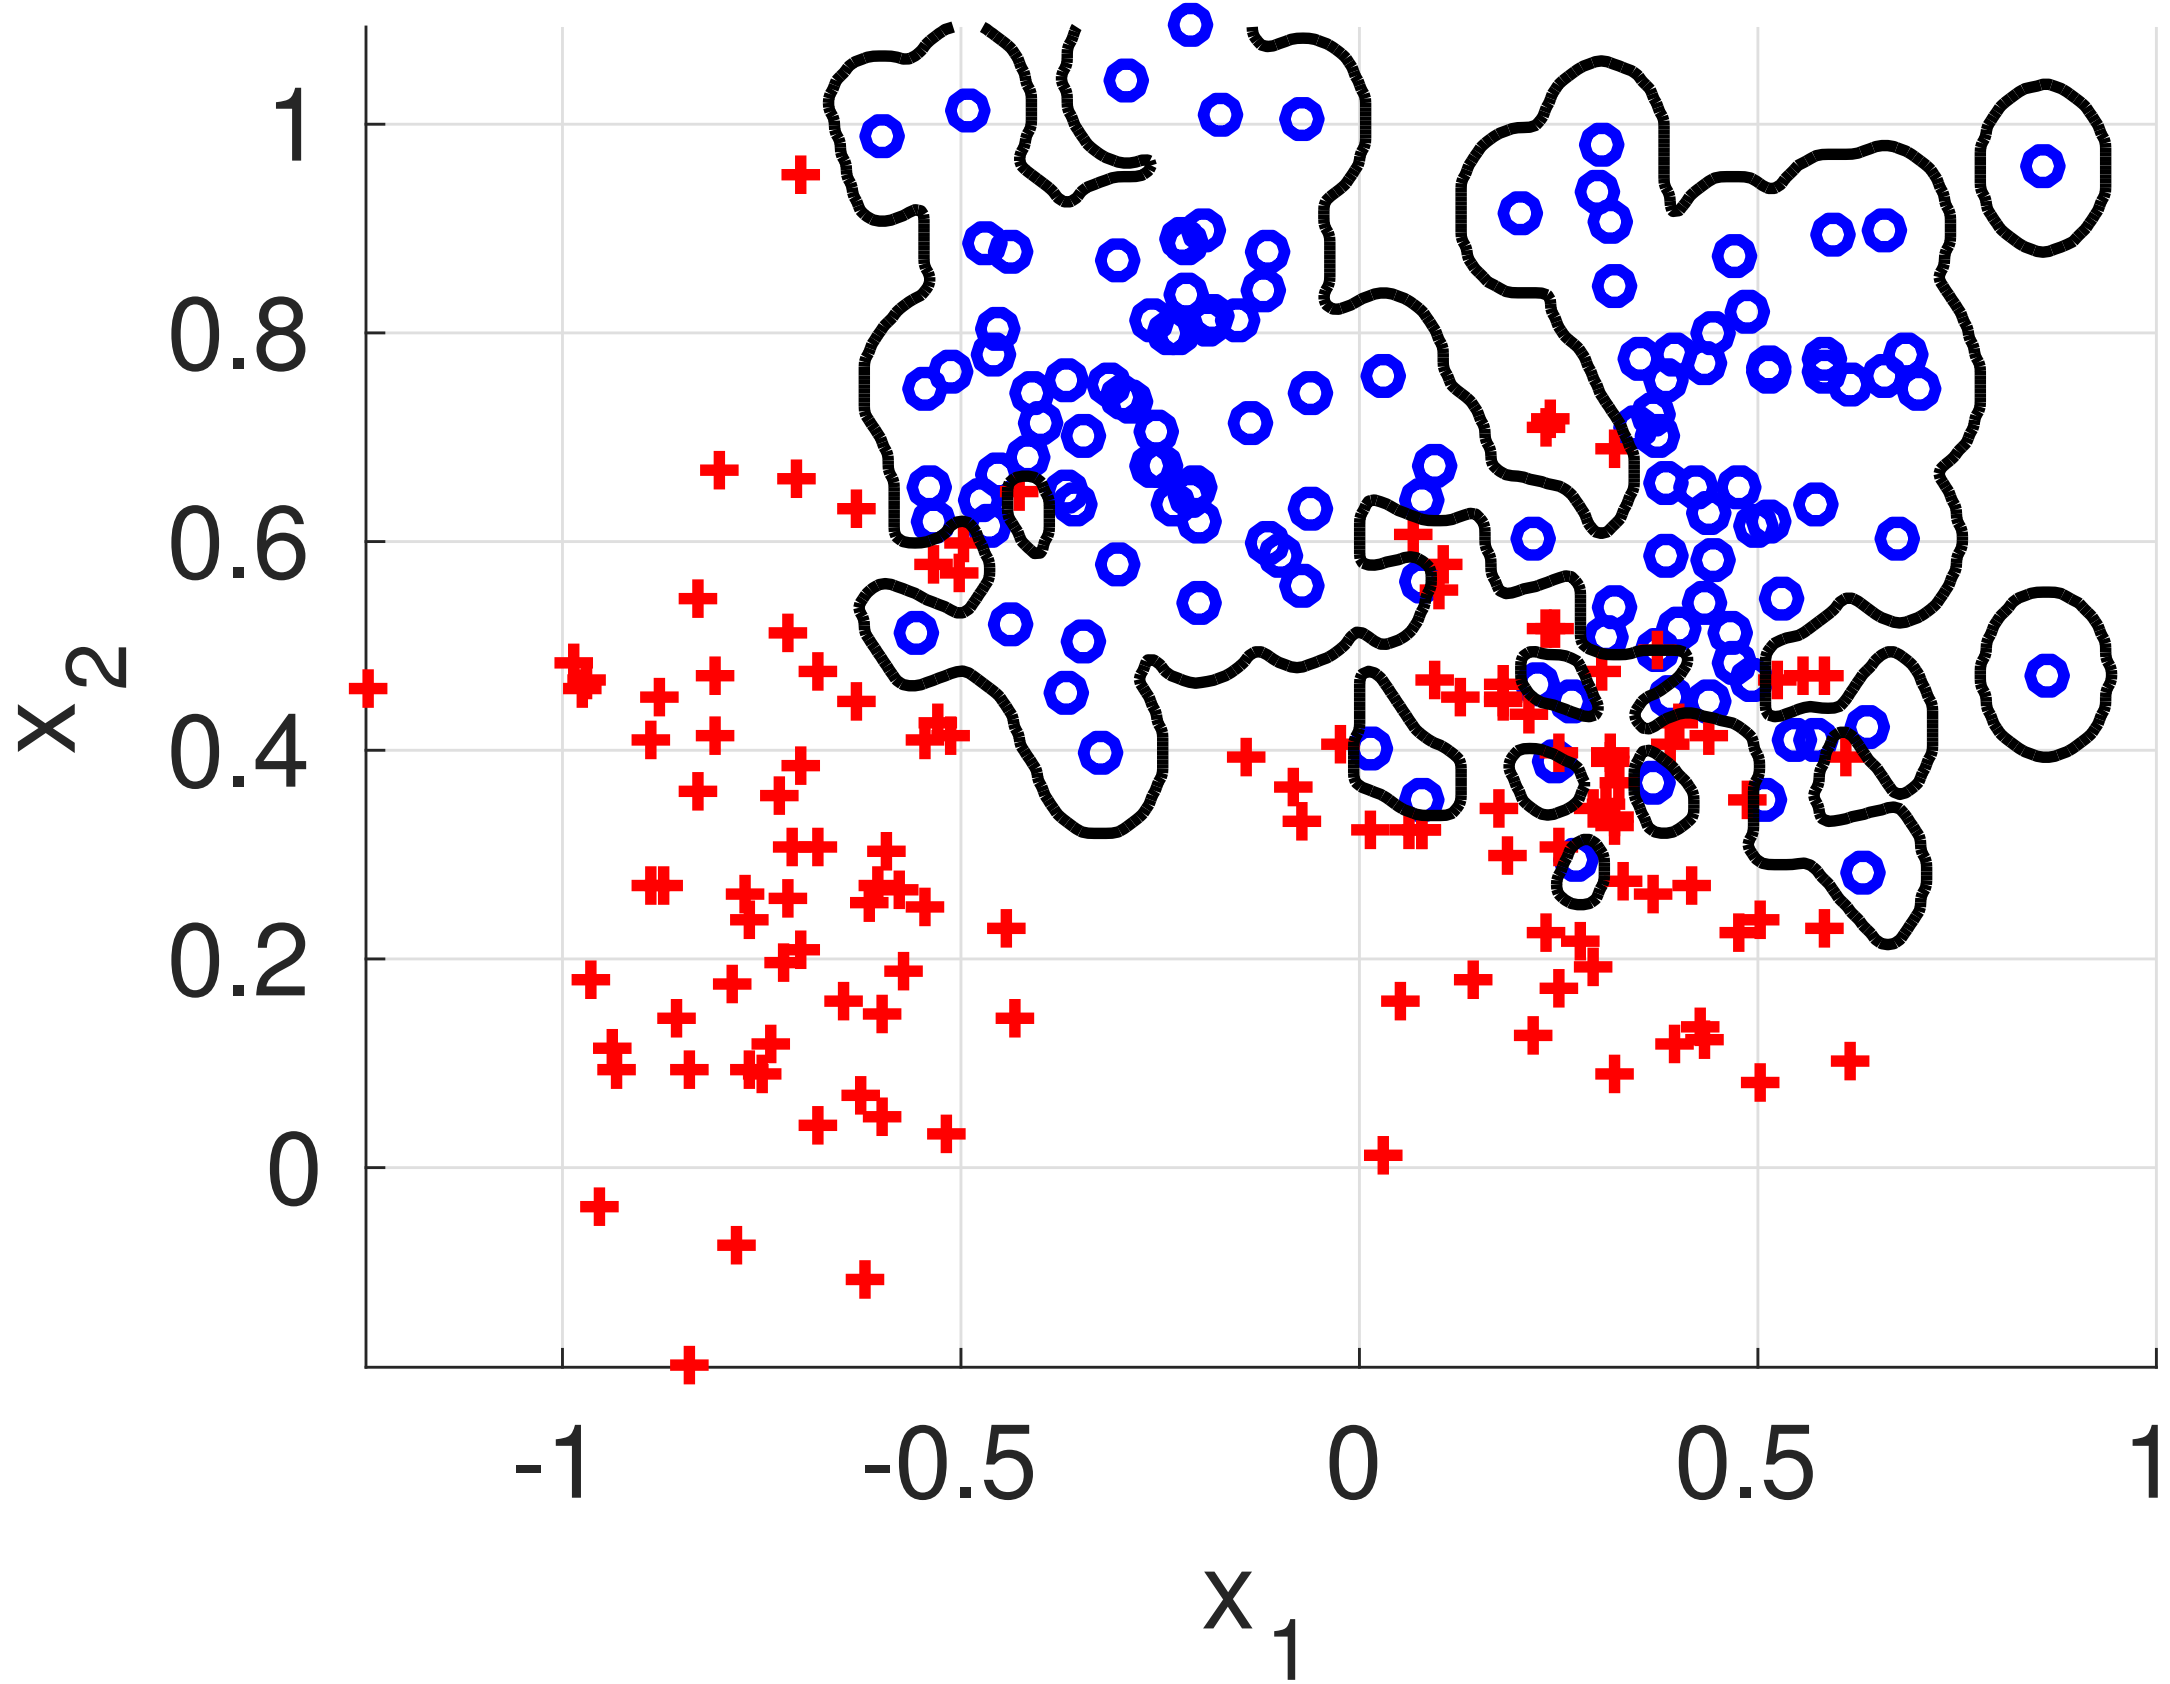
err=8.4%



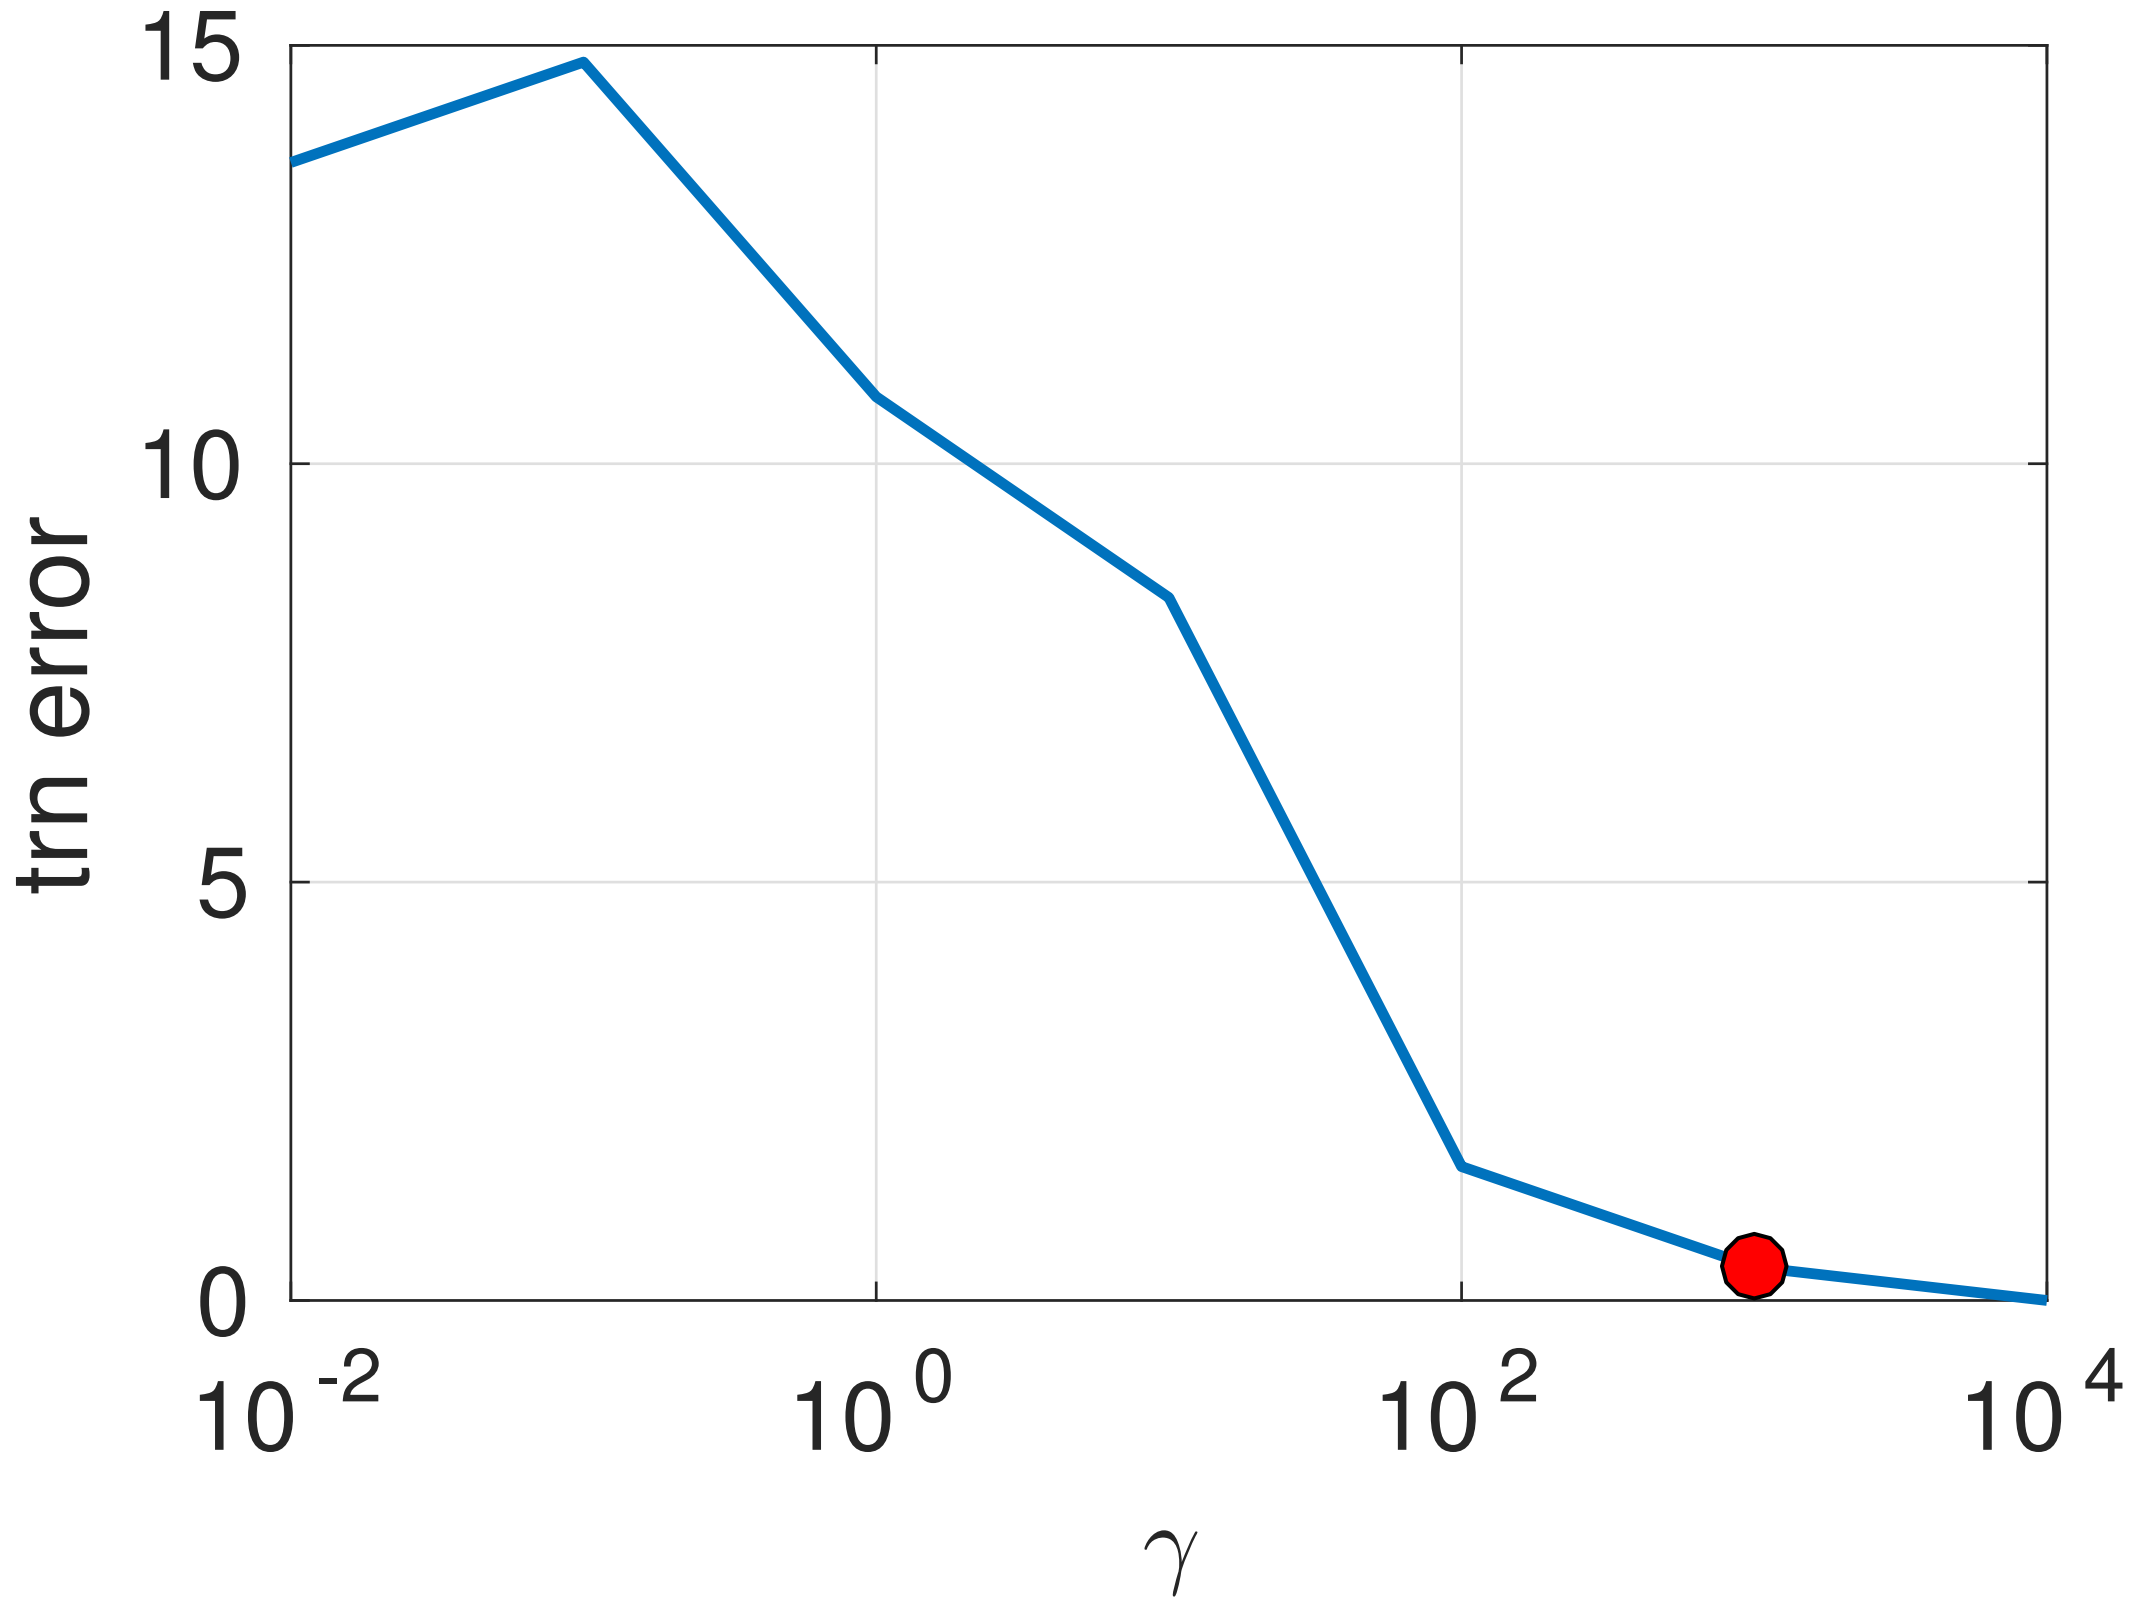


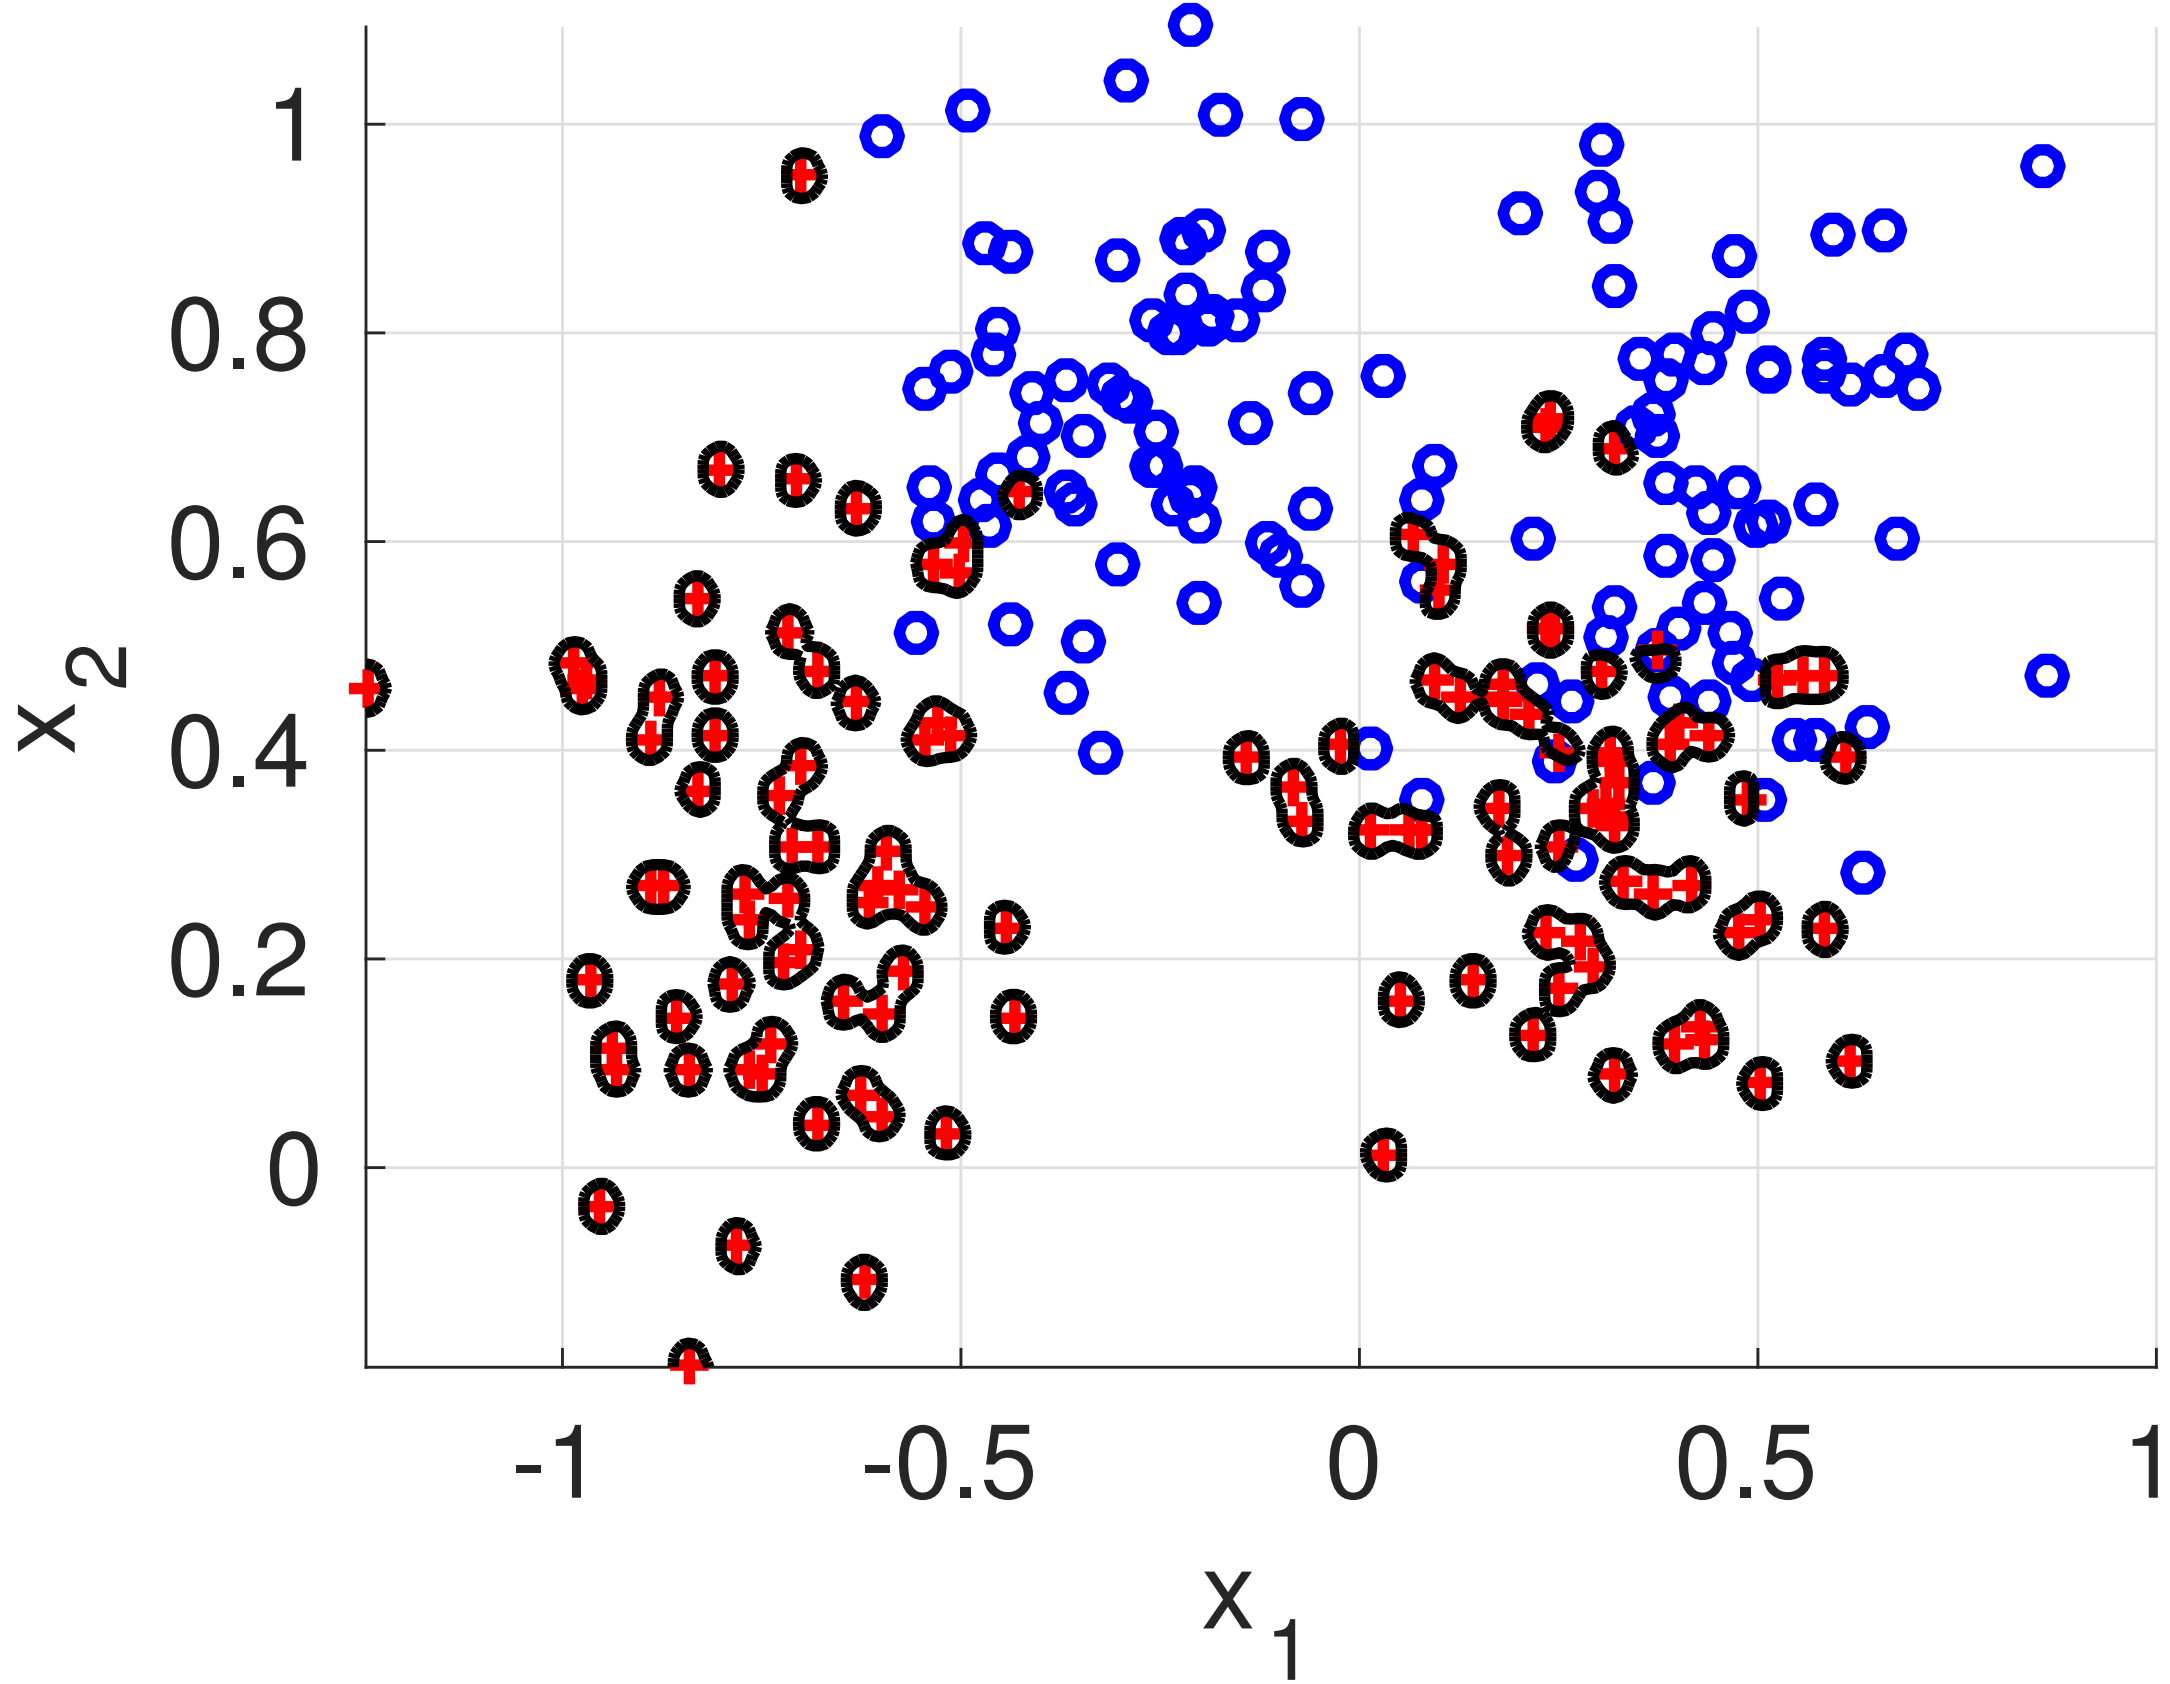
err=1.6%



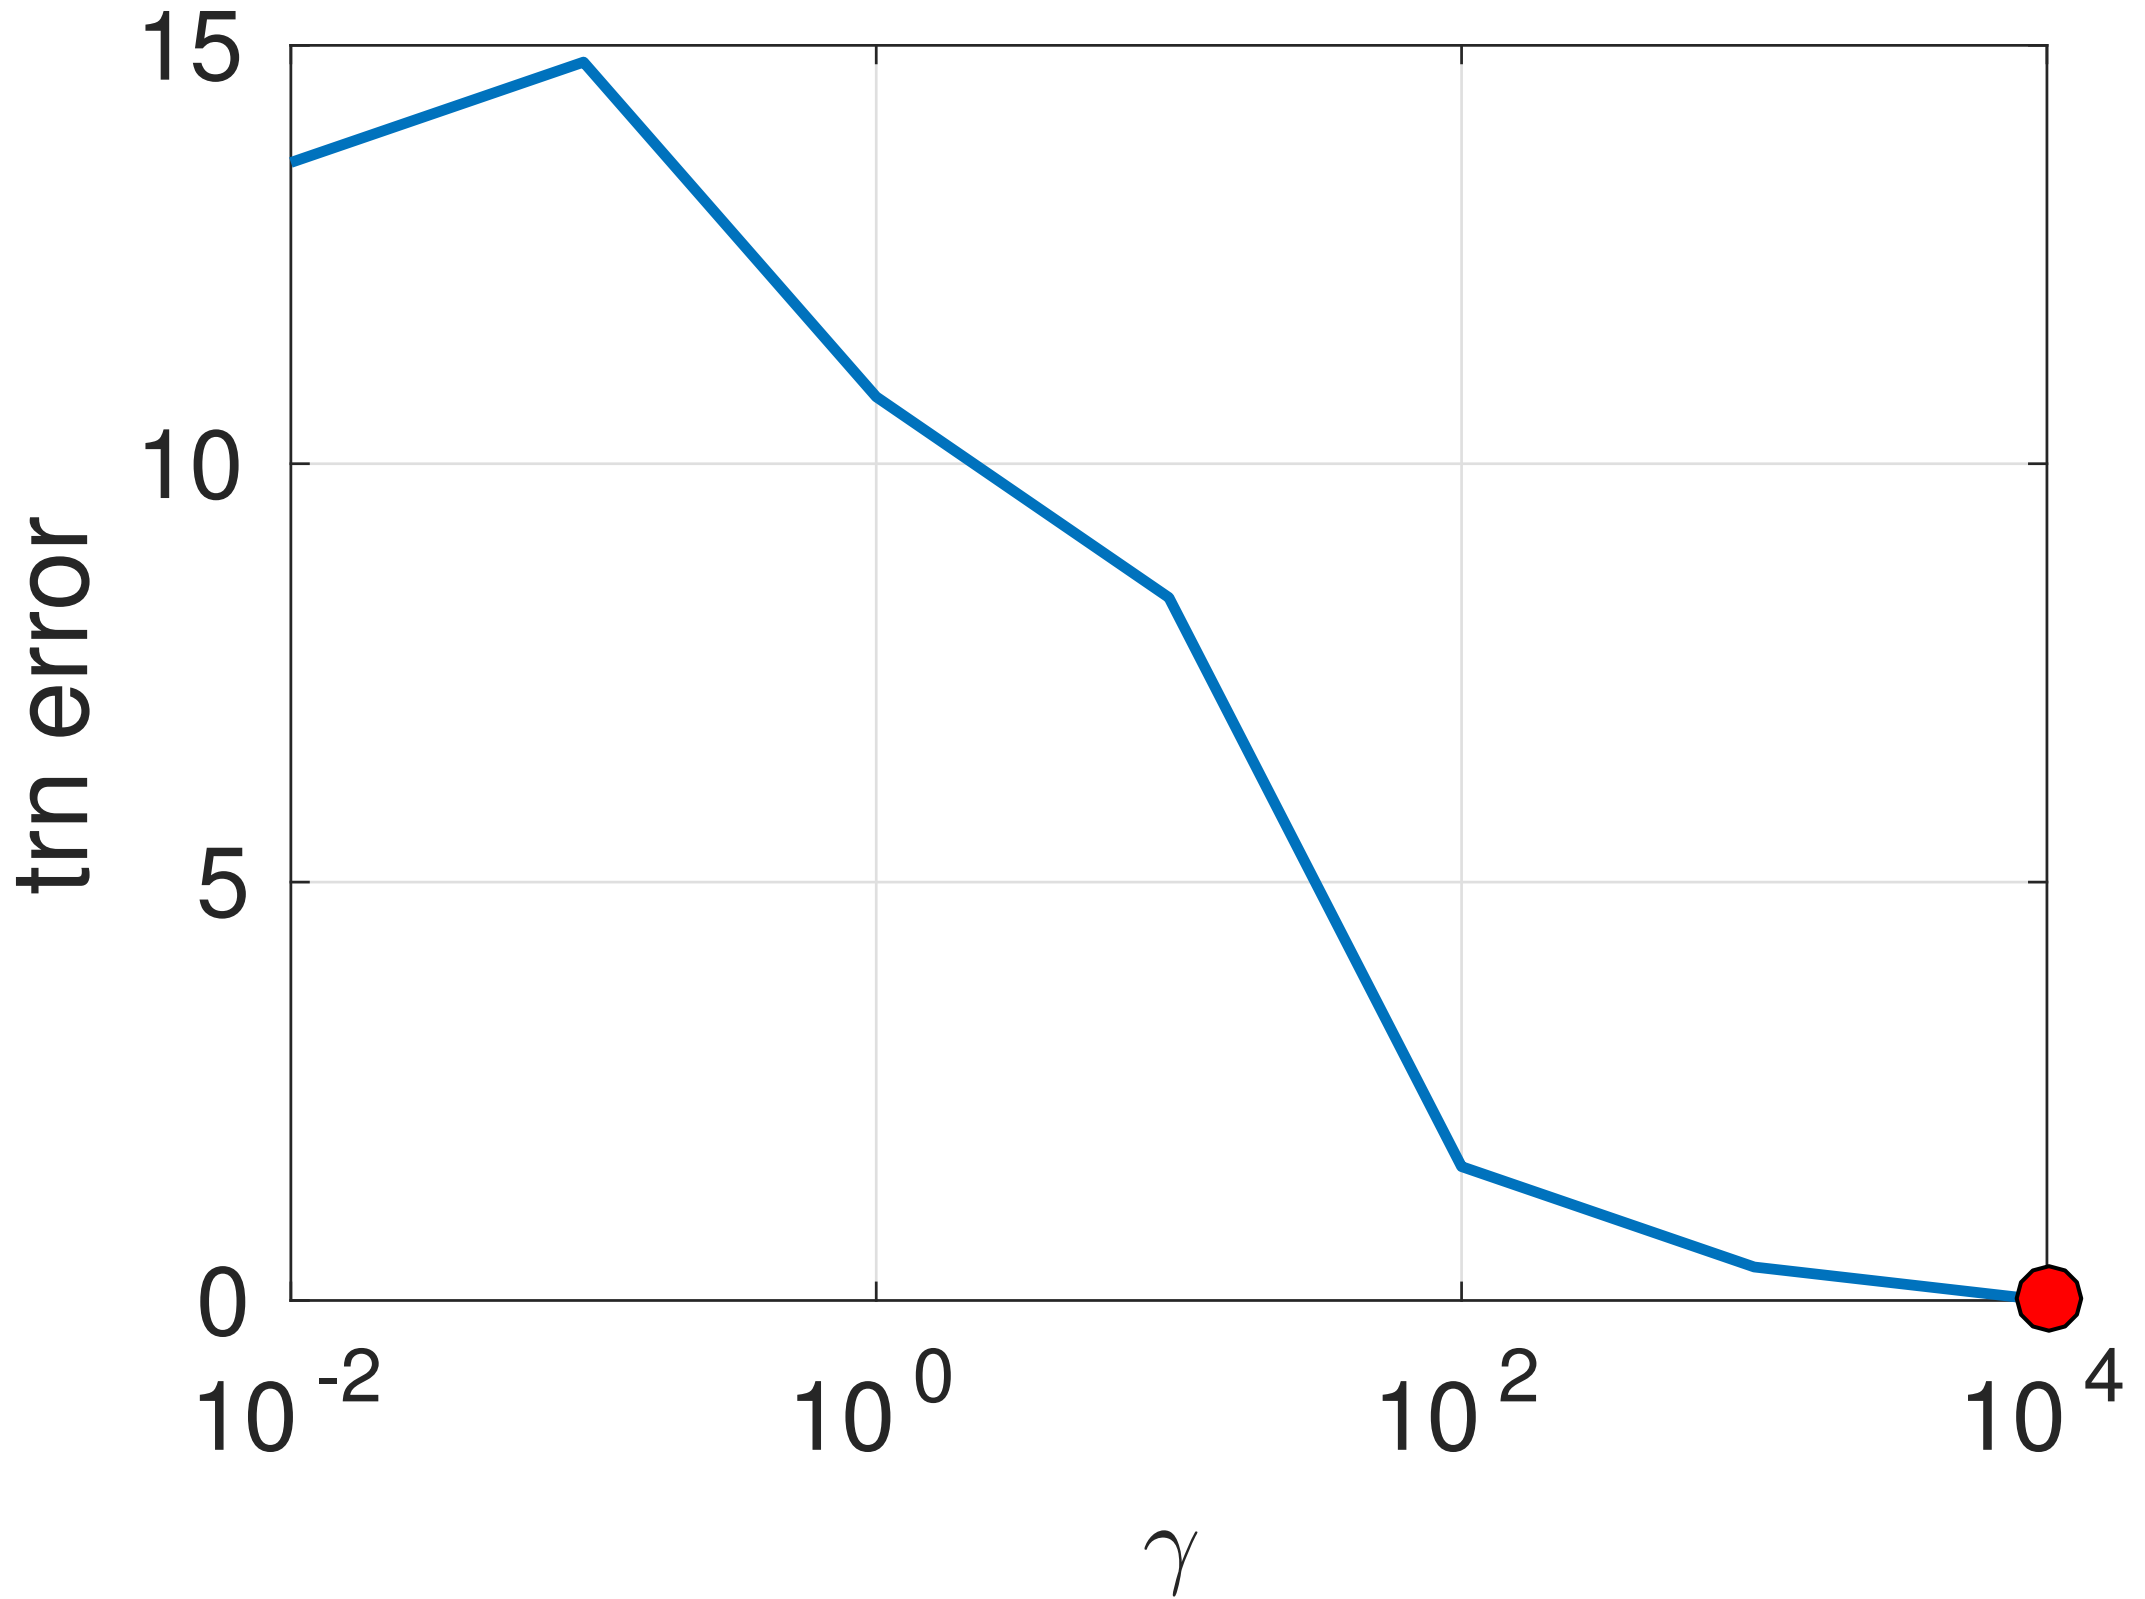


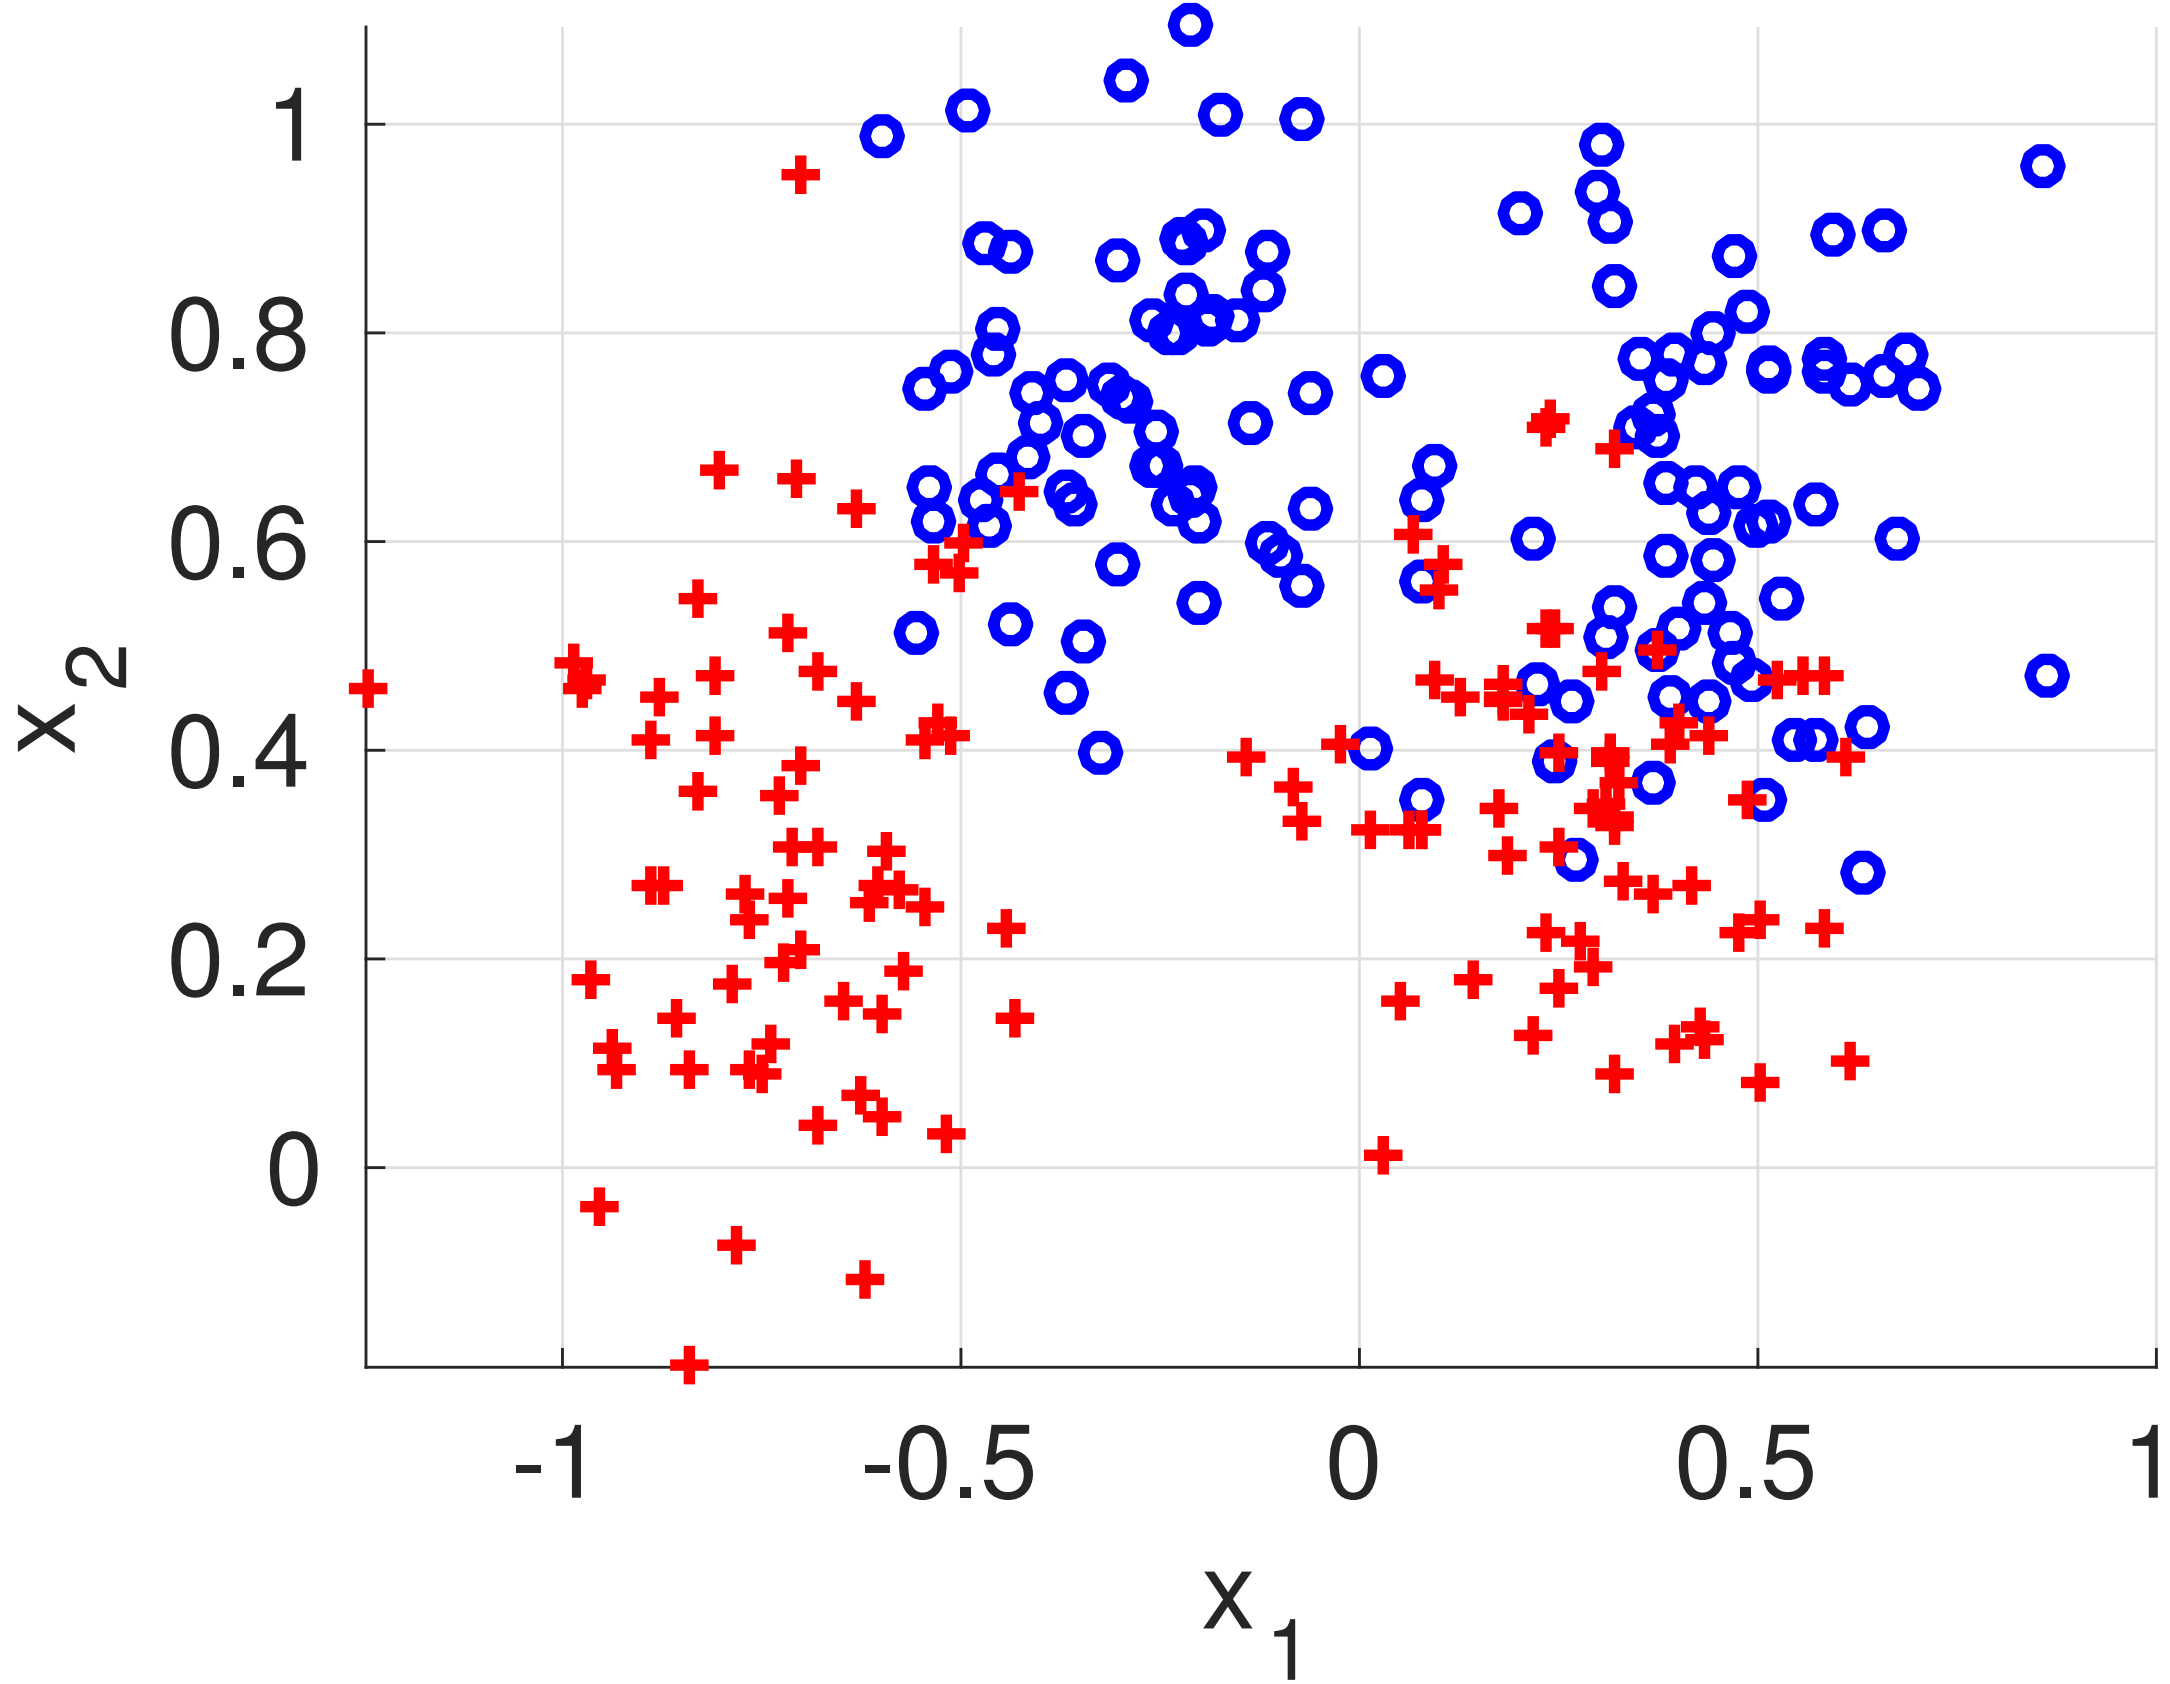
err=0.4%

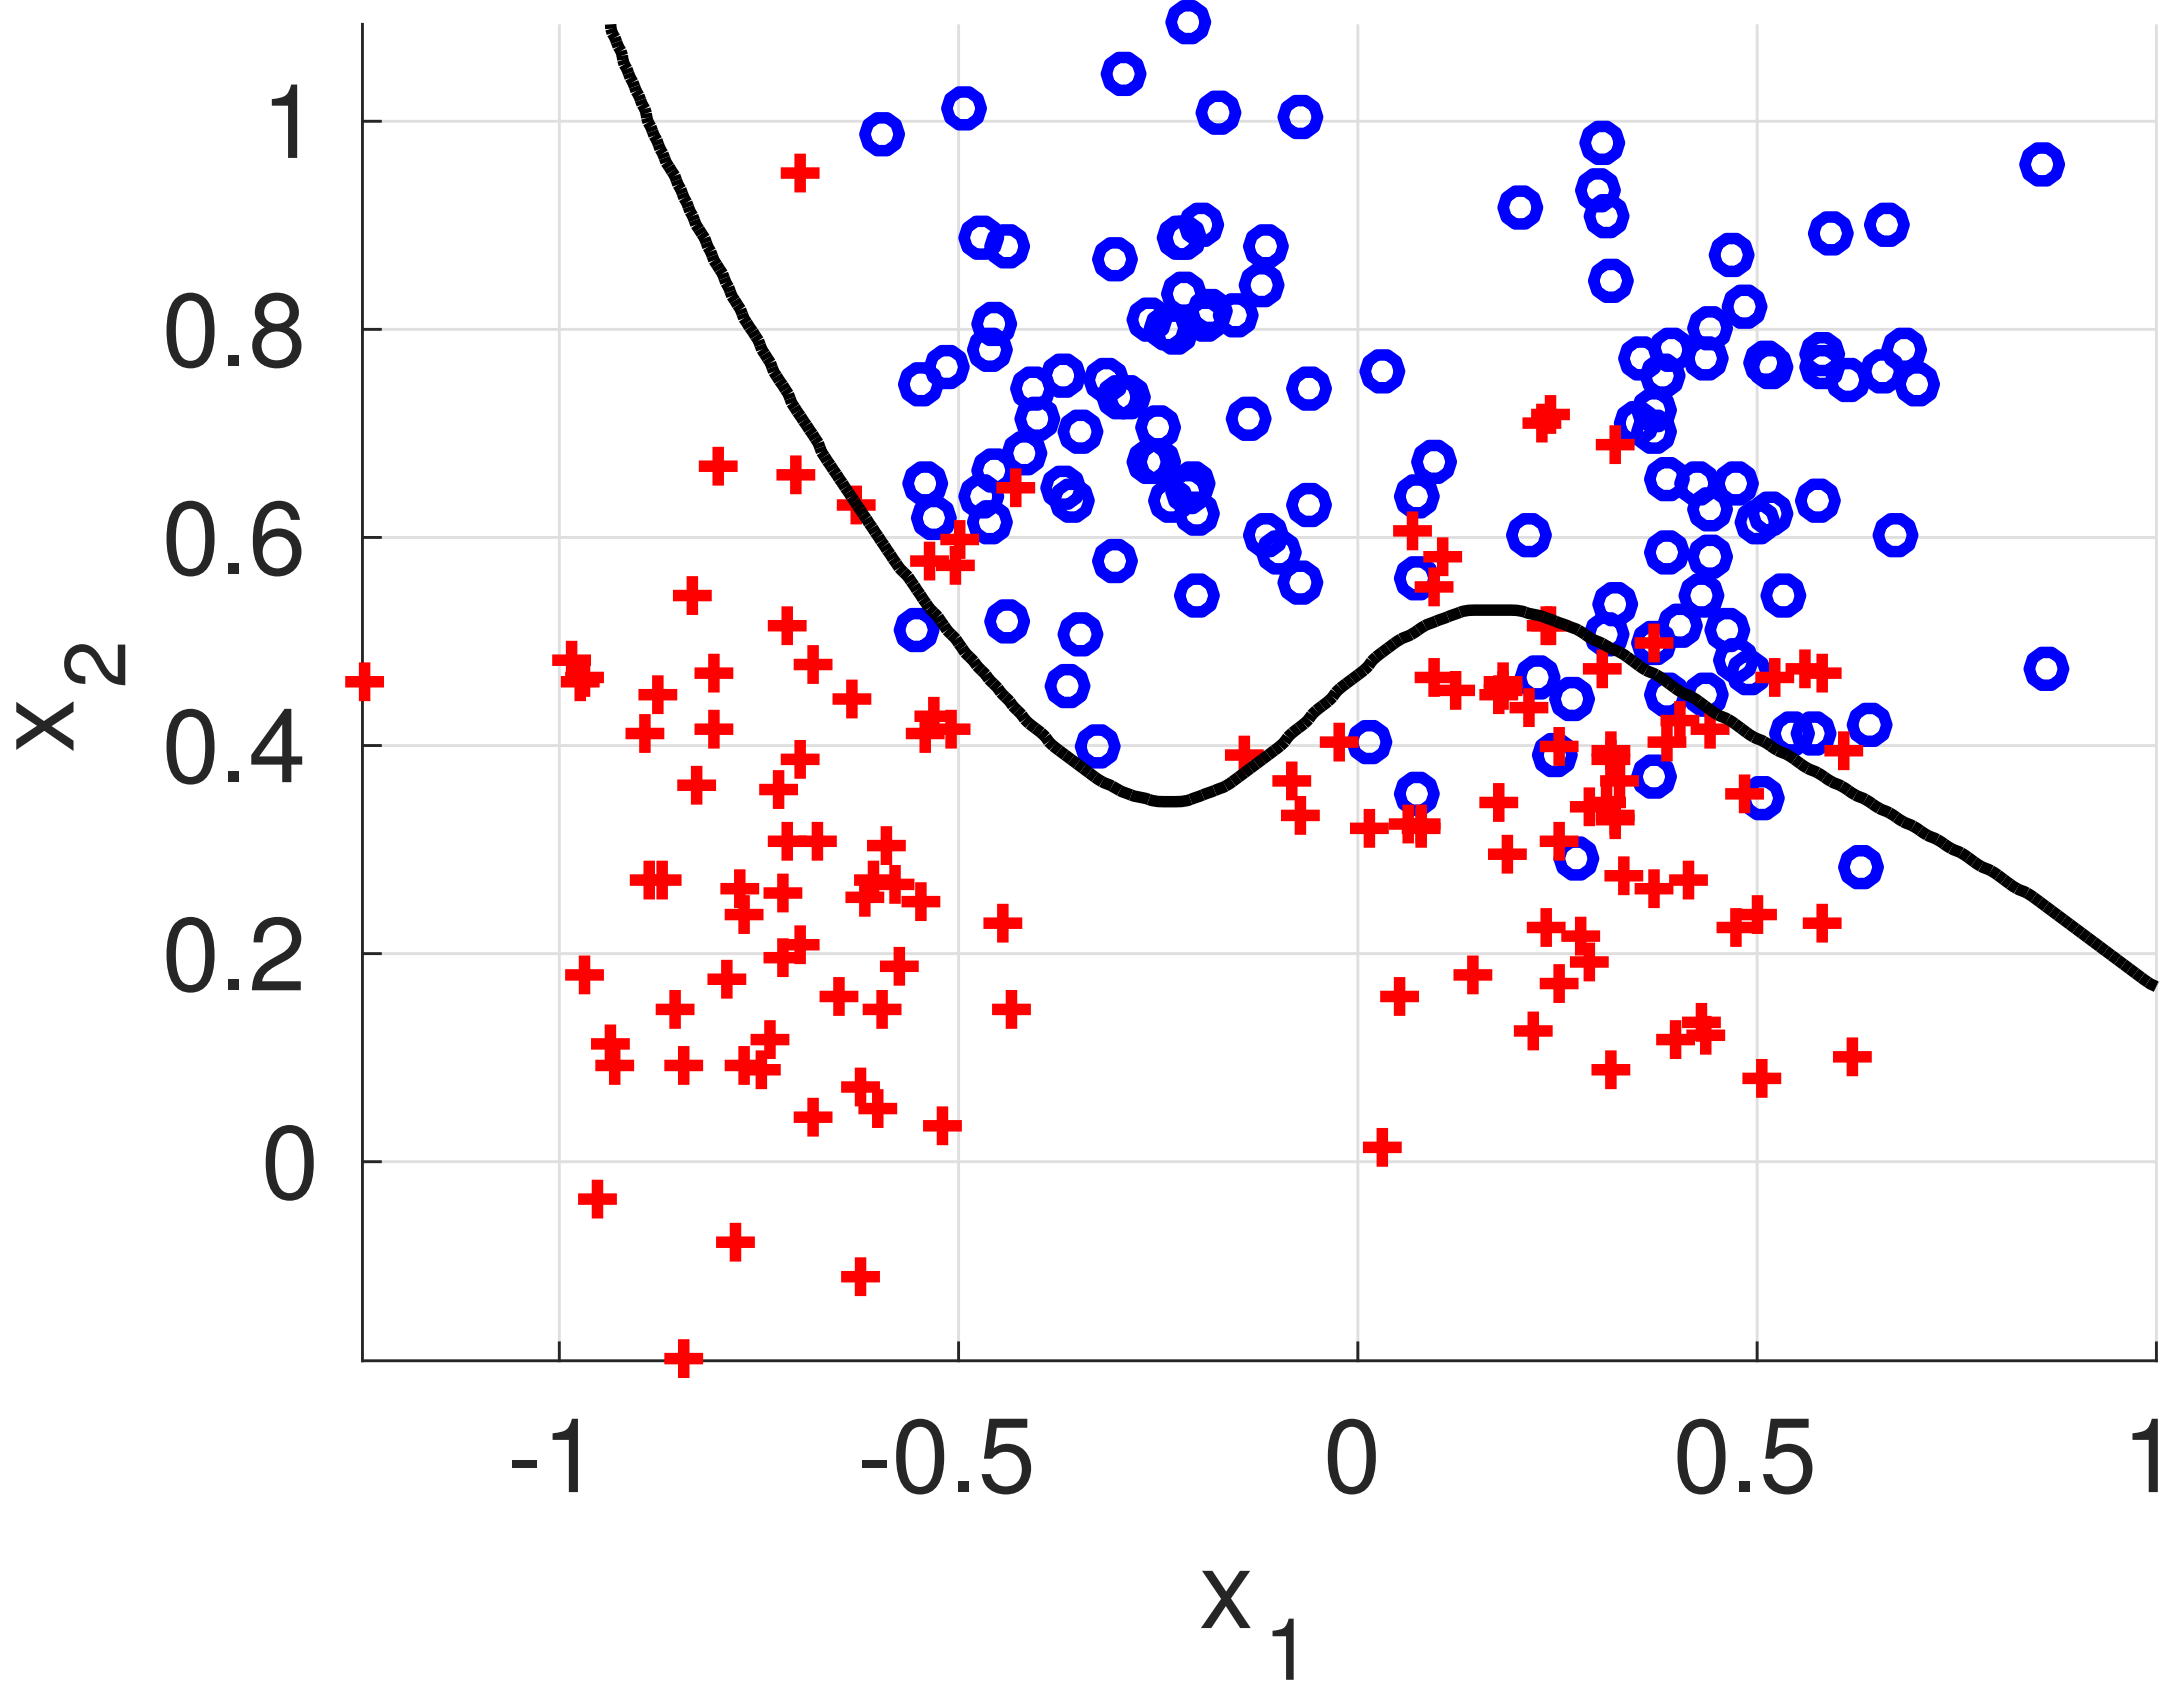




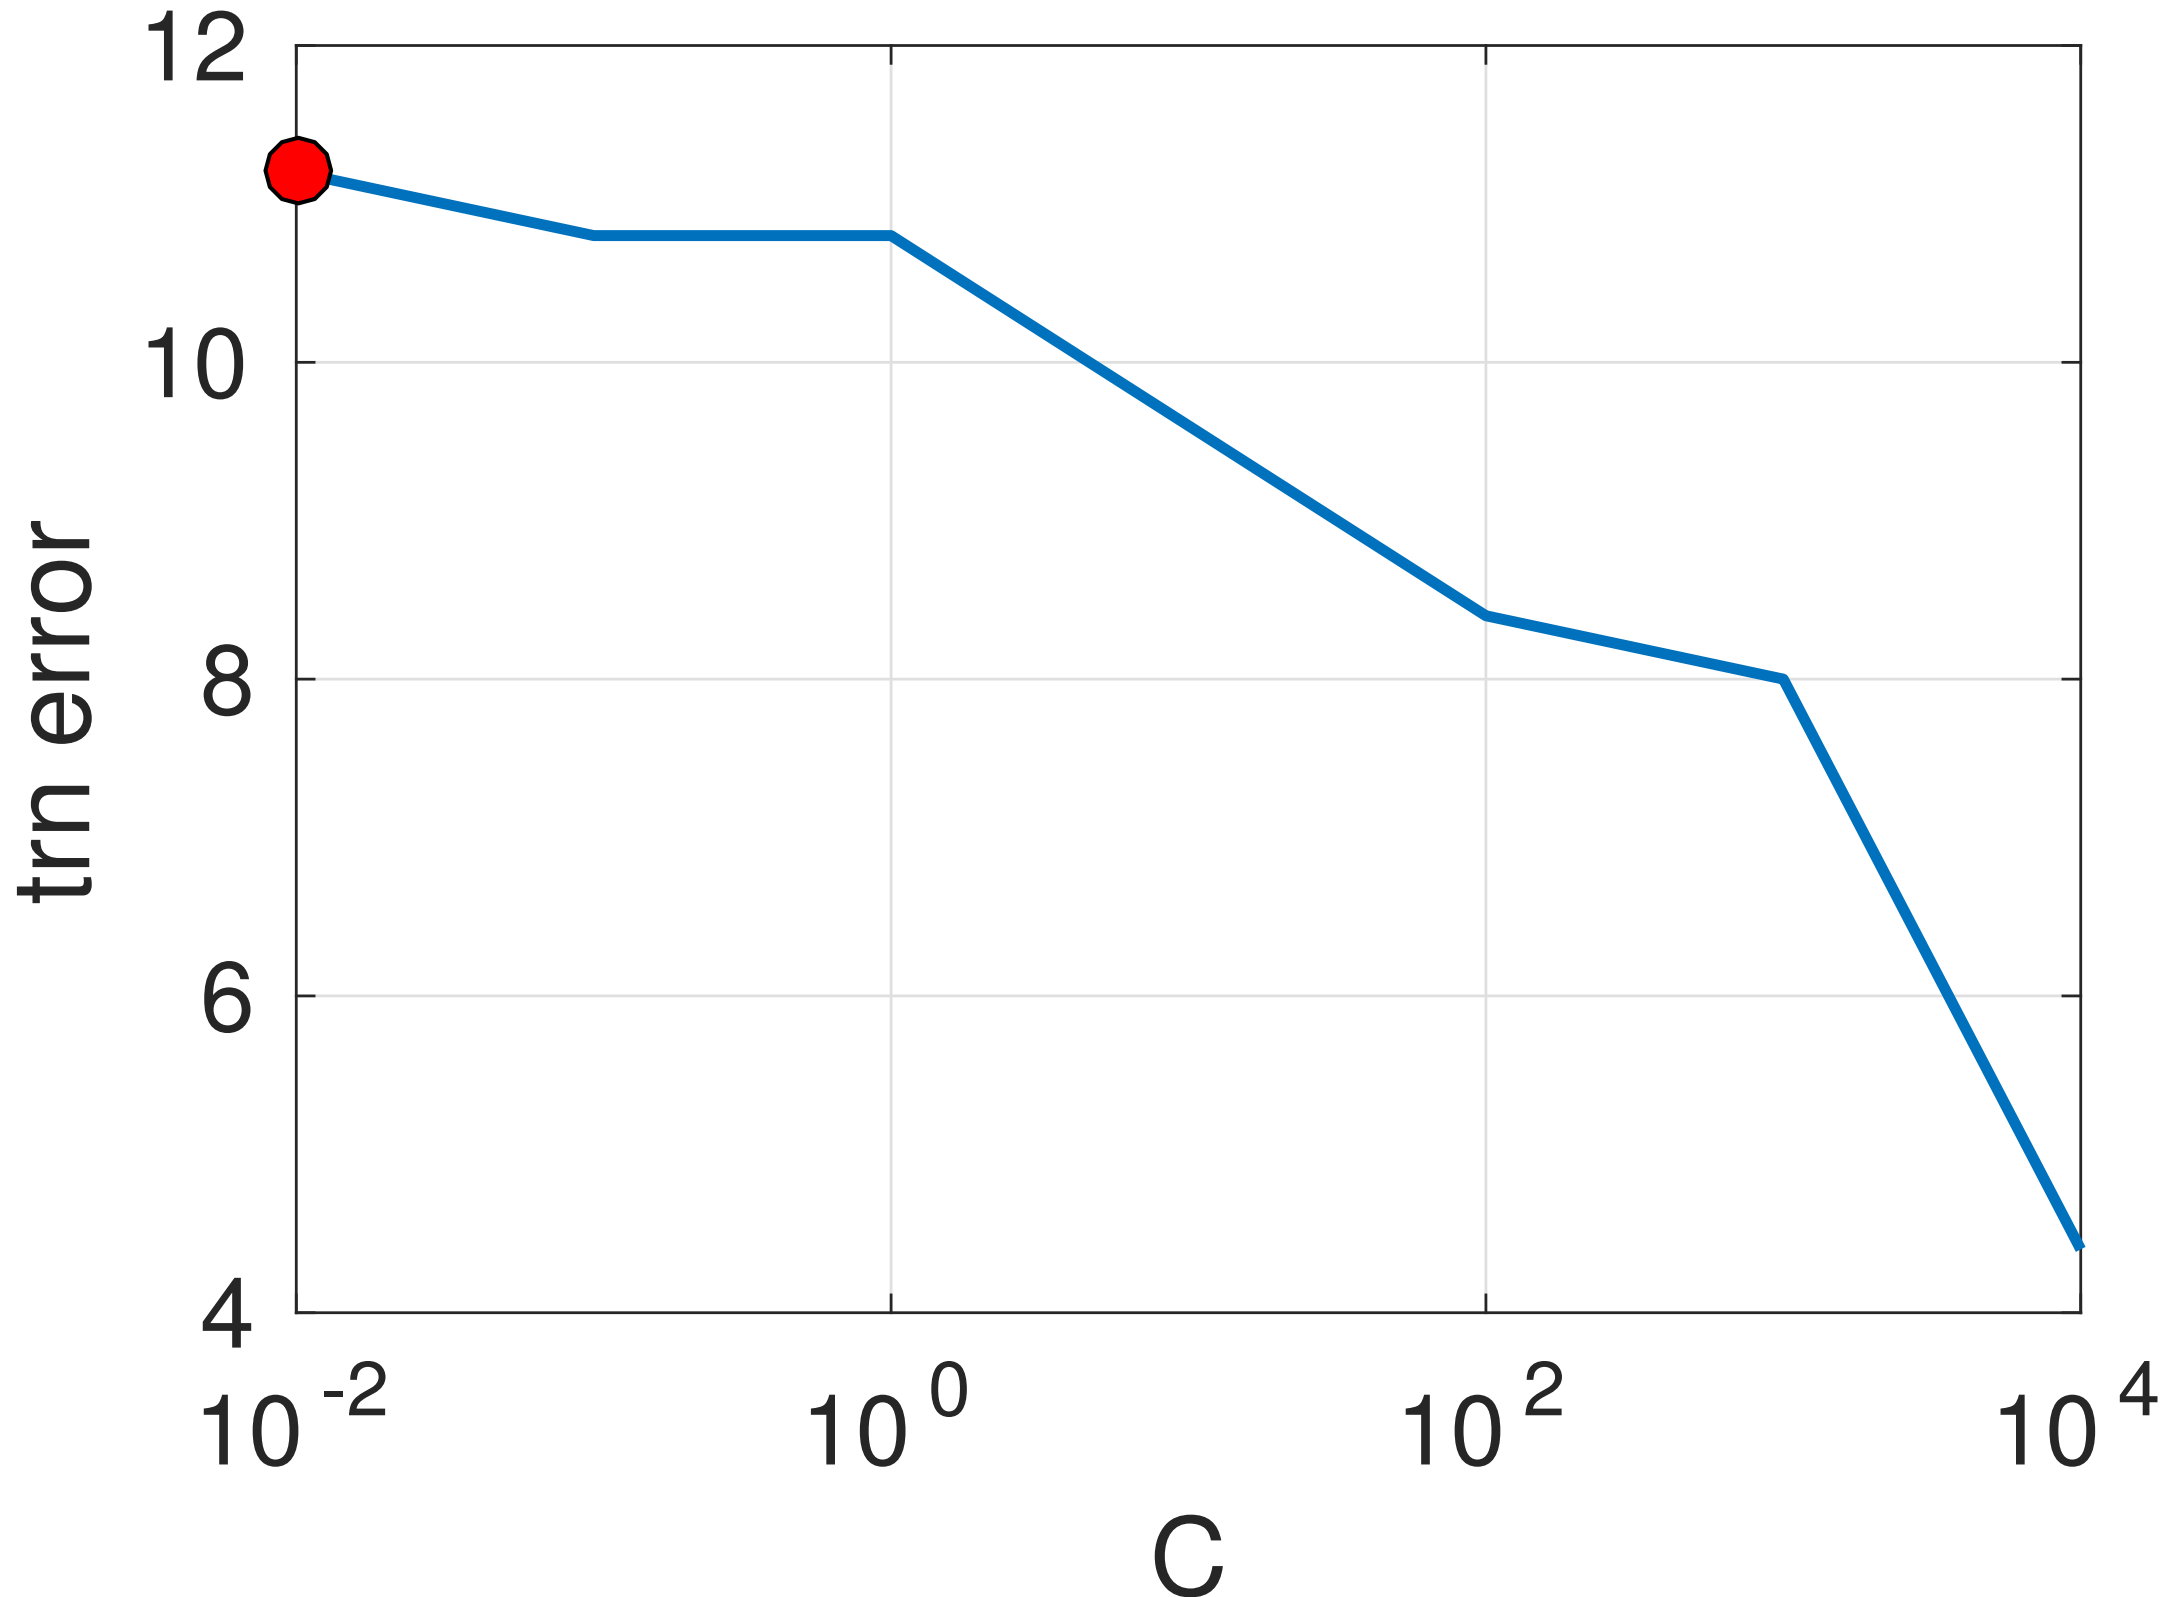
err=0.0%

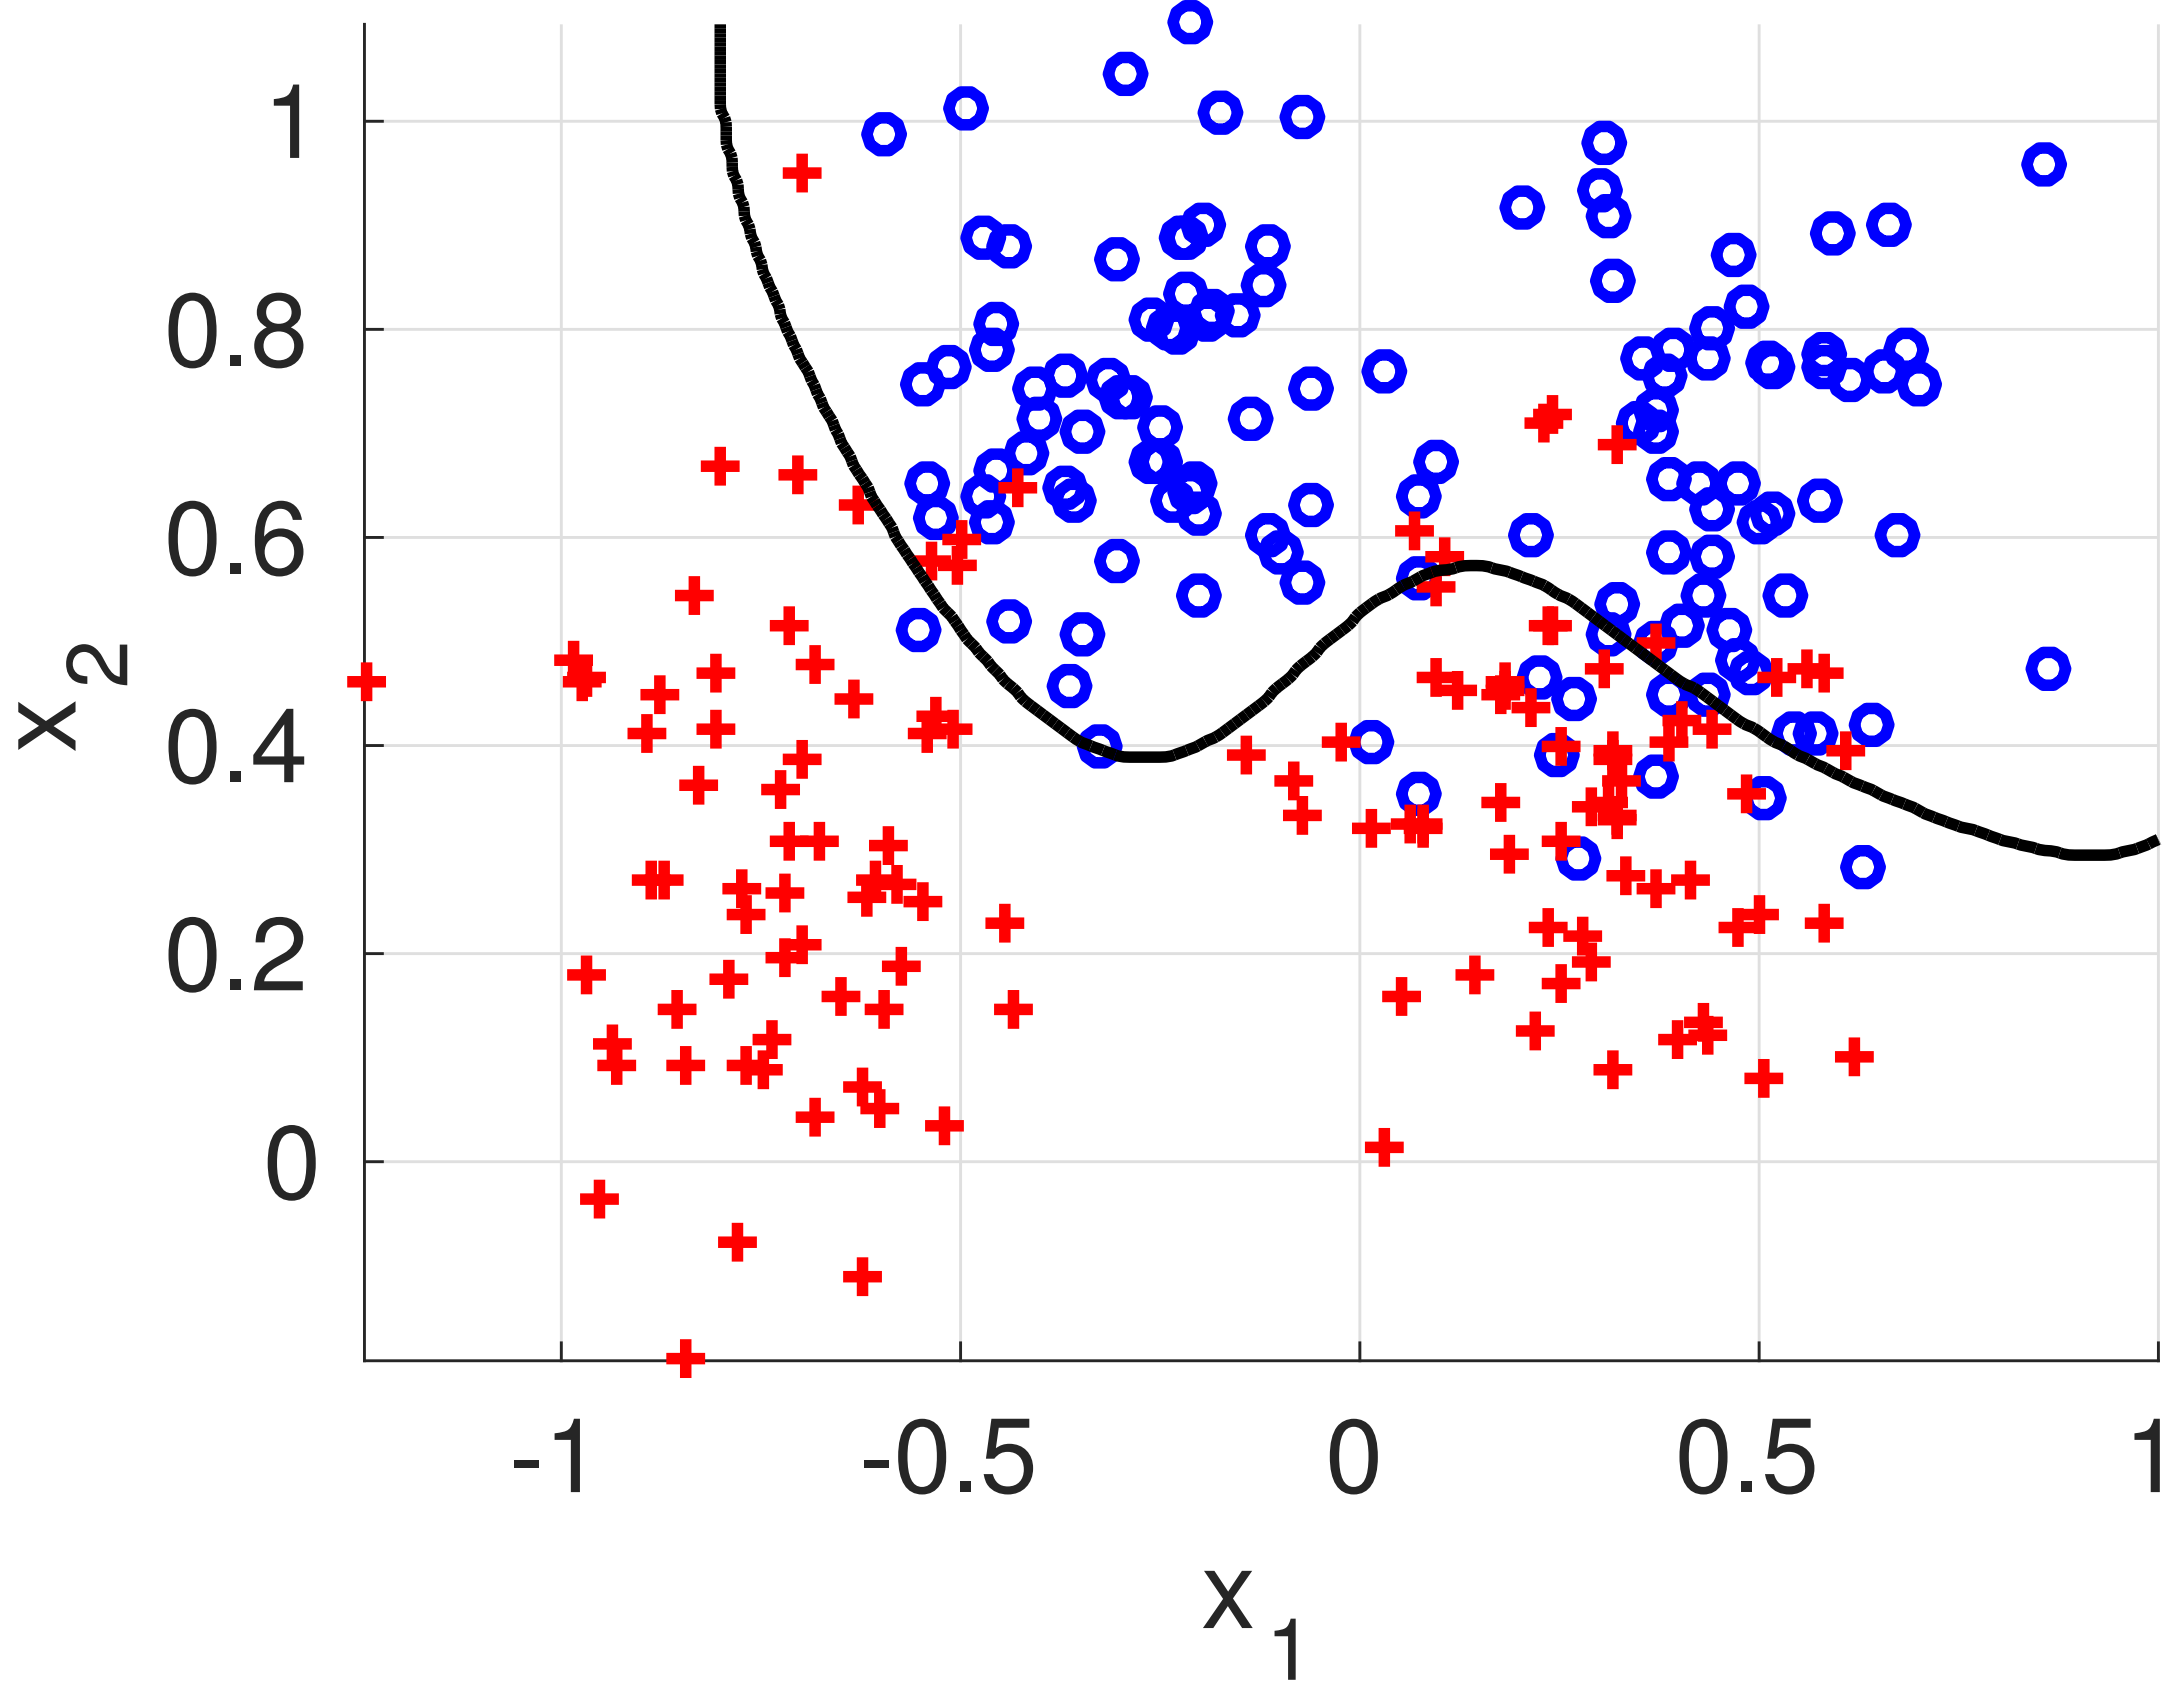




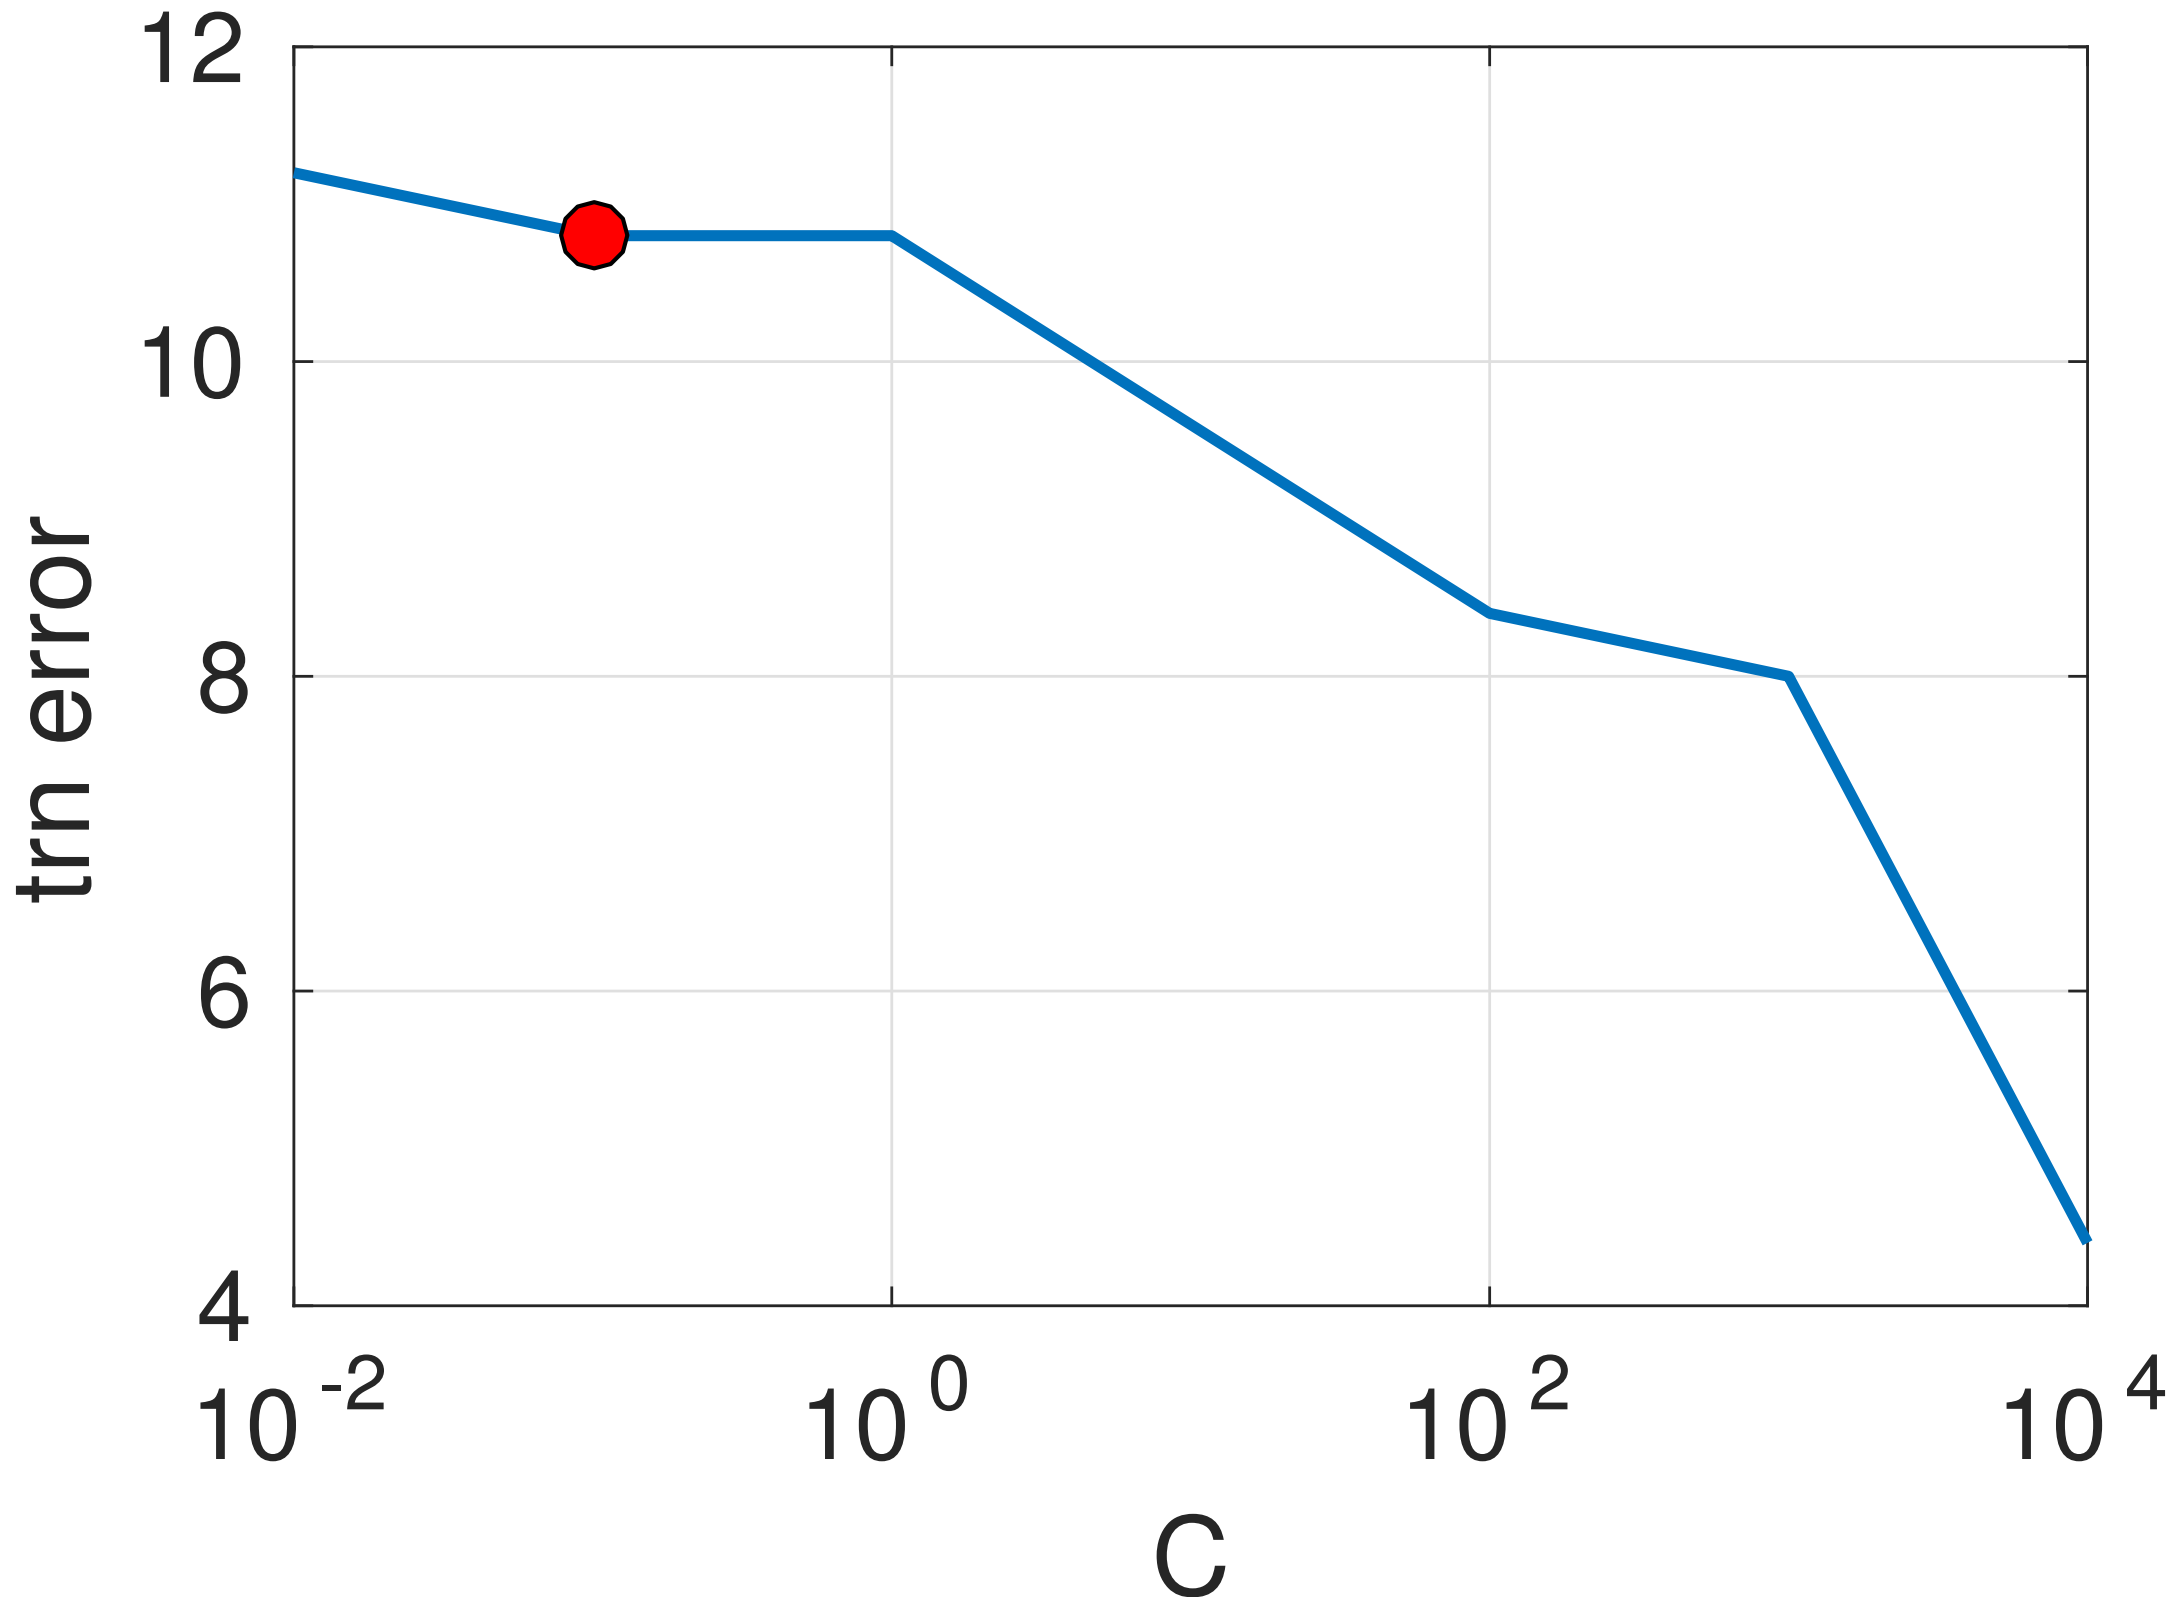


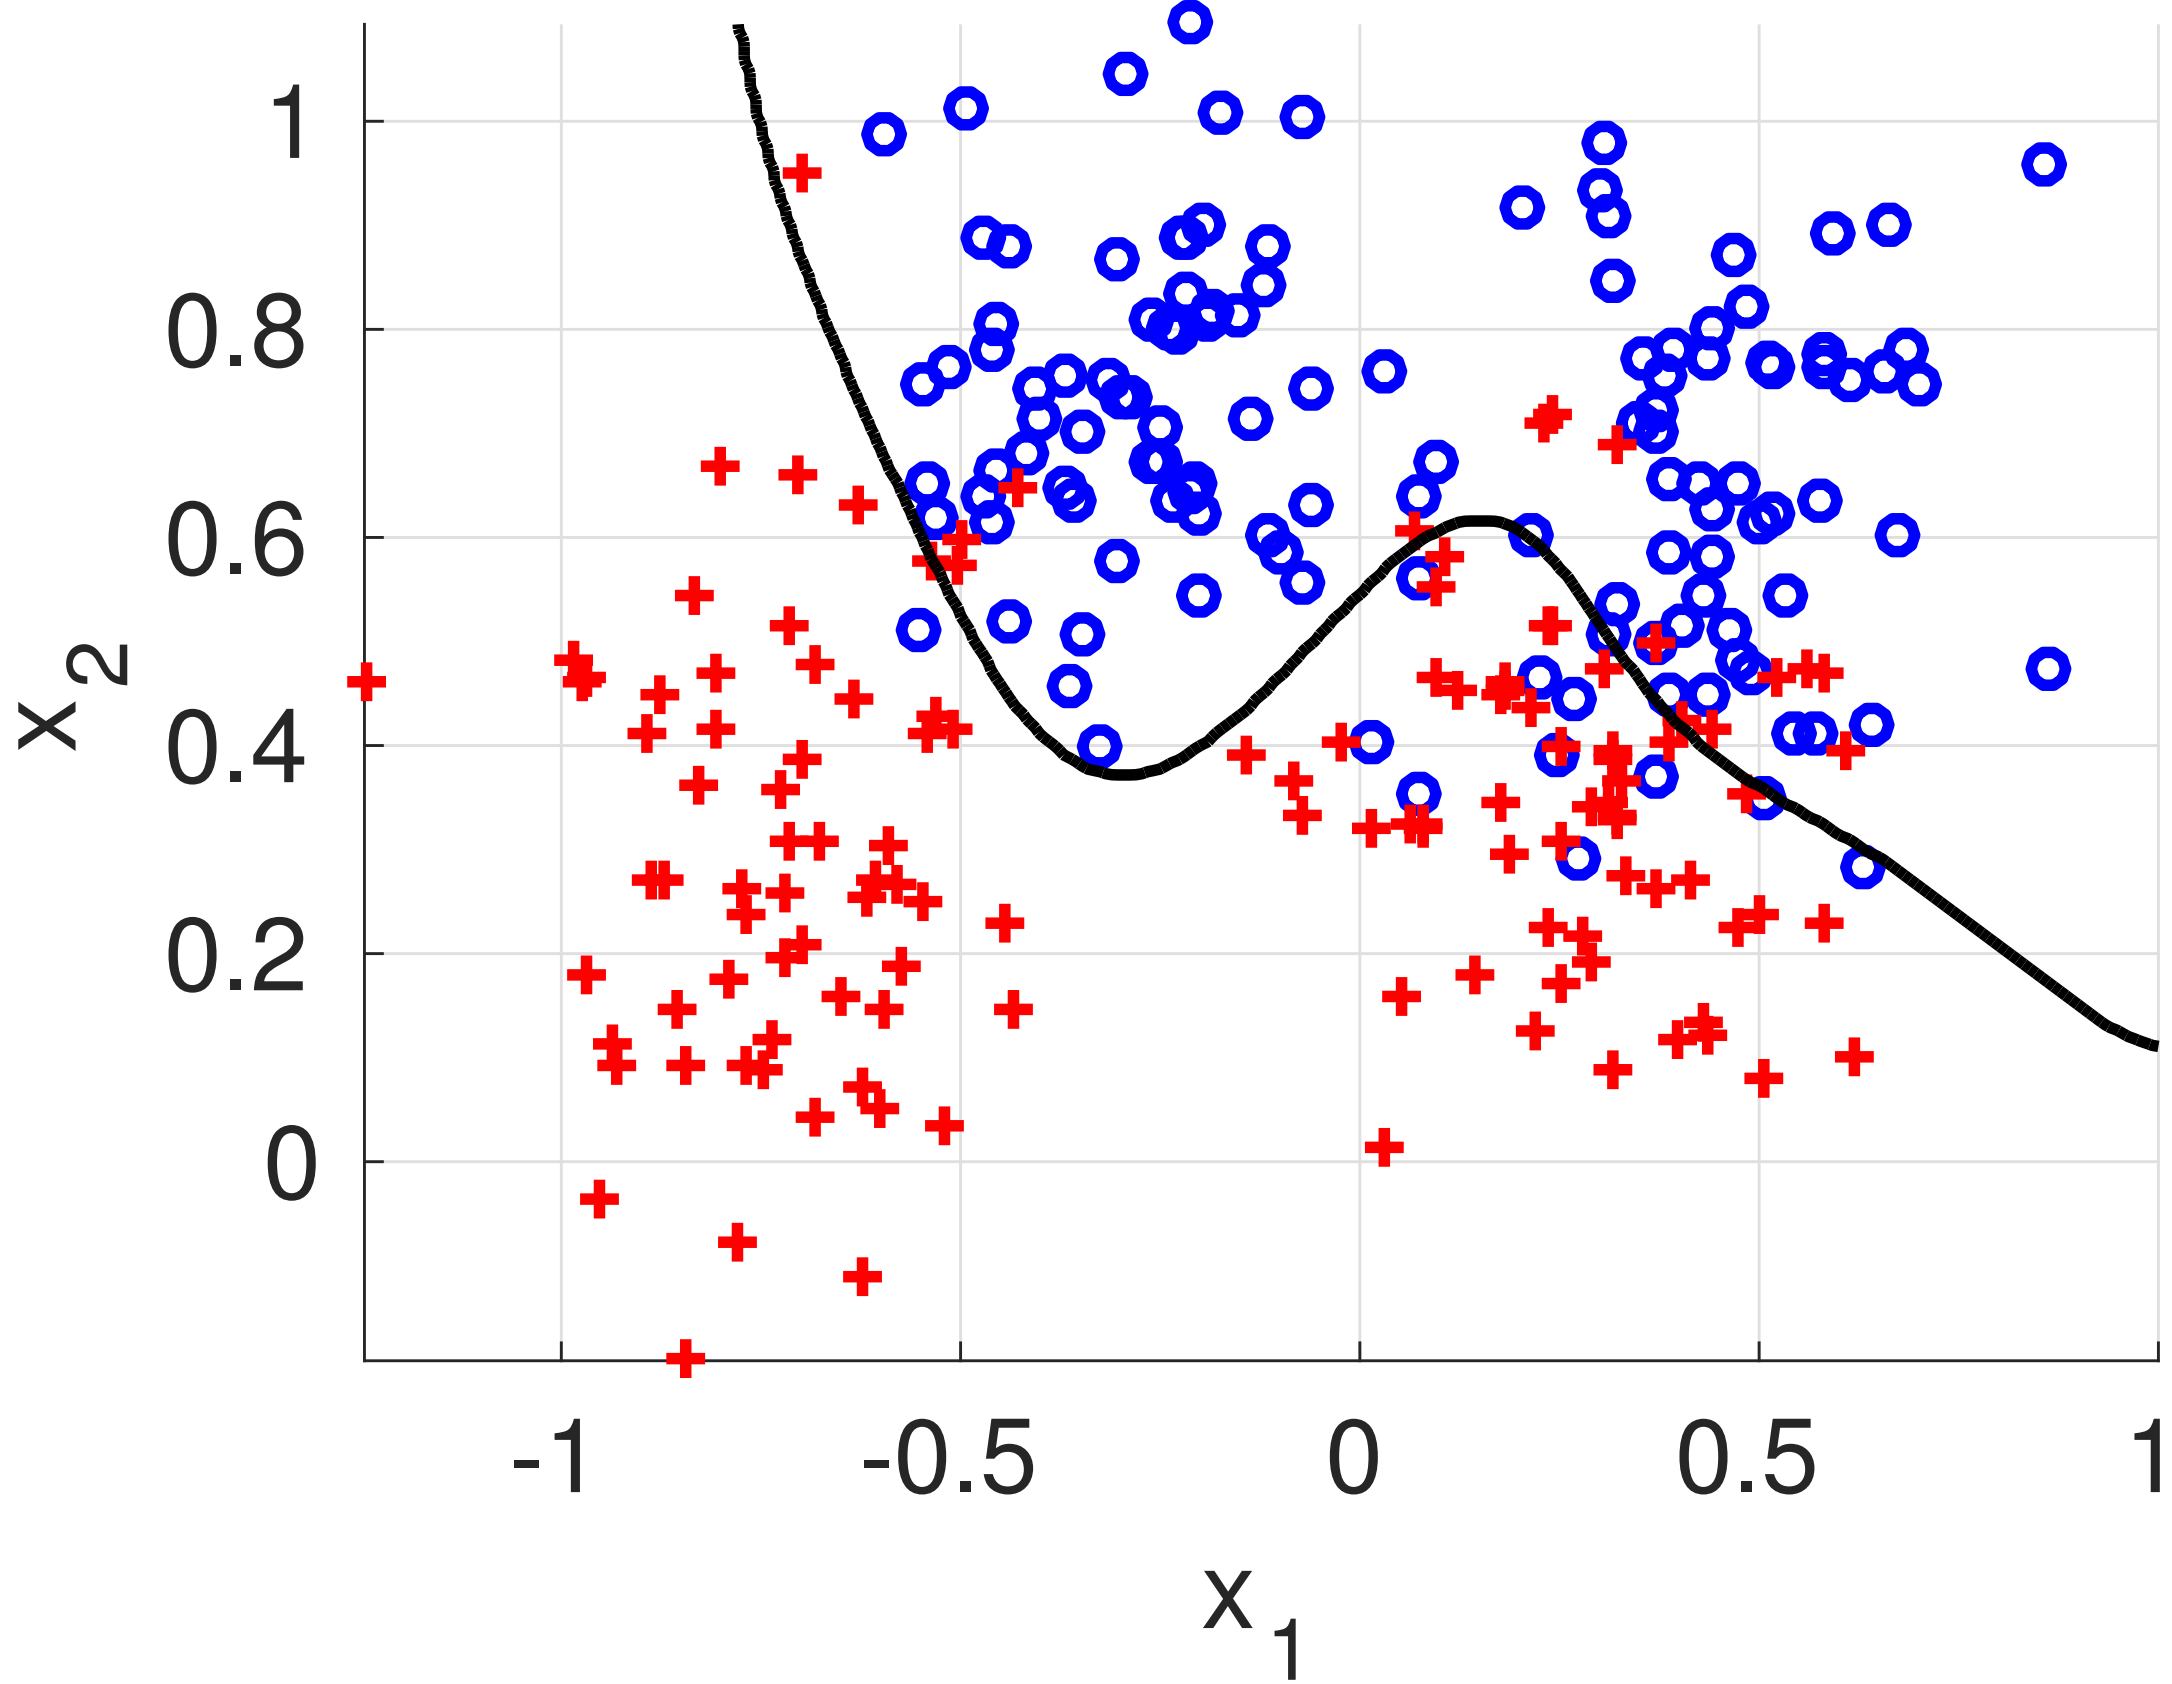
err=11.2%



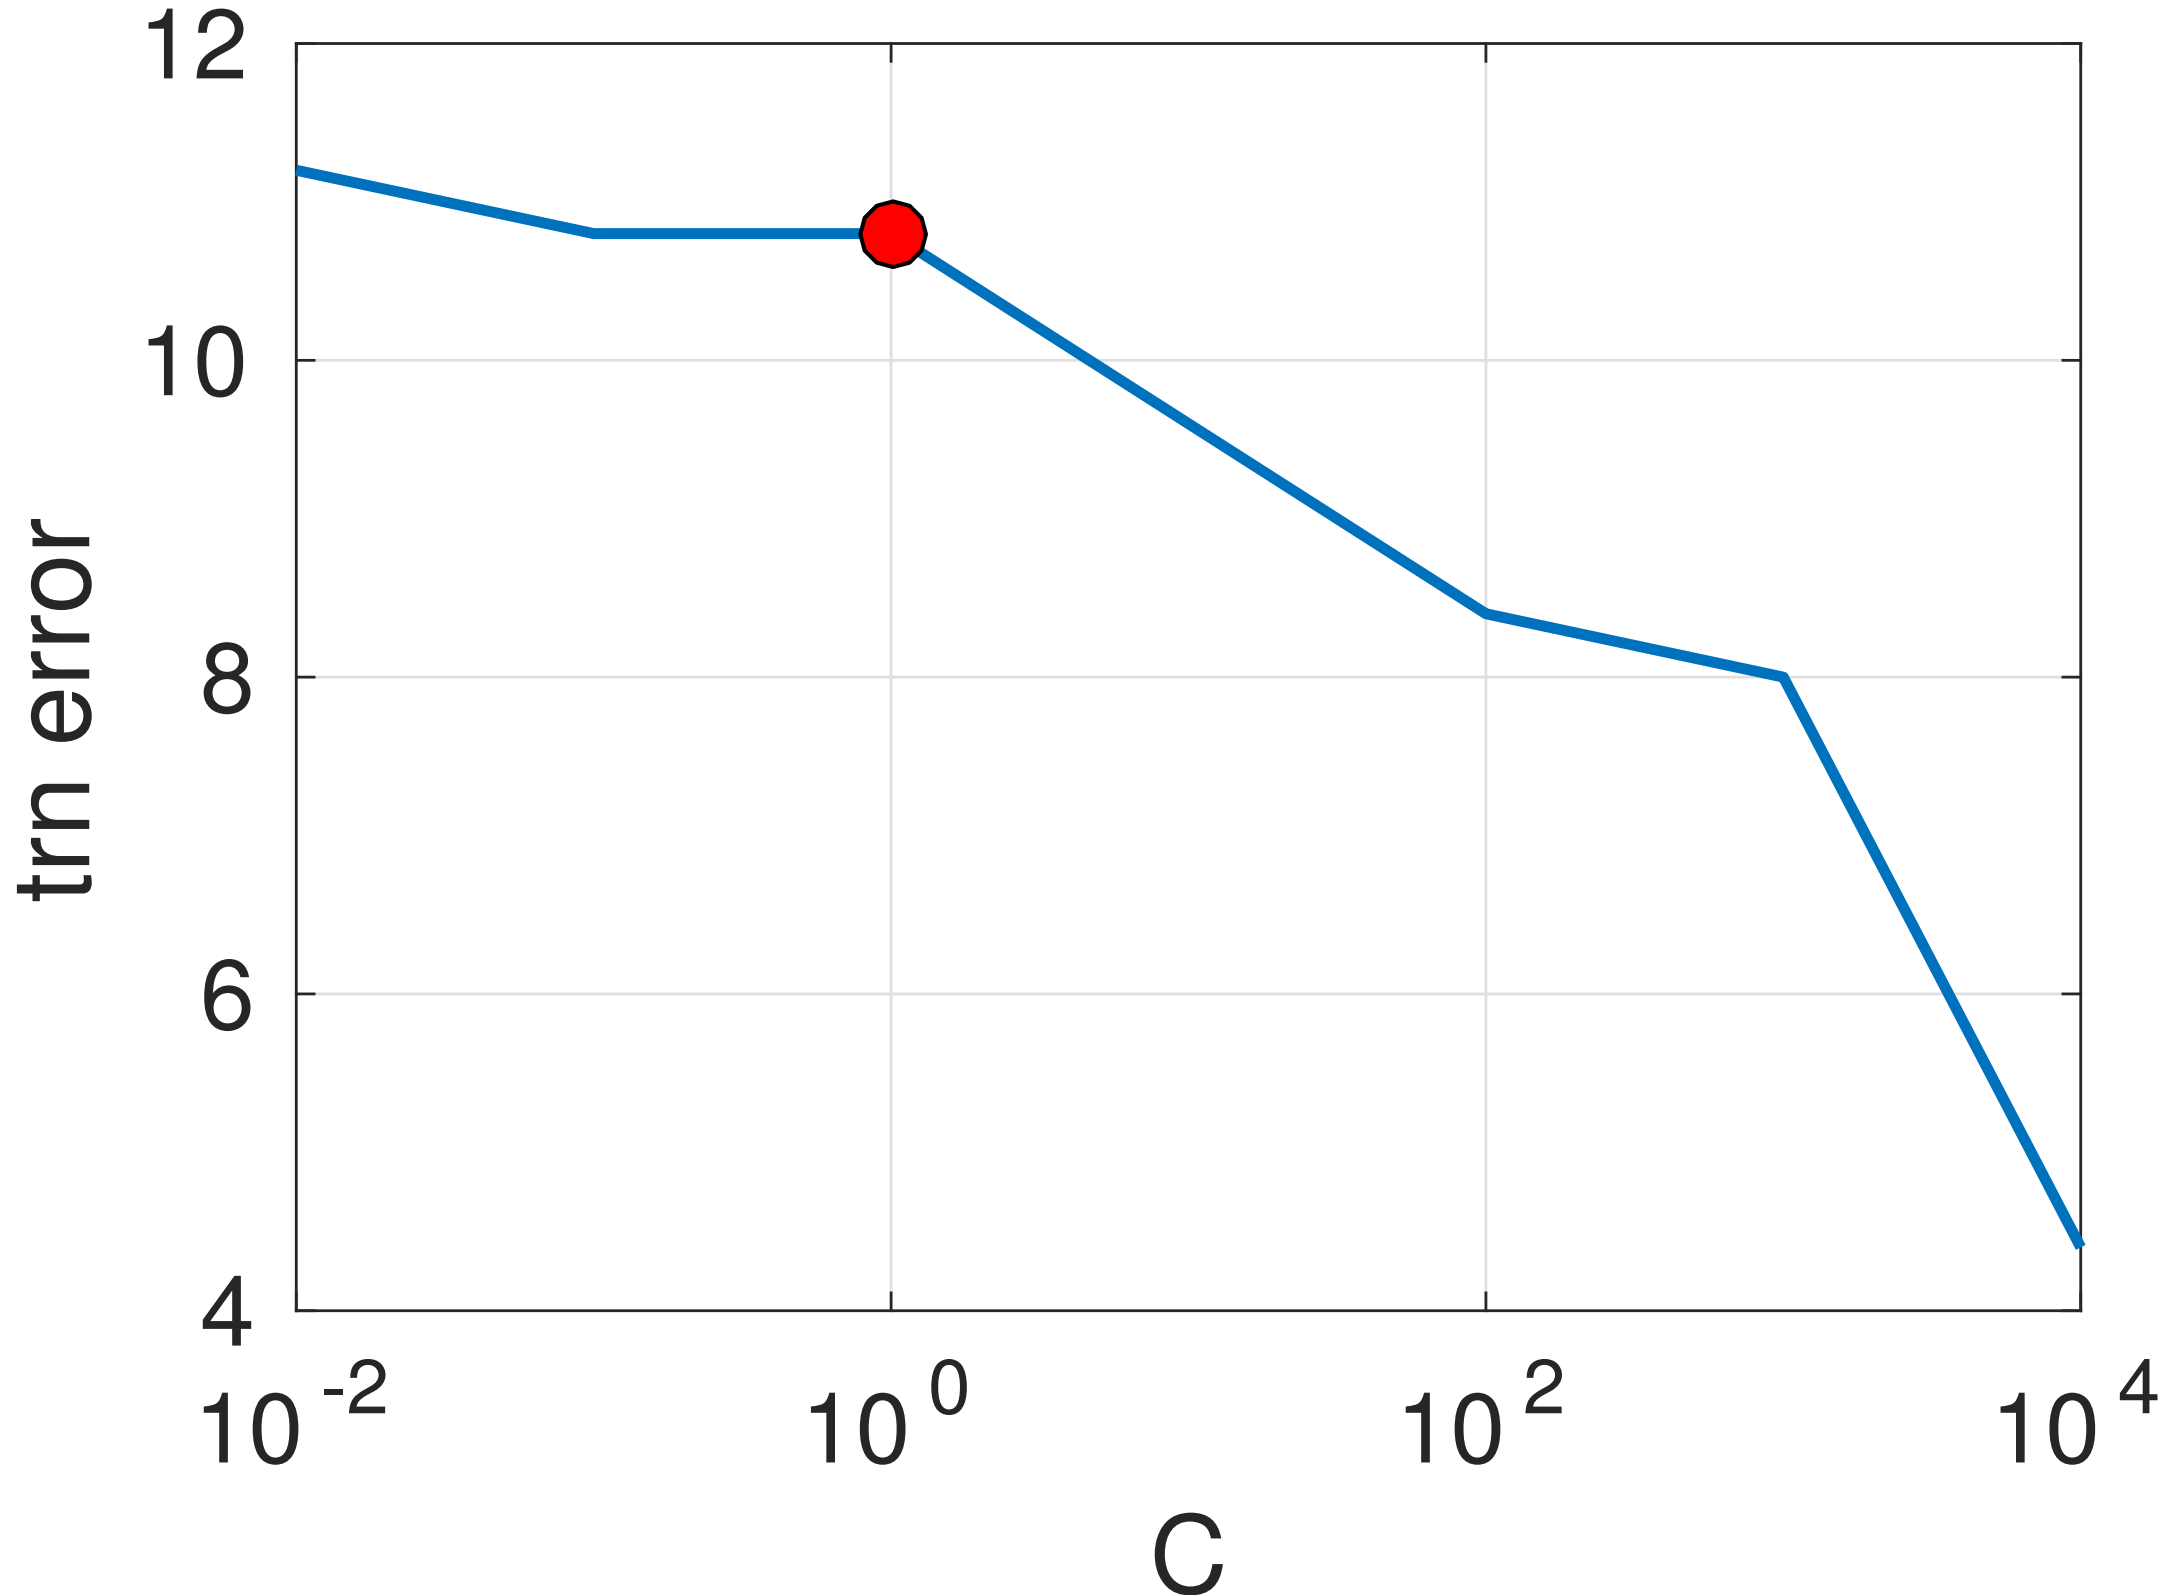


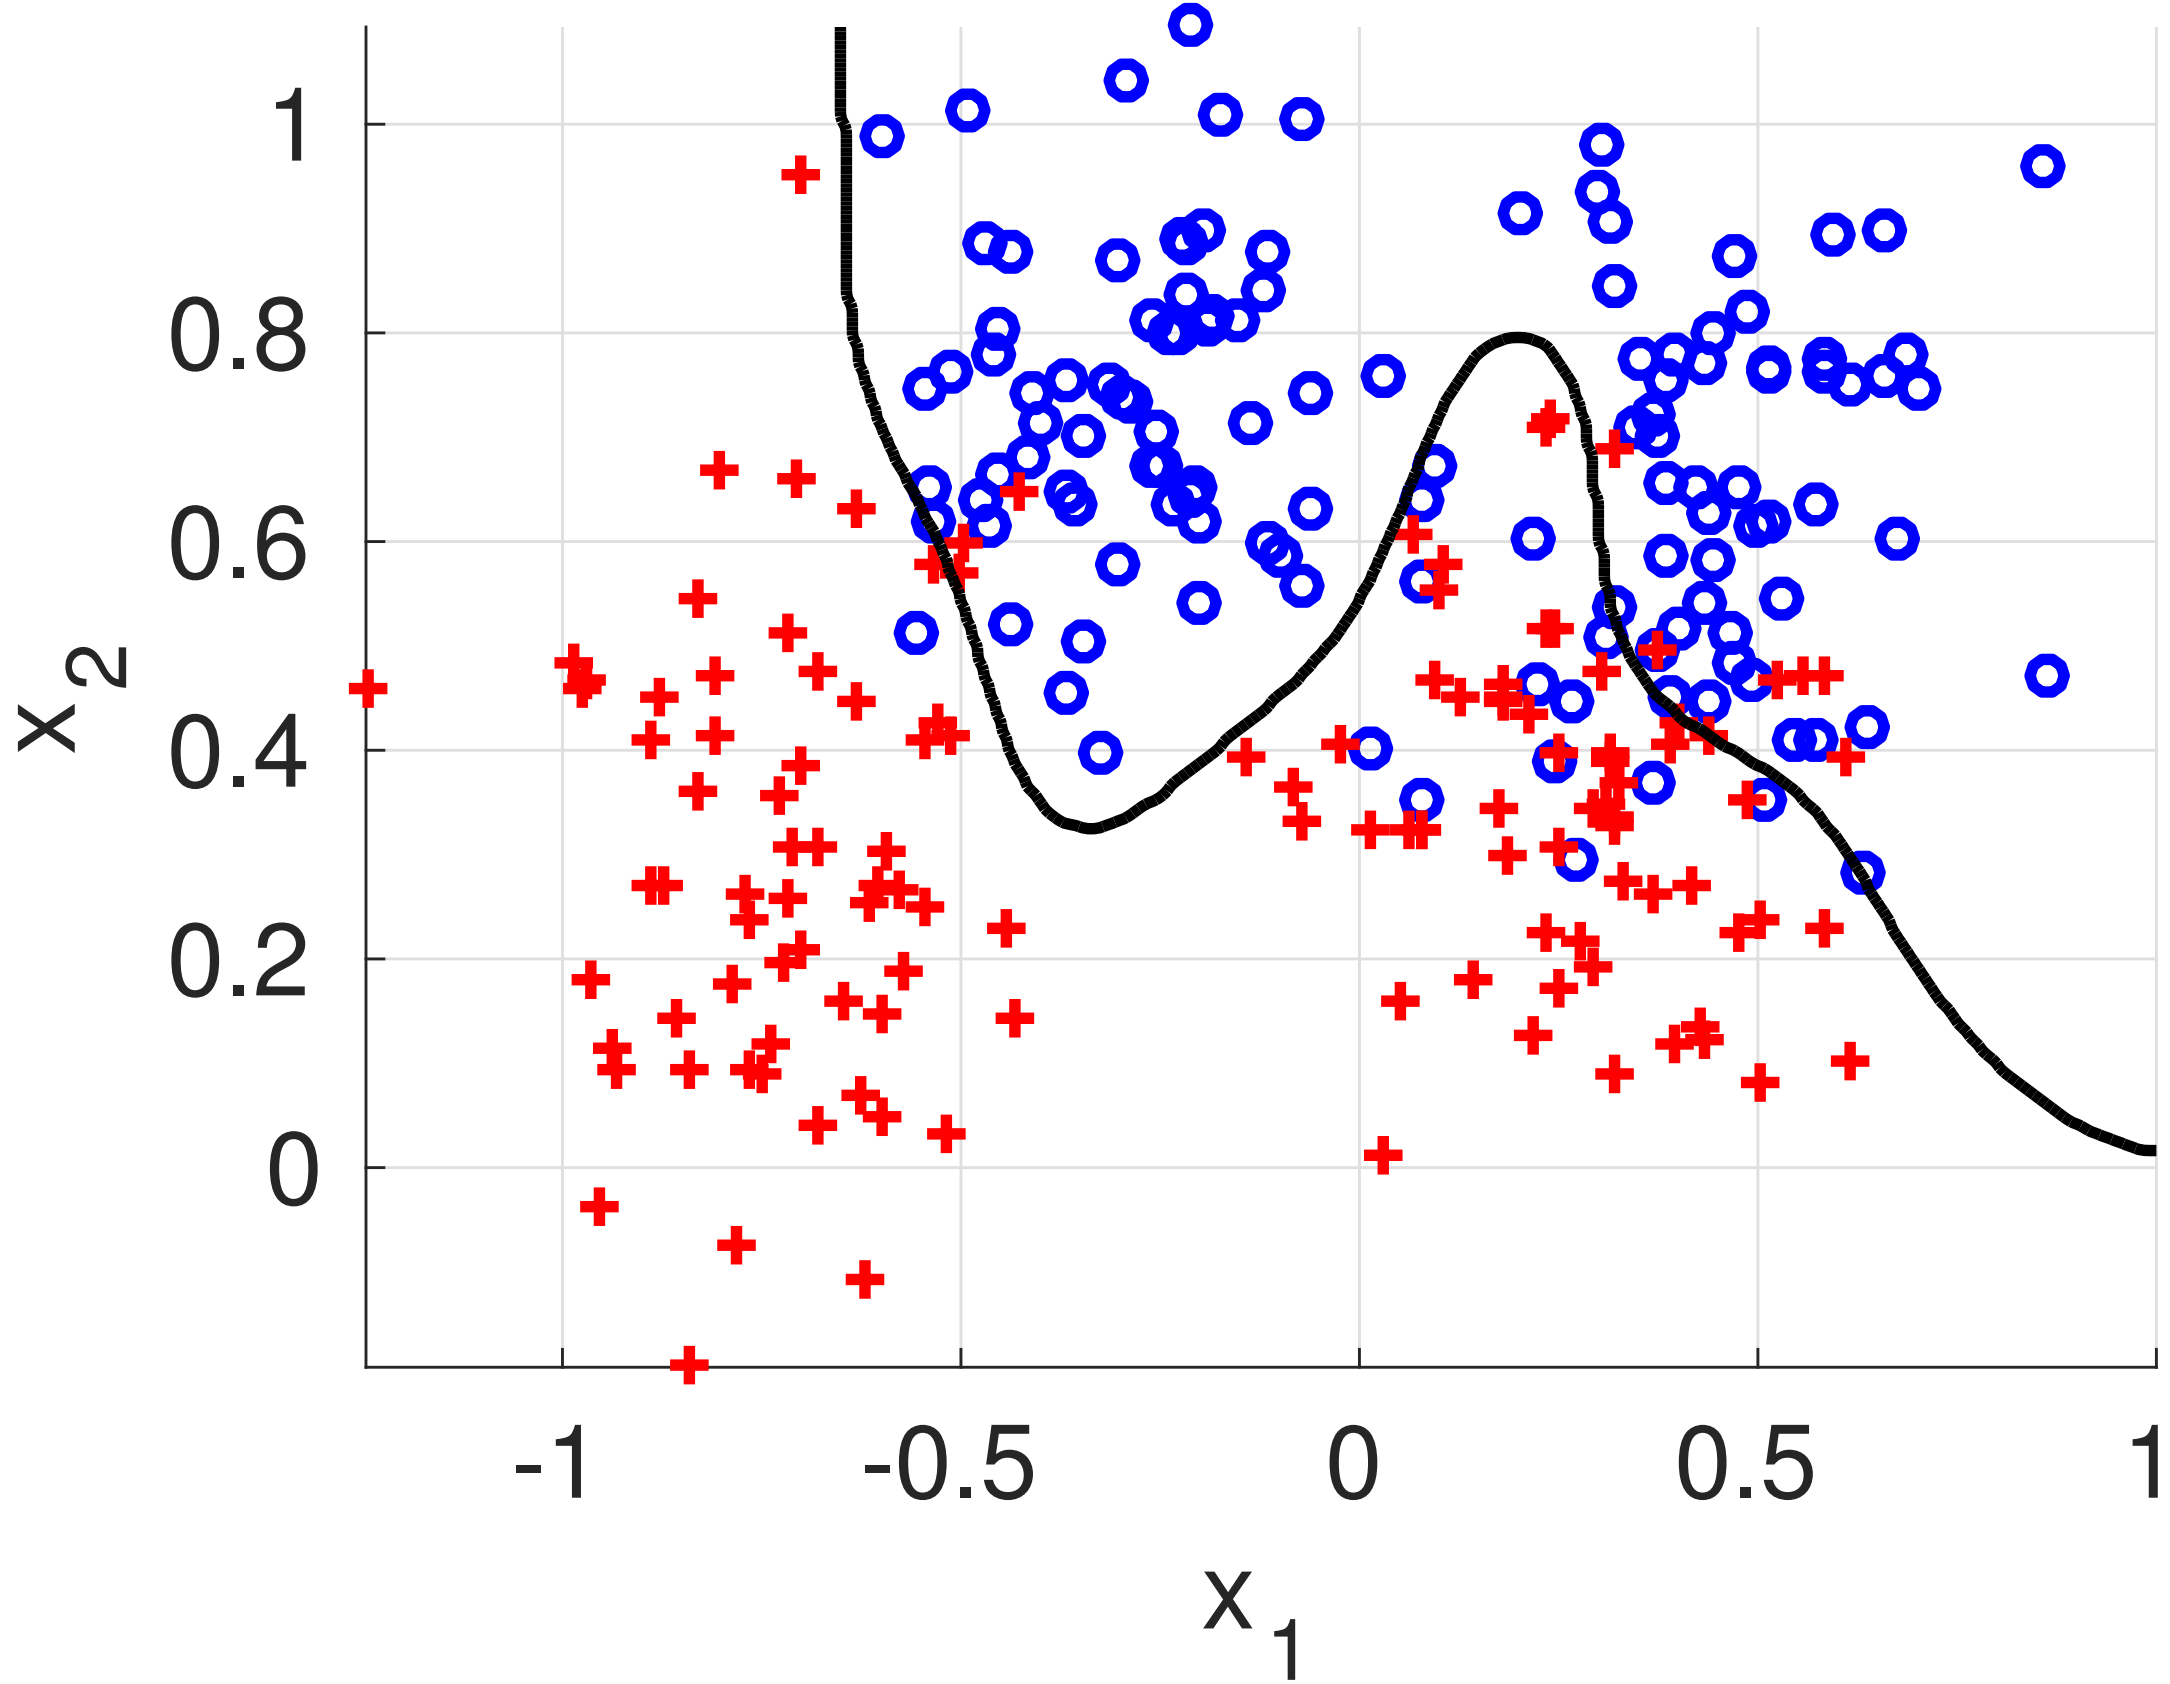
err=10.8%



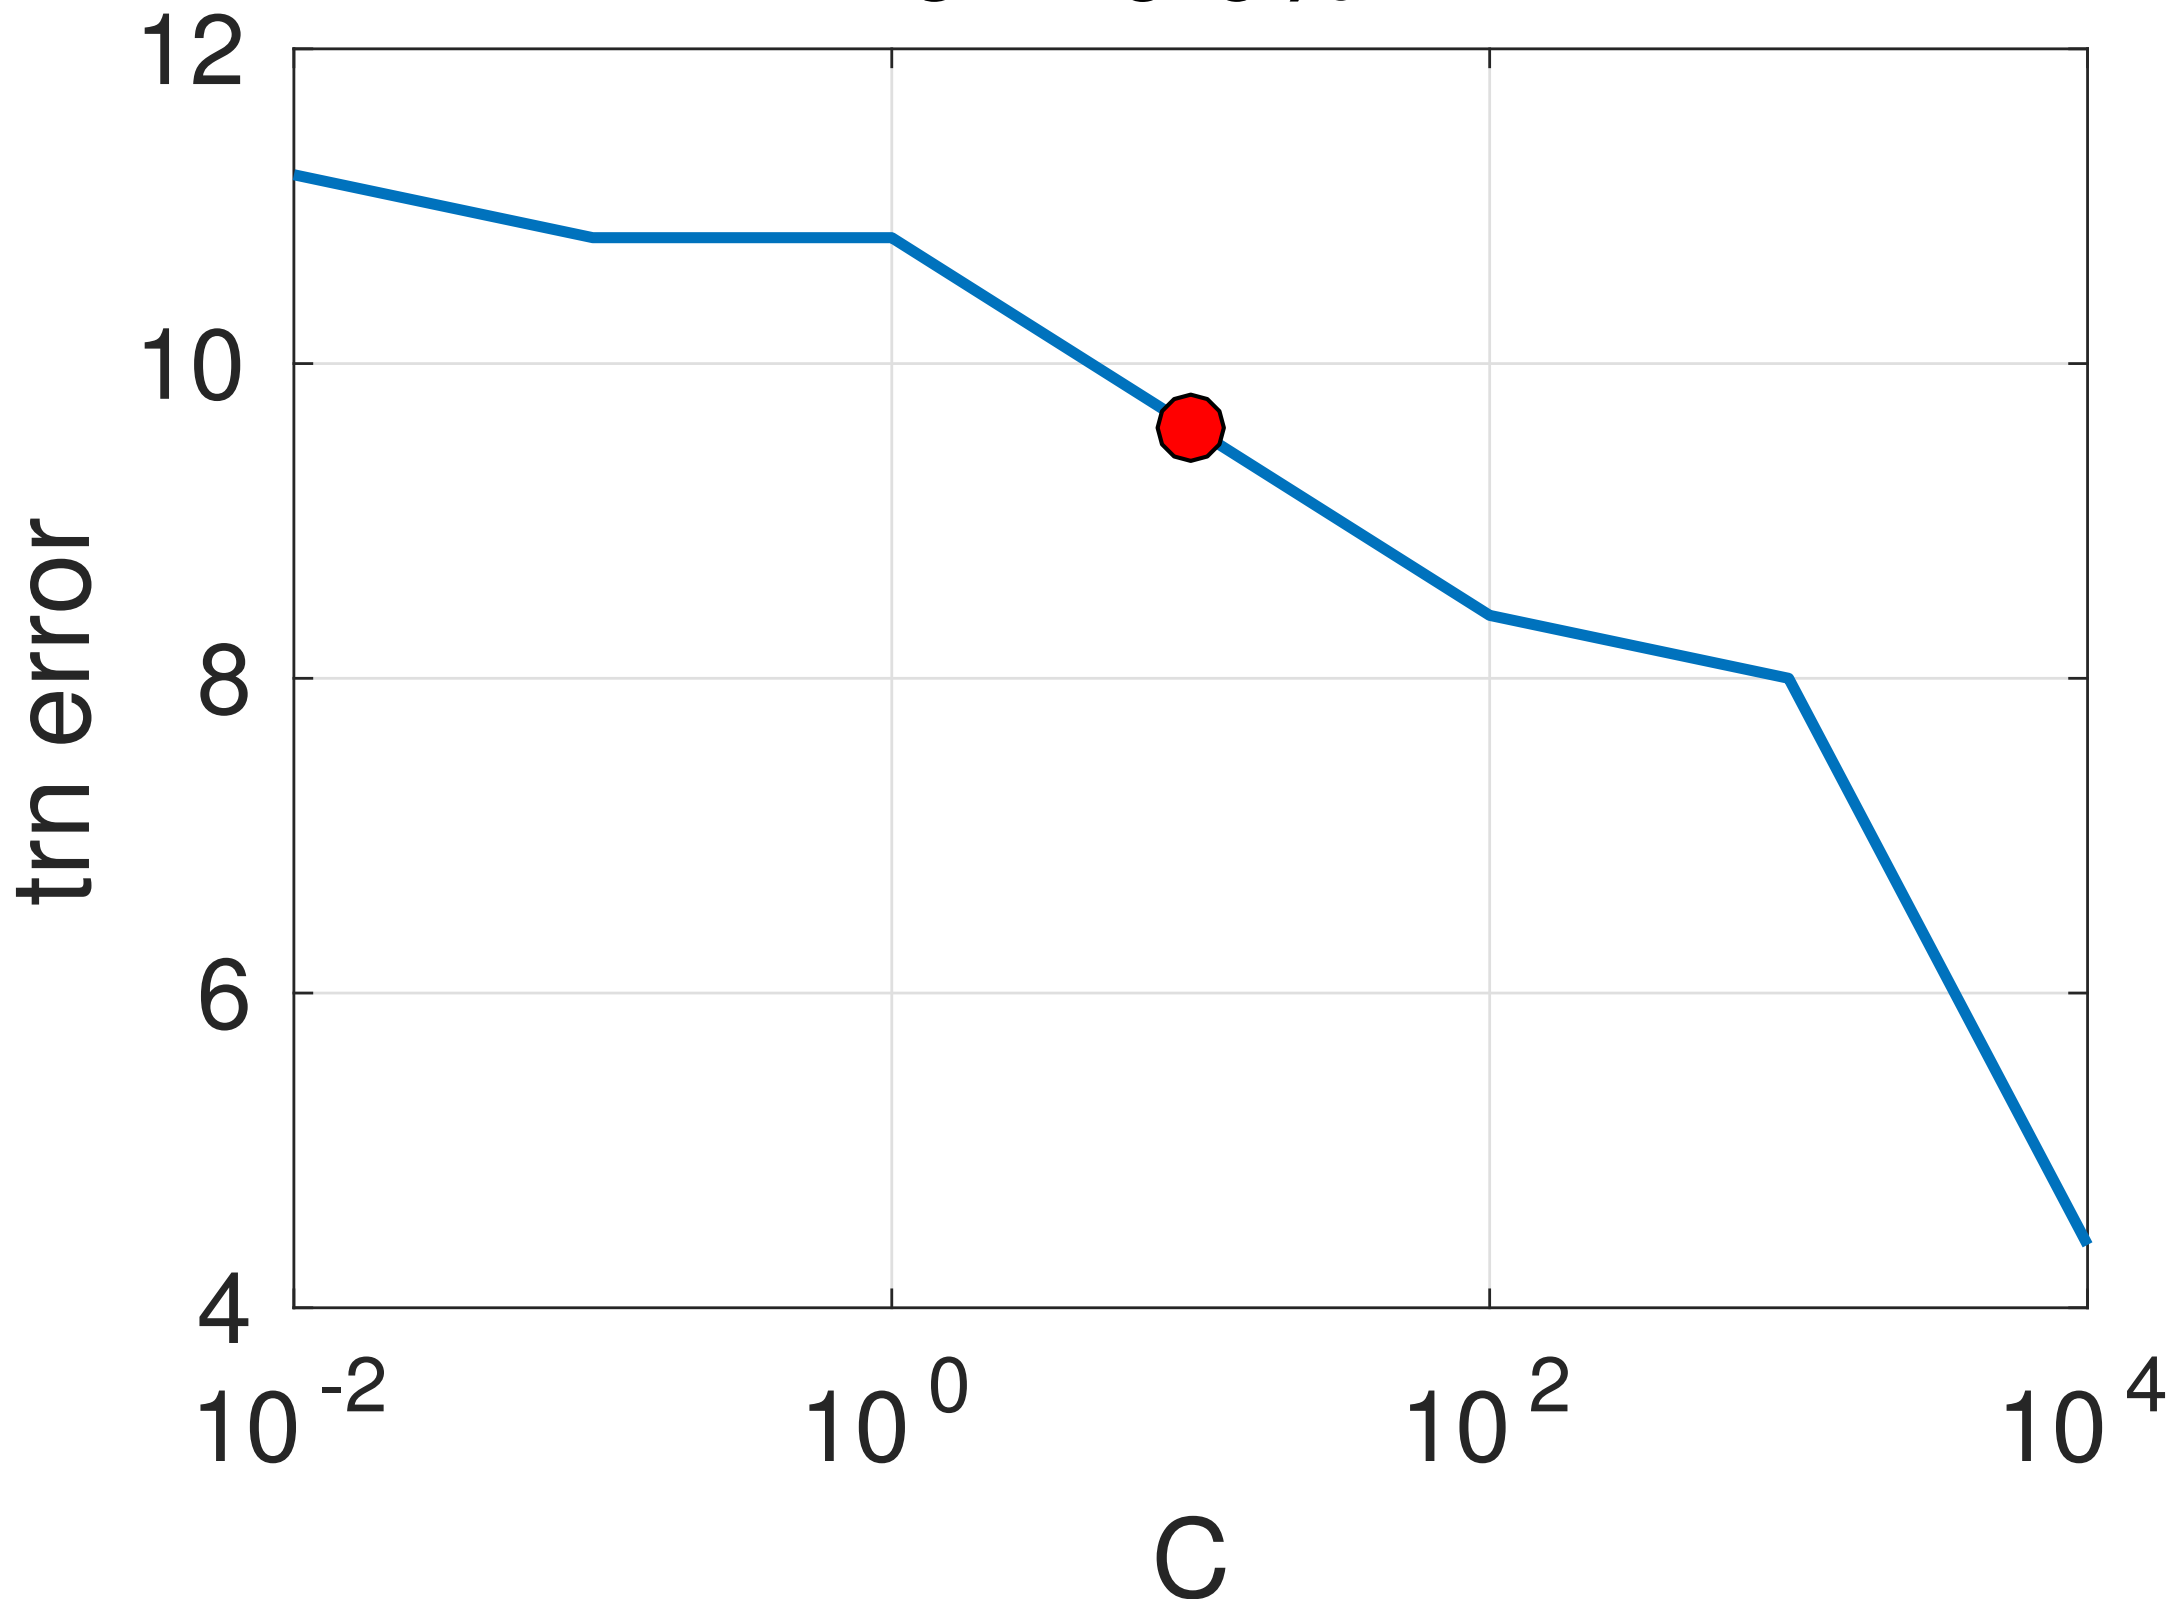


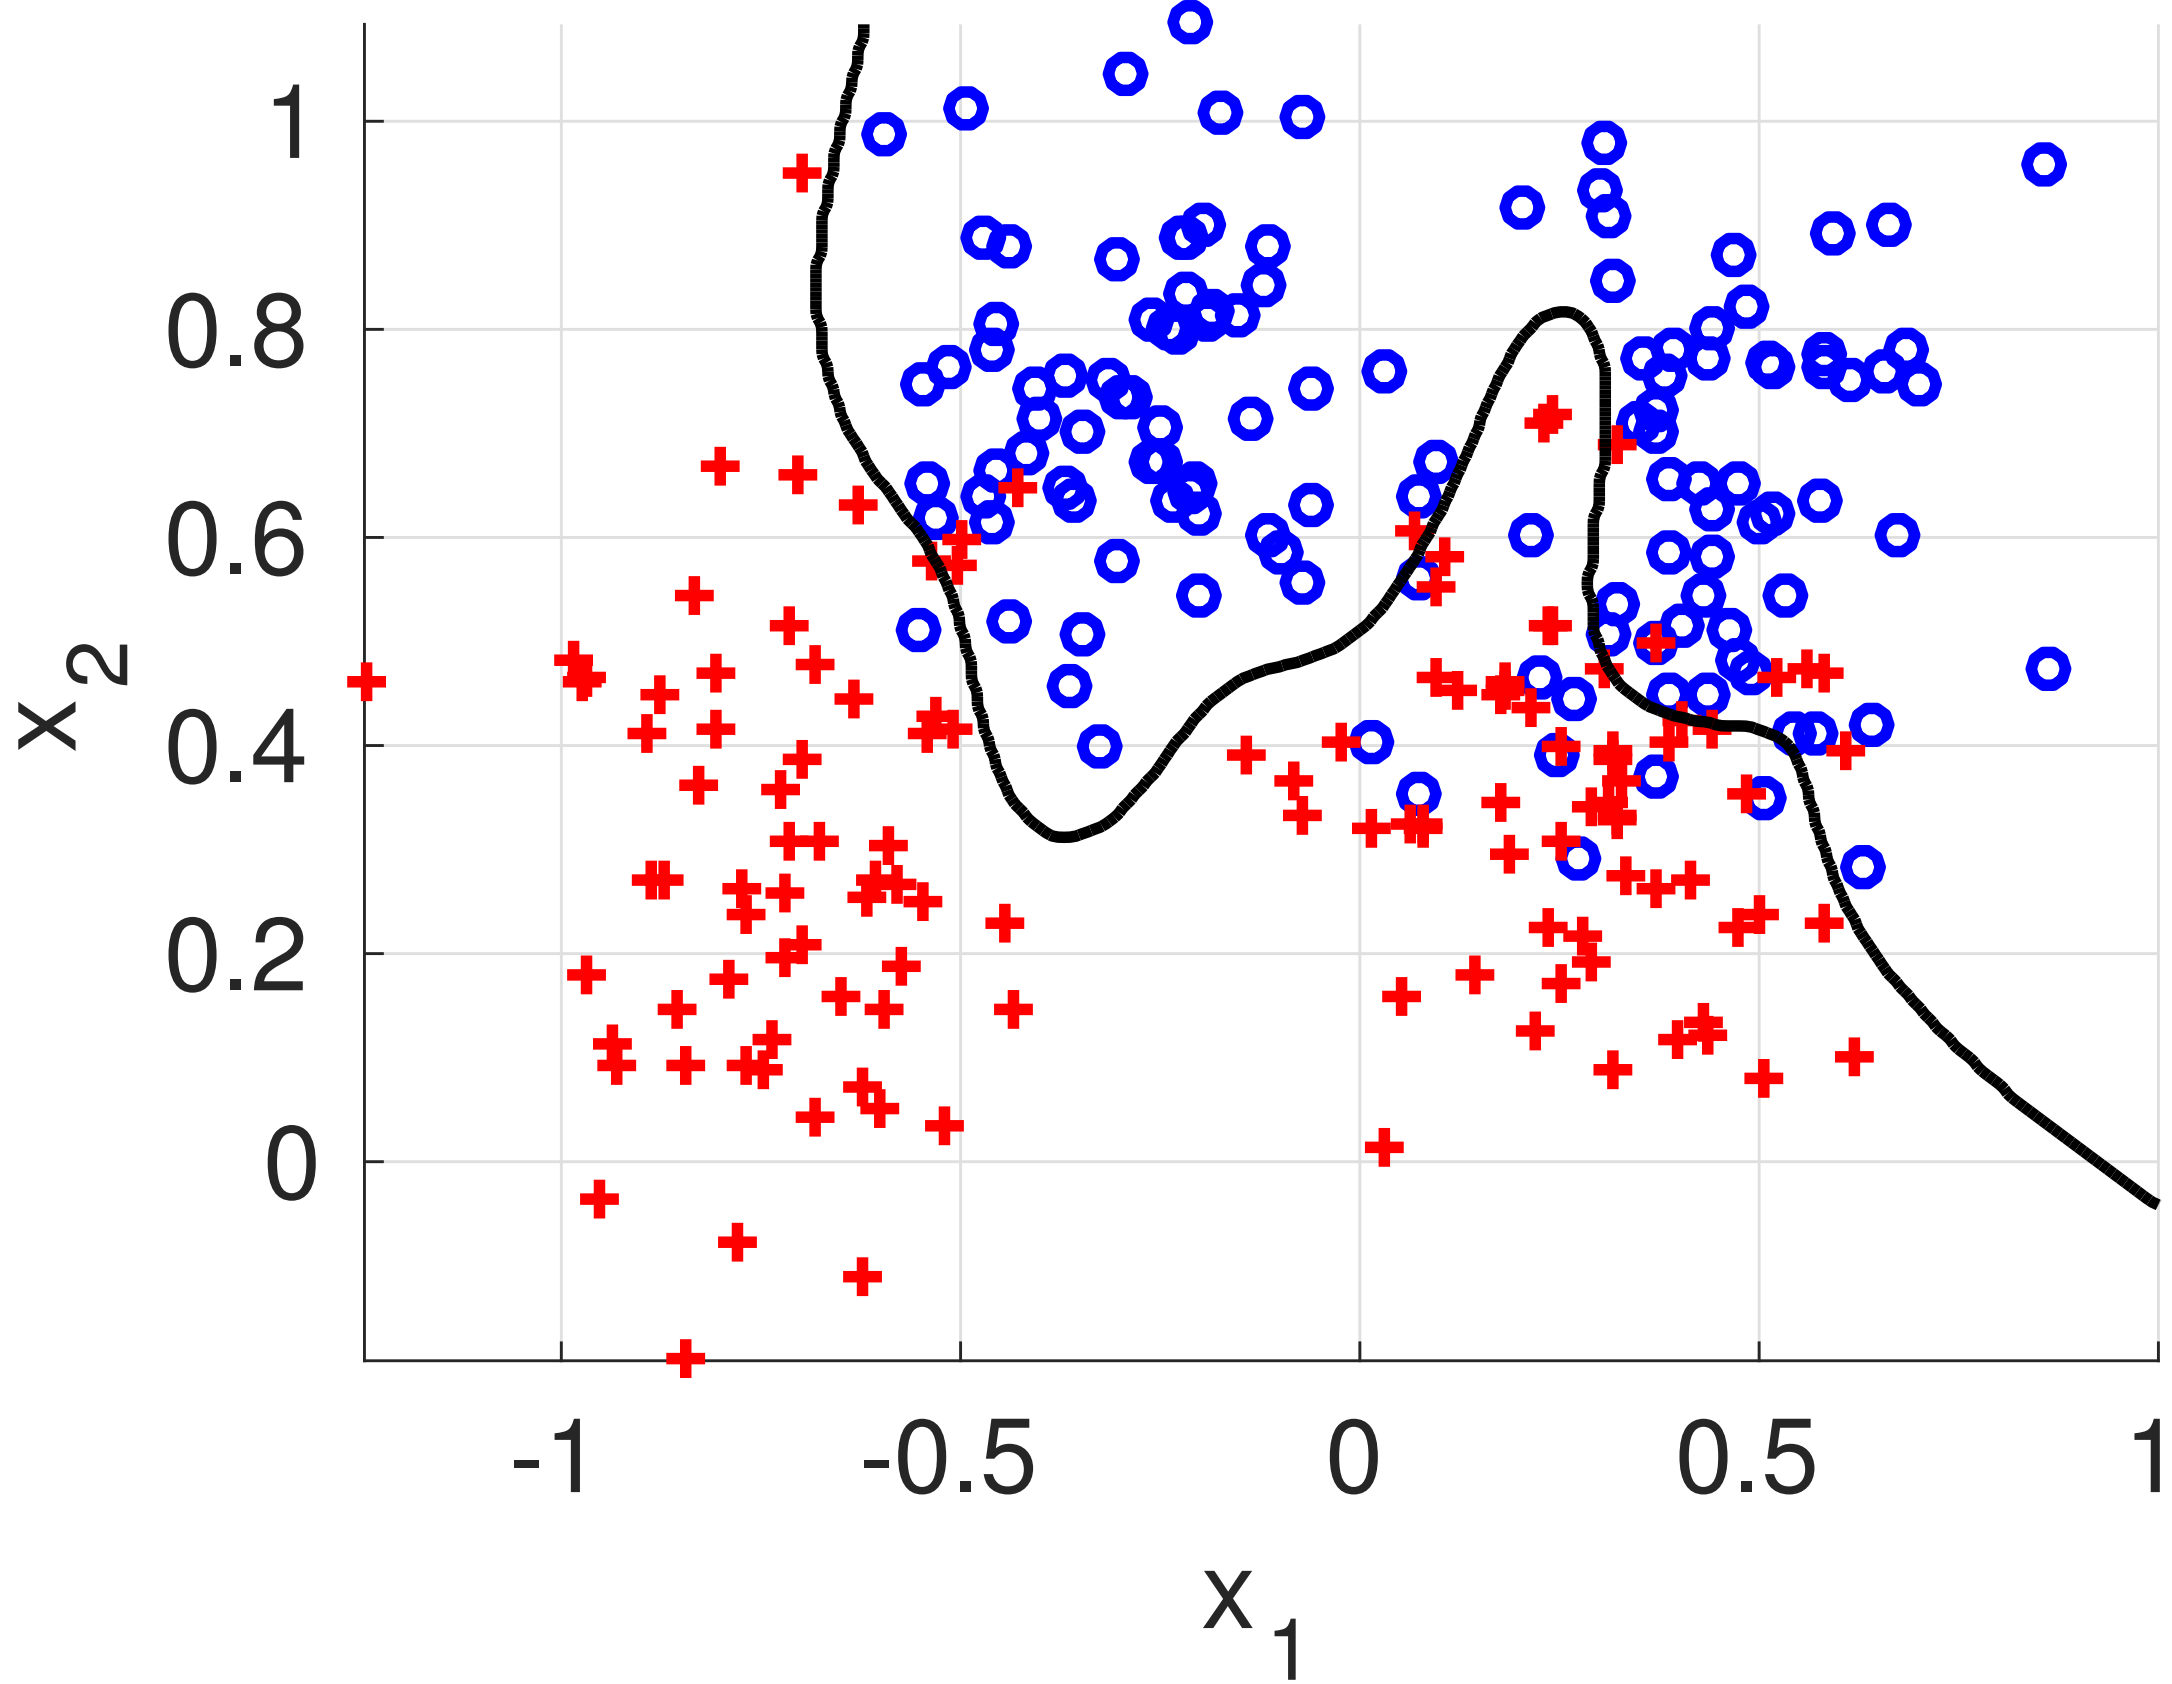
err=10.8%



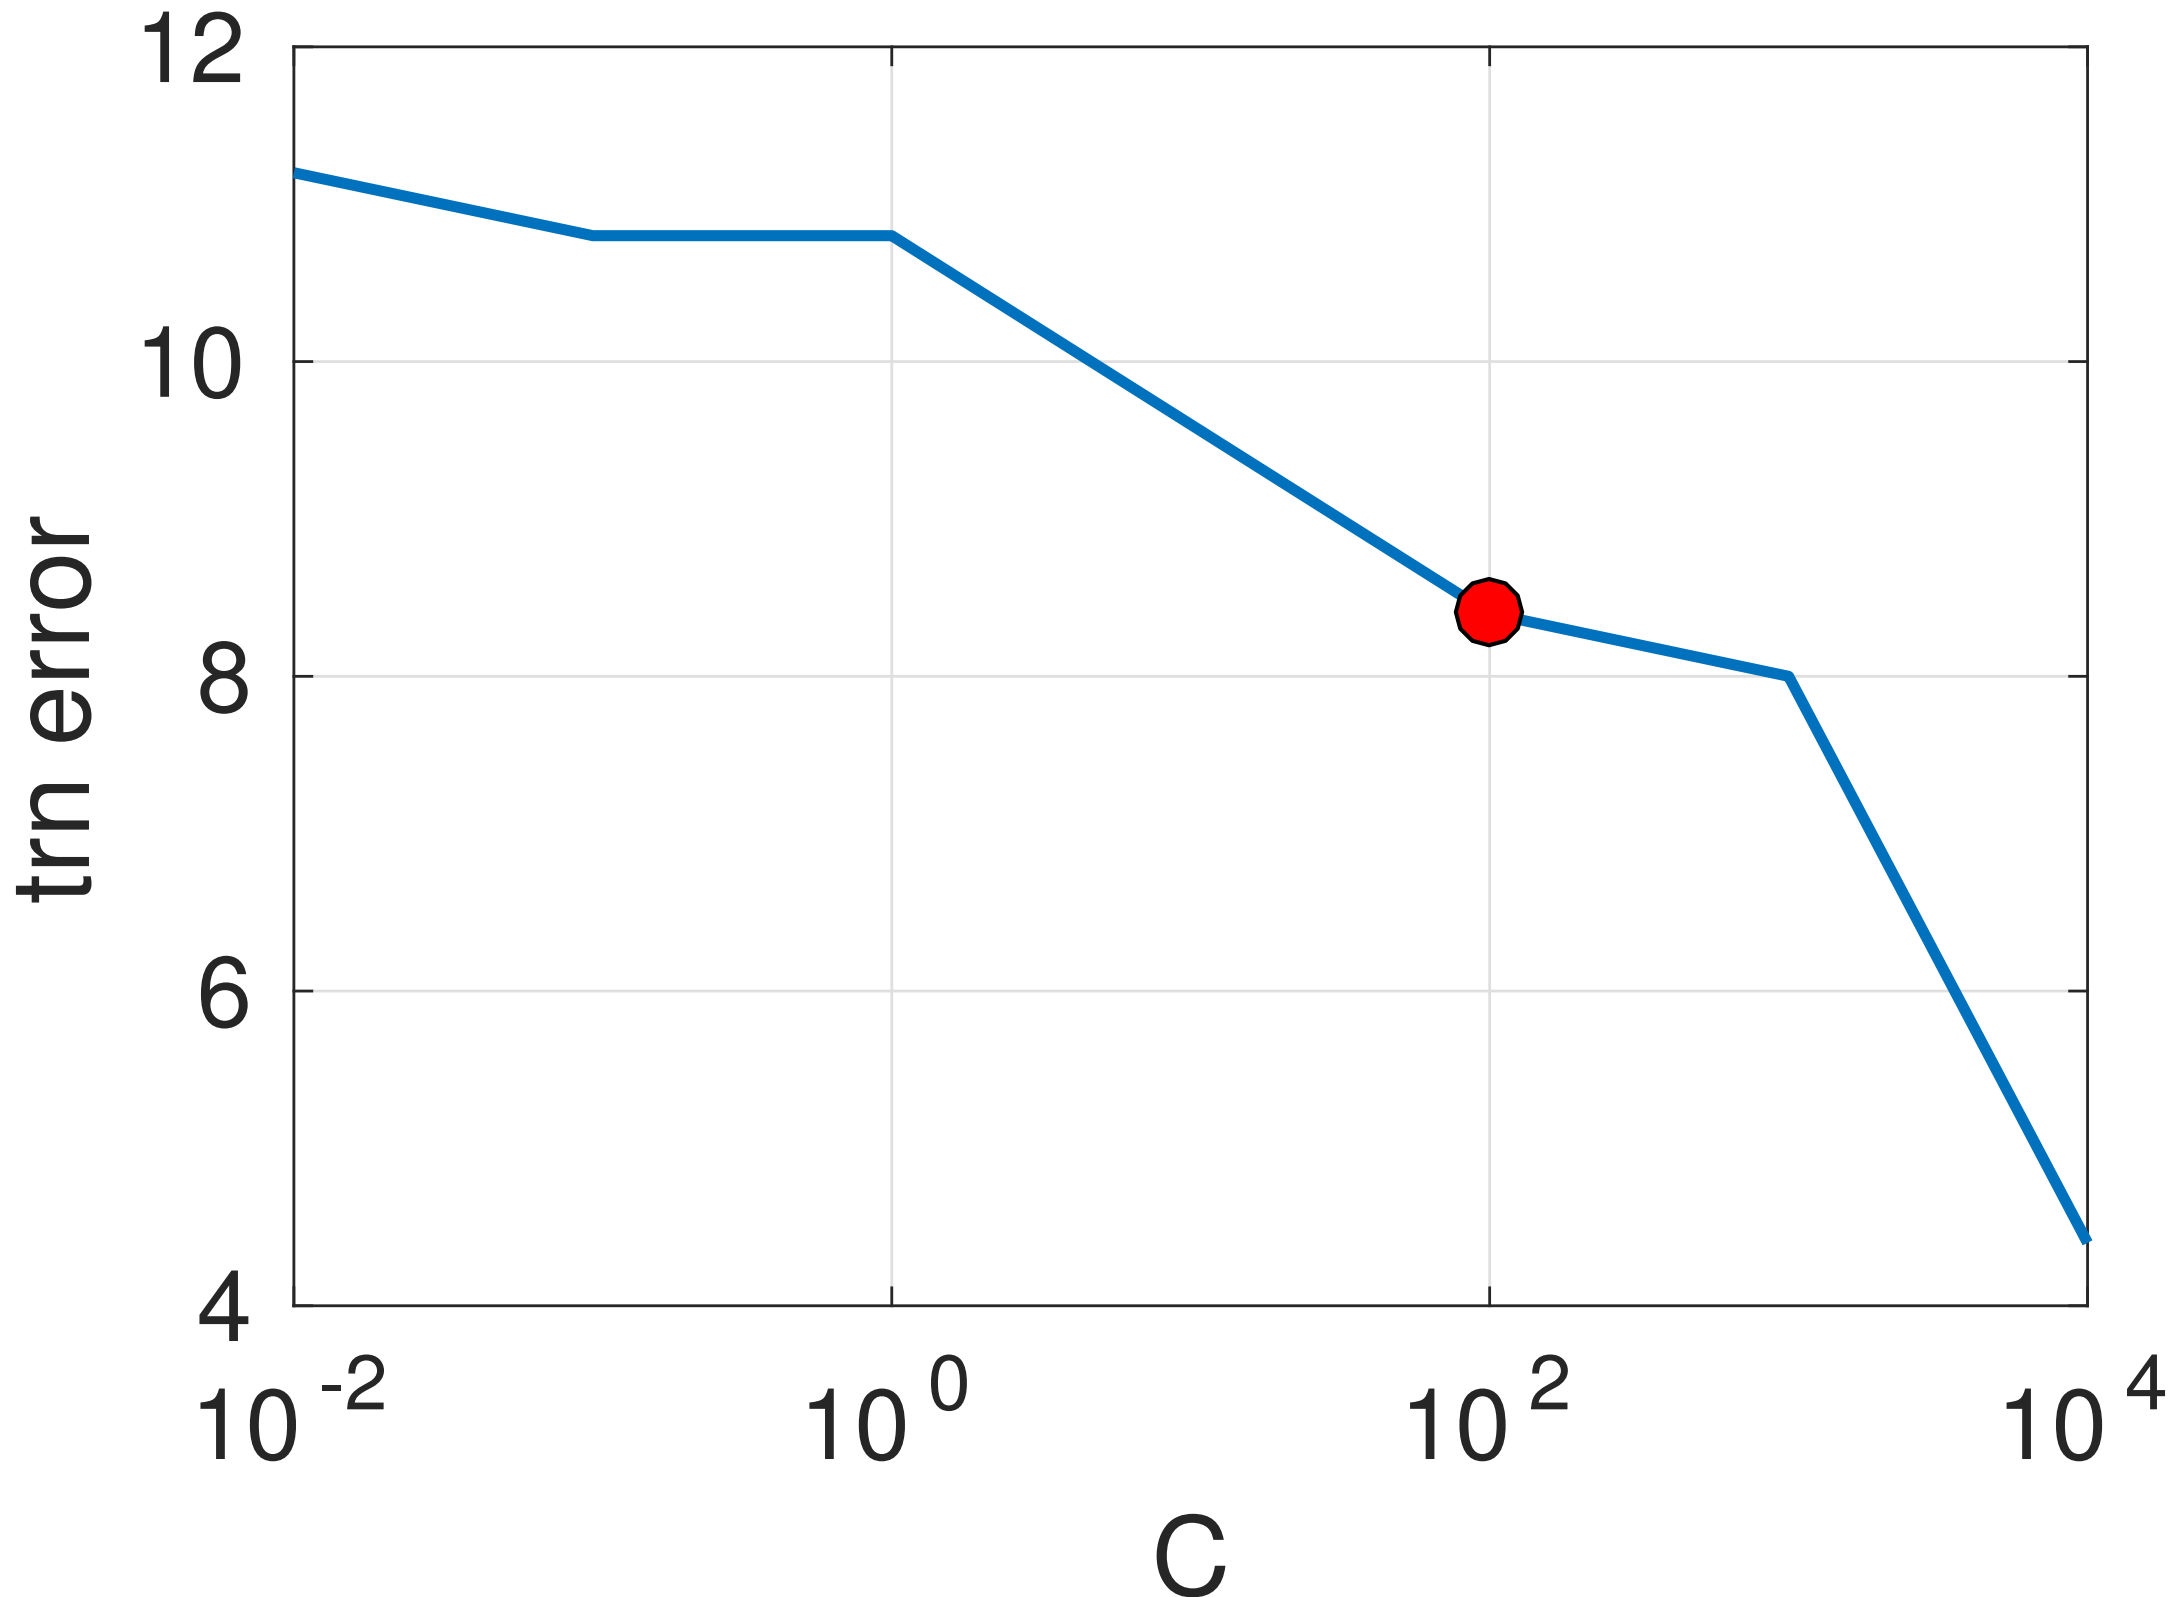


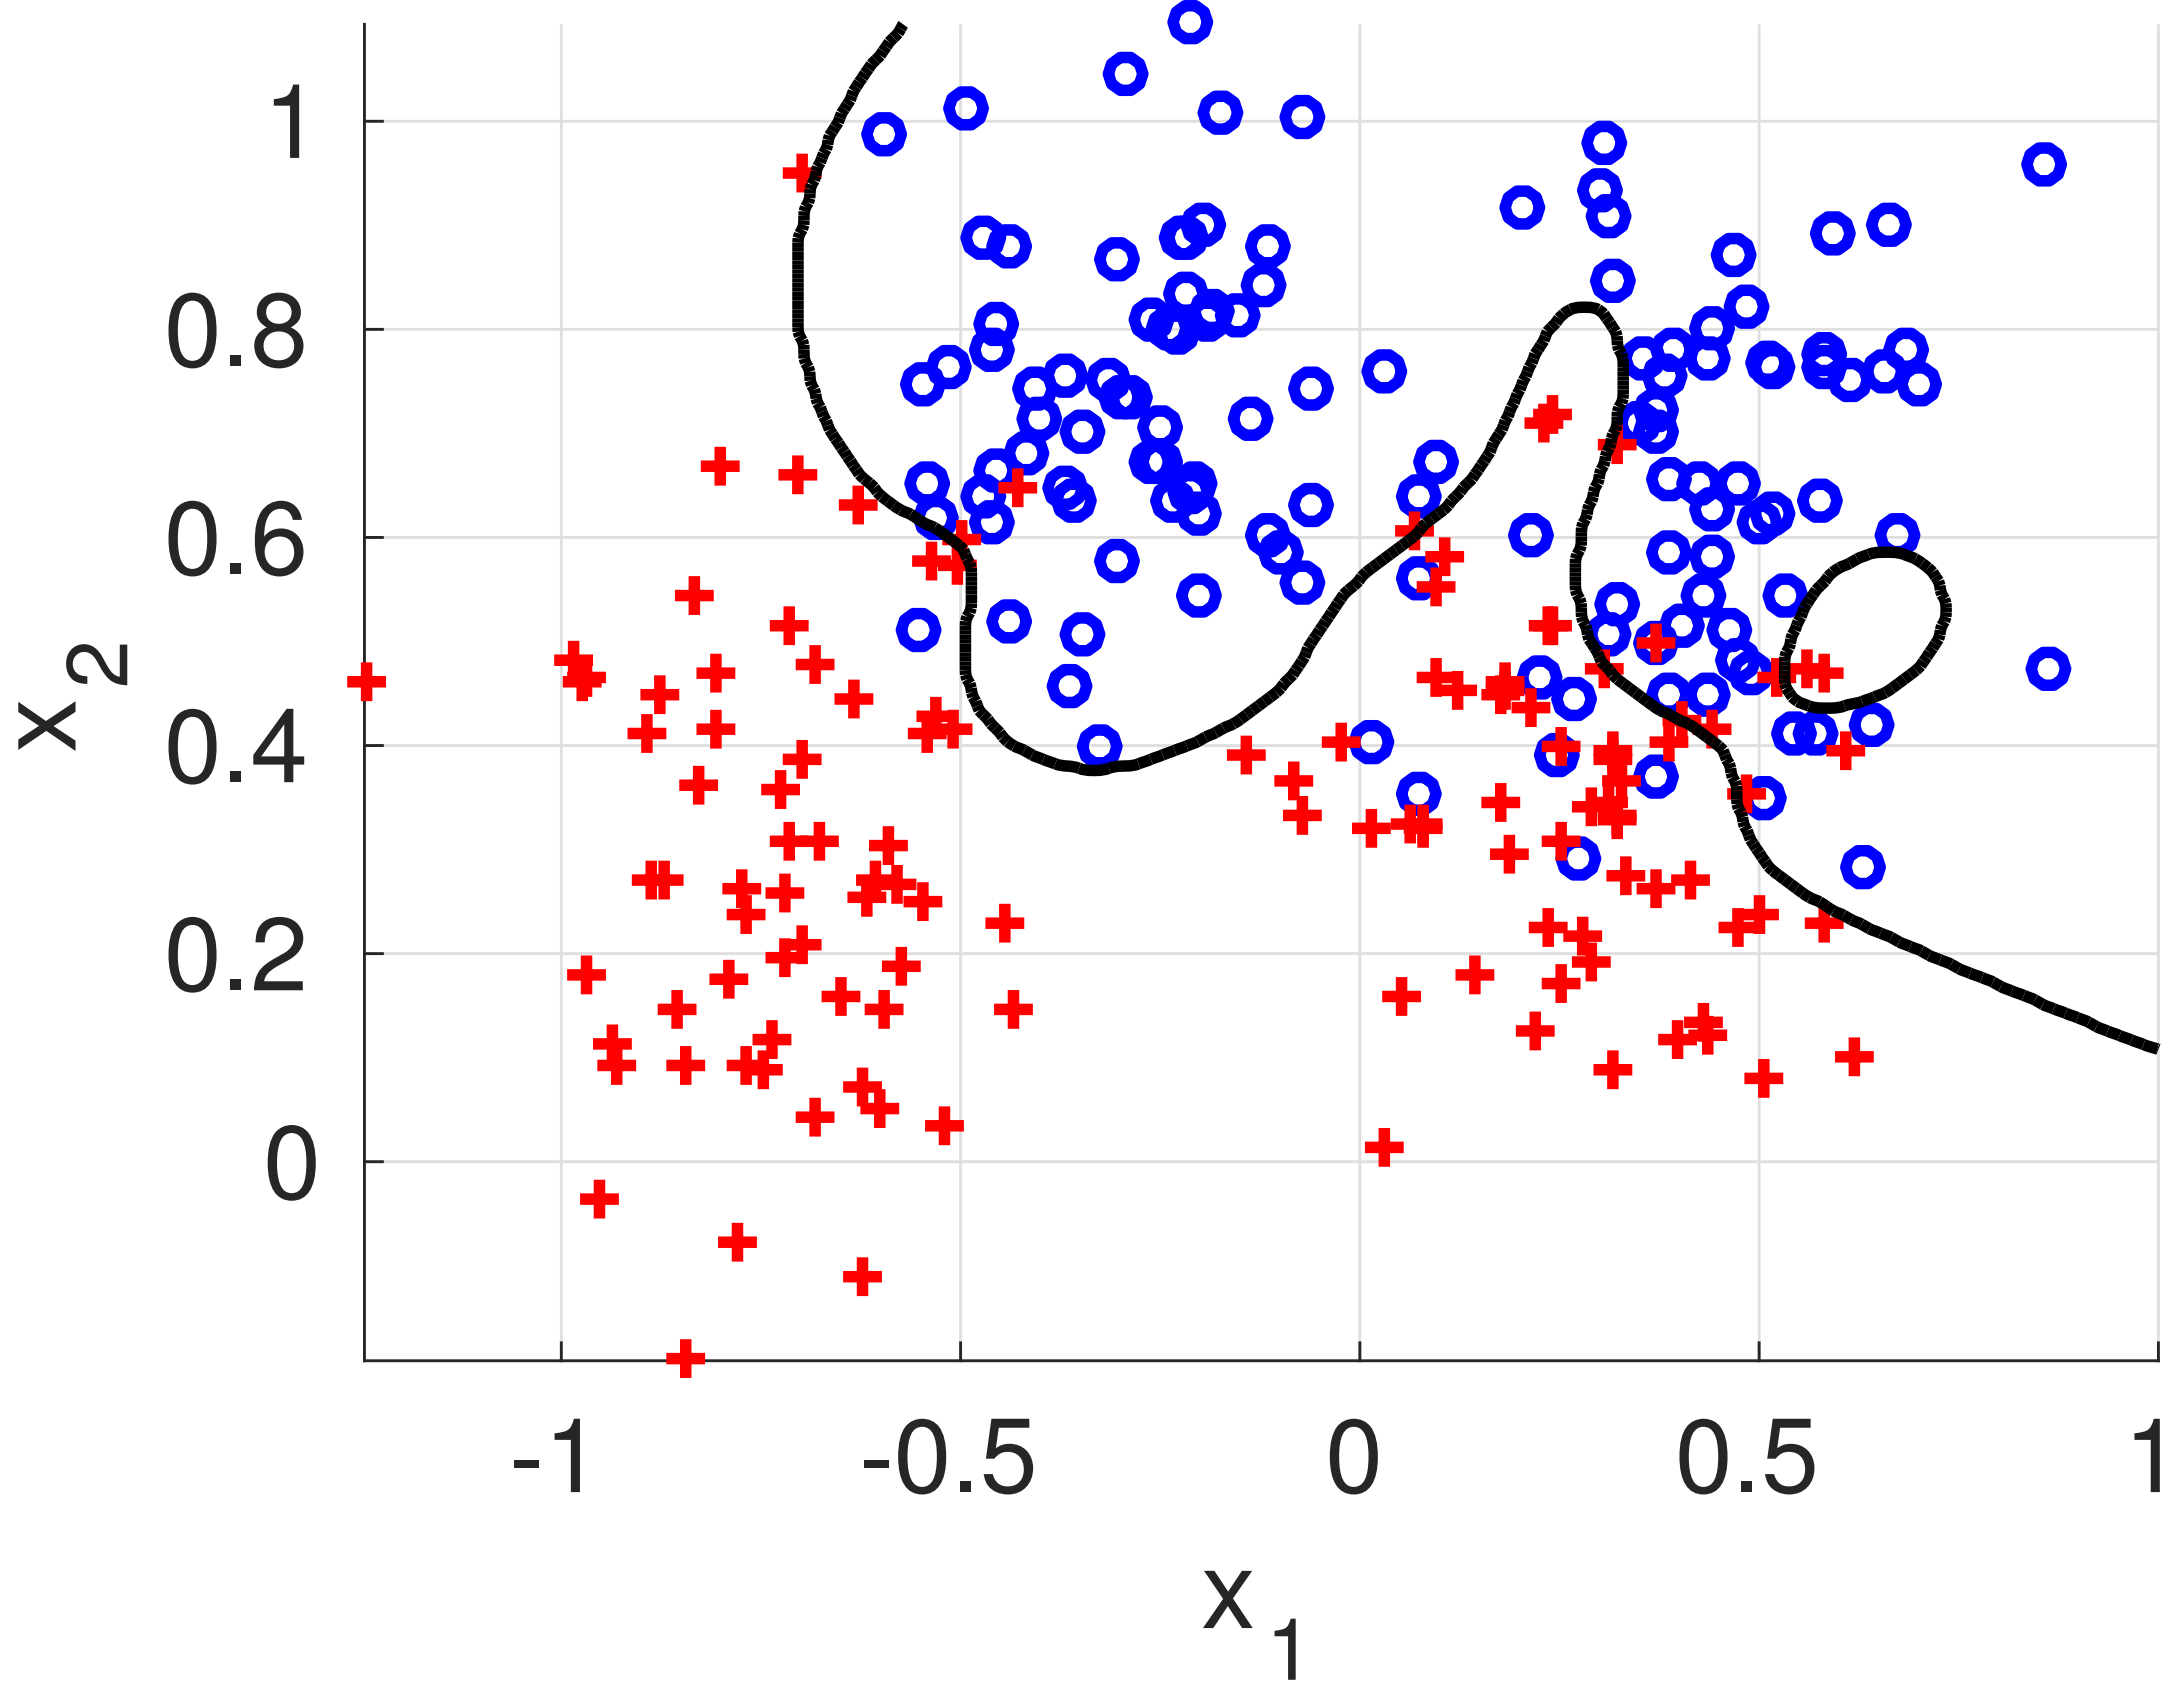
err=9.6%



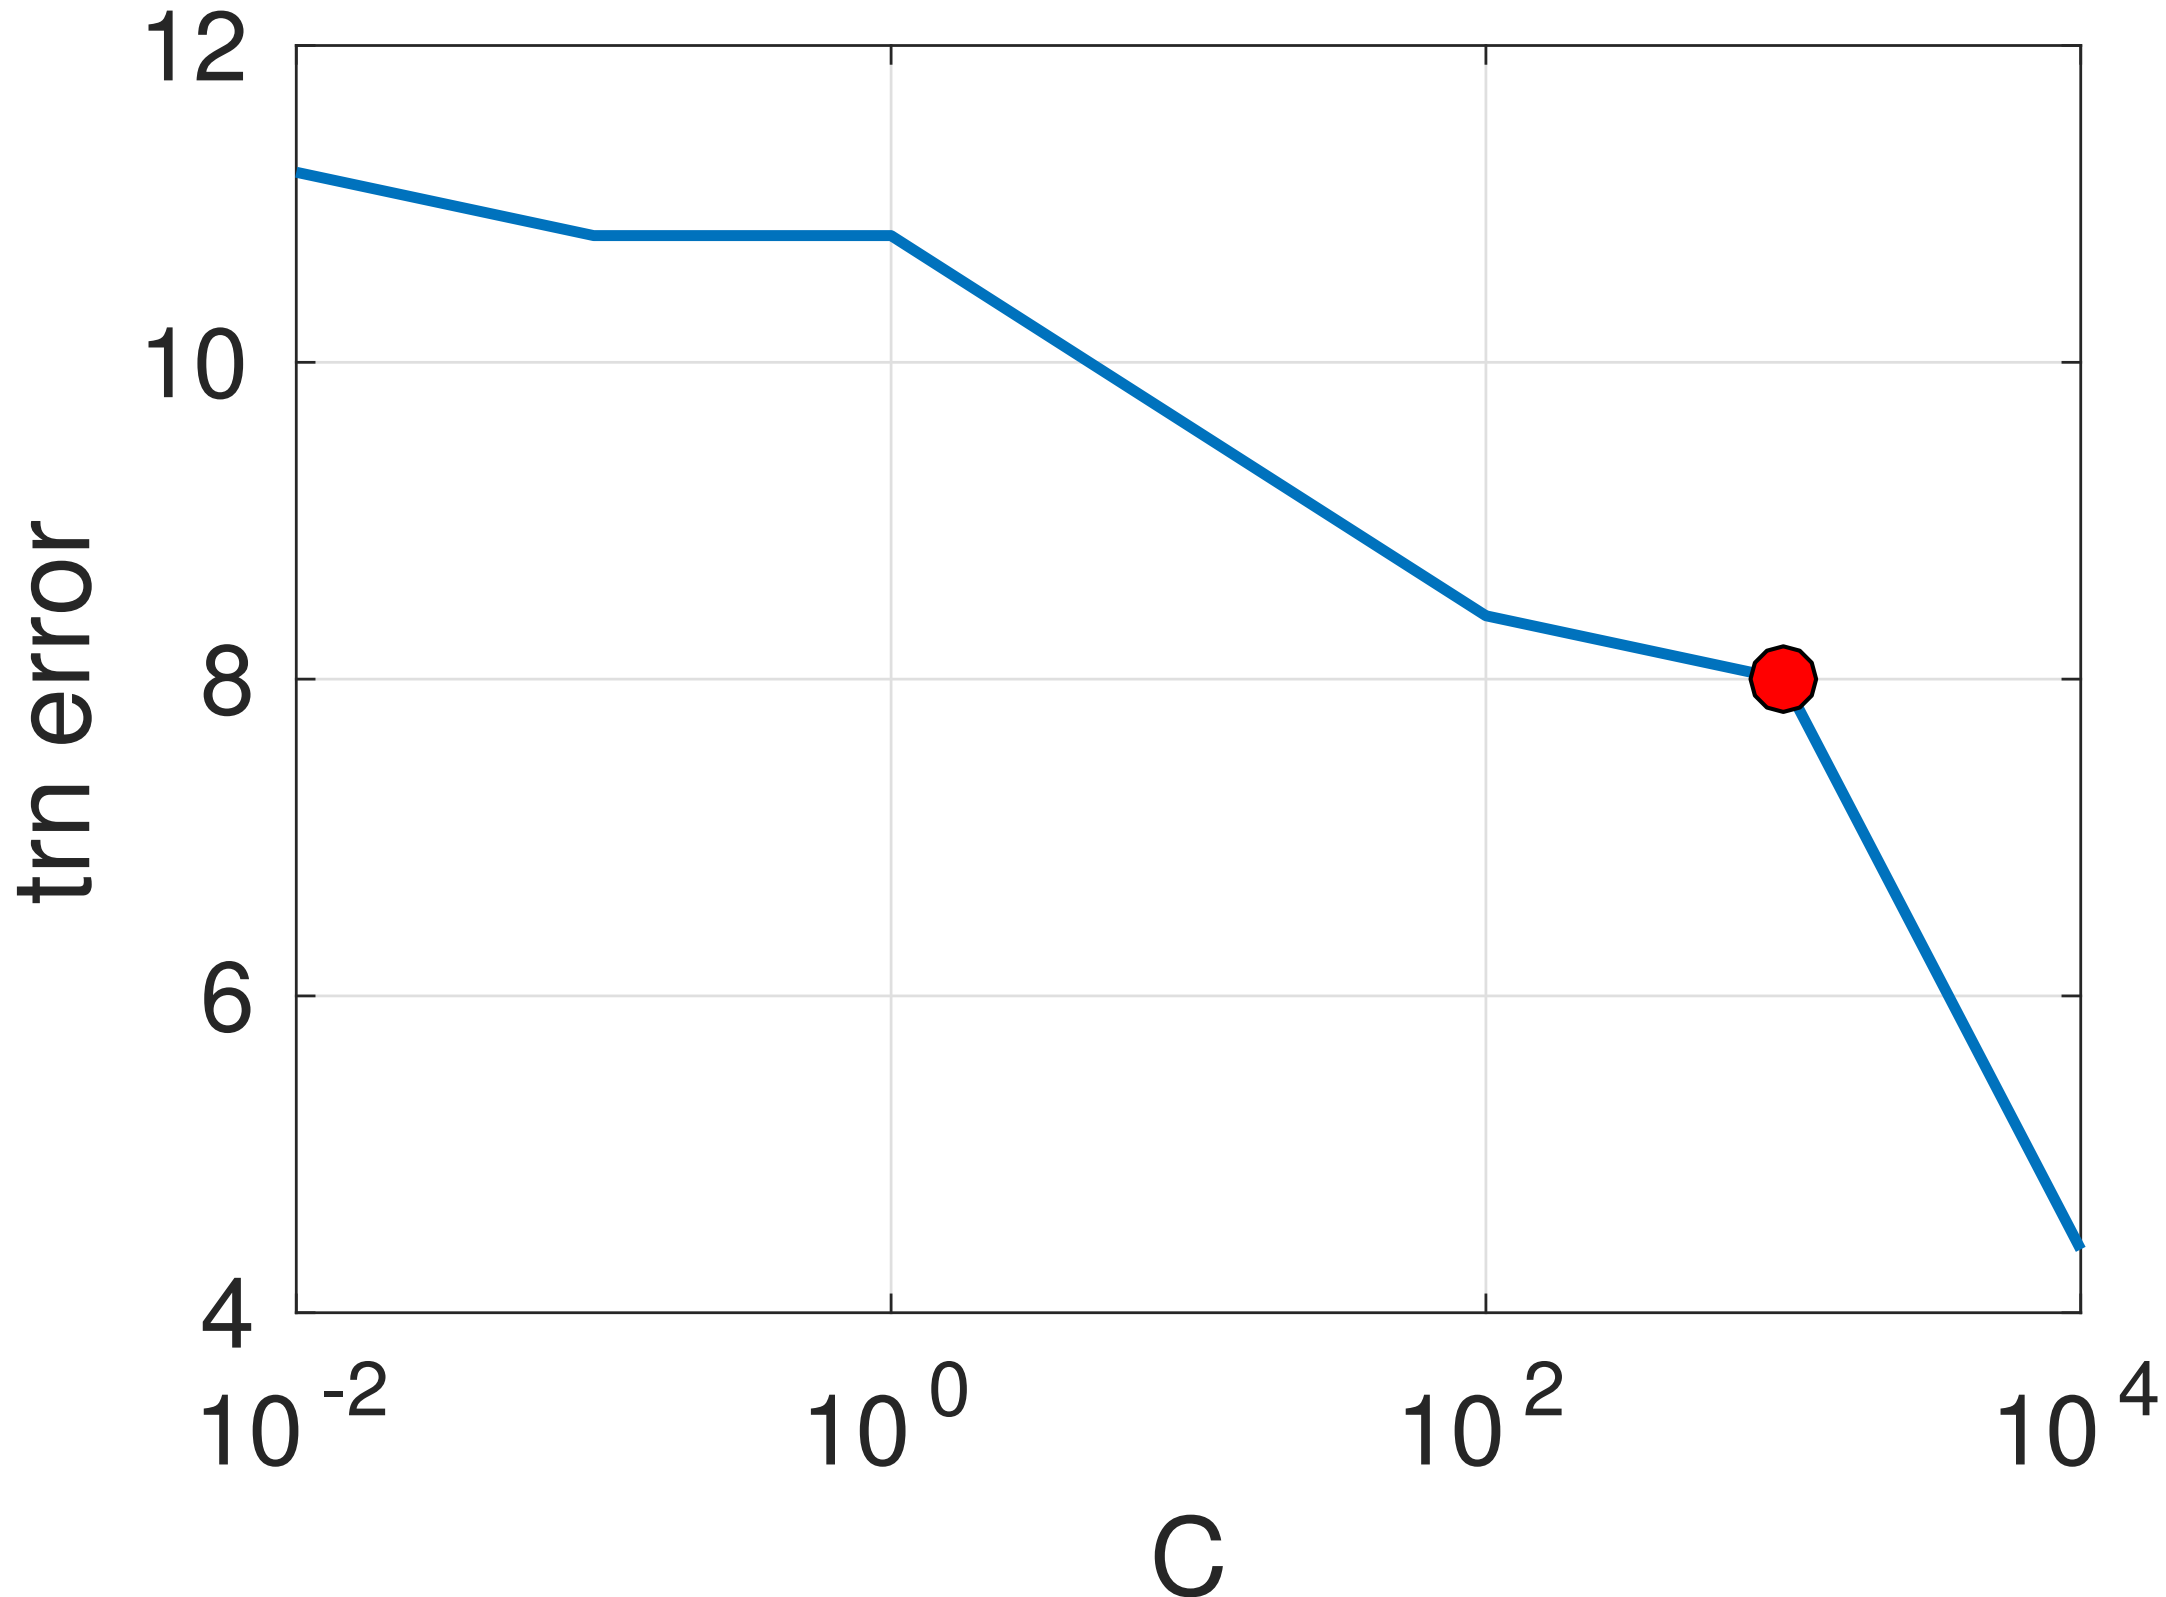


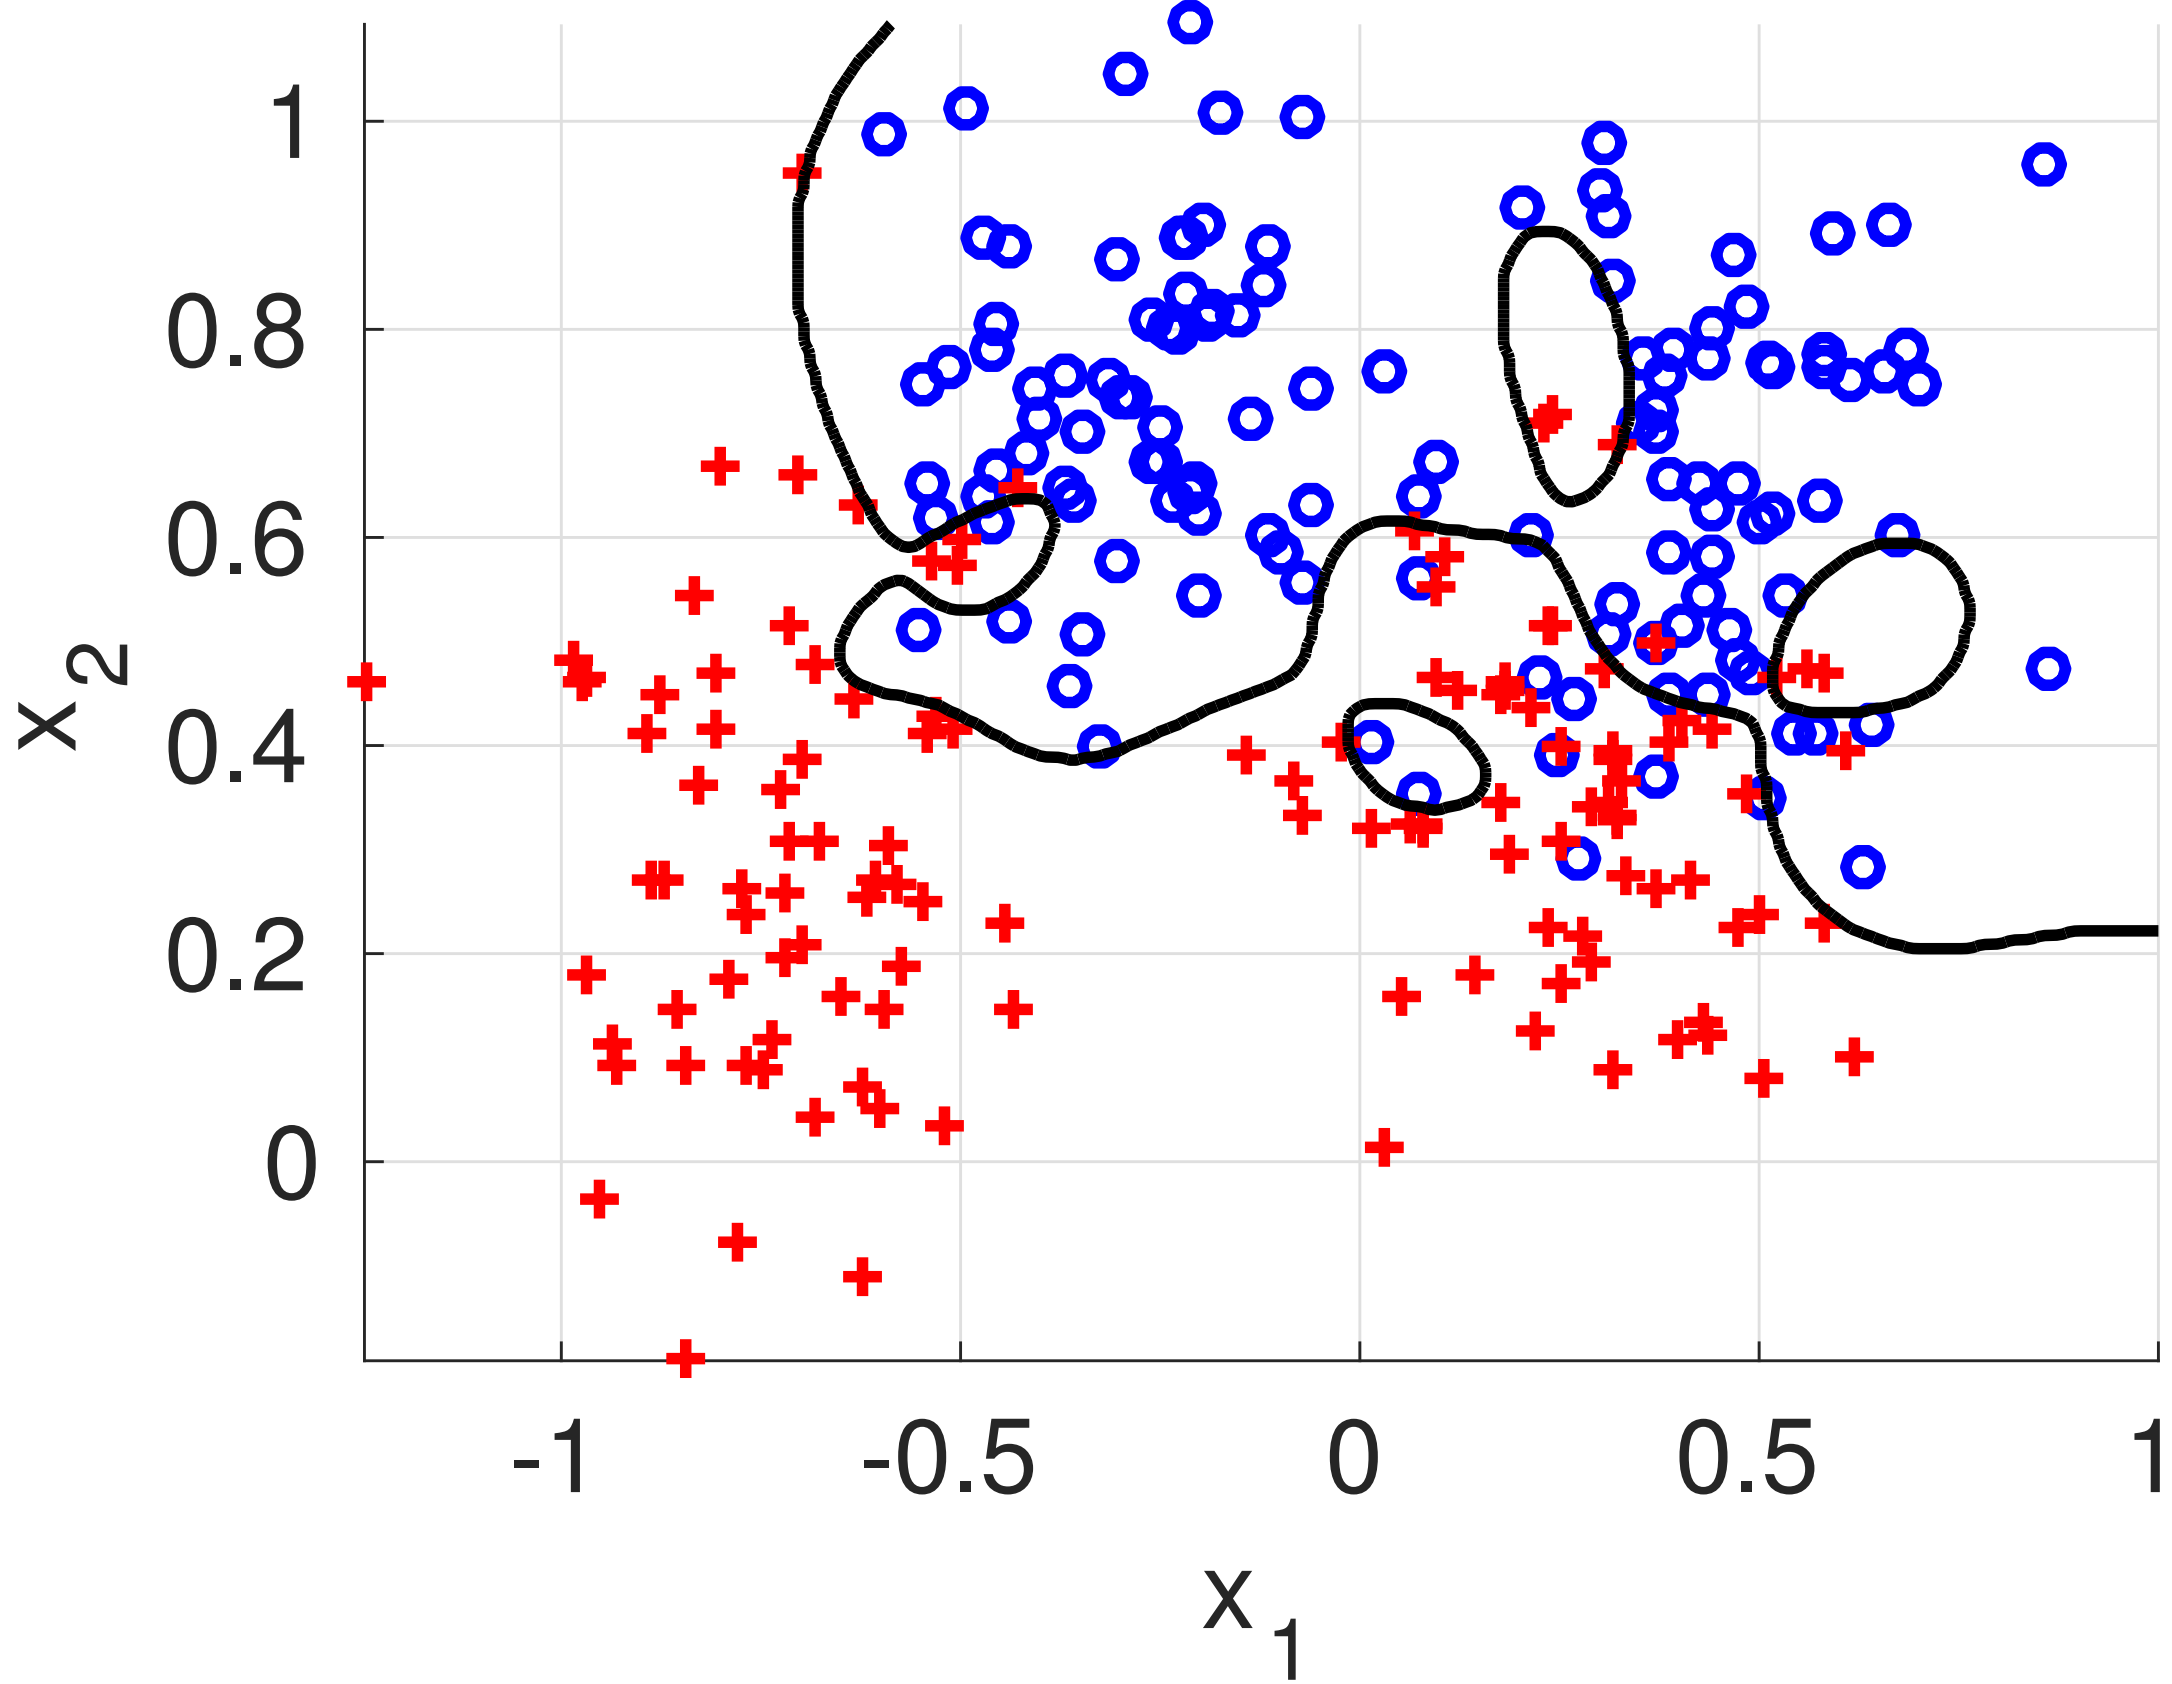
err=8.4%



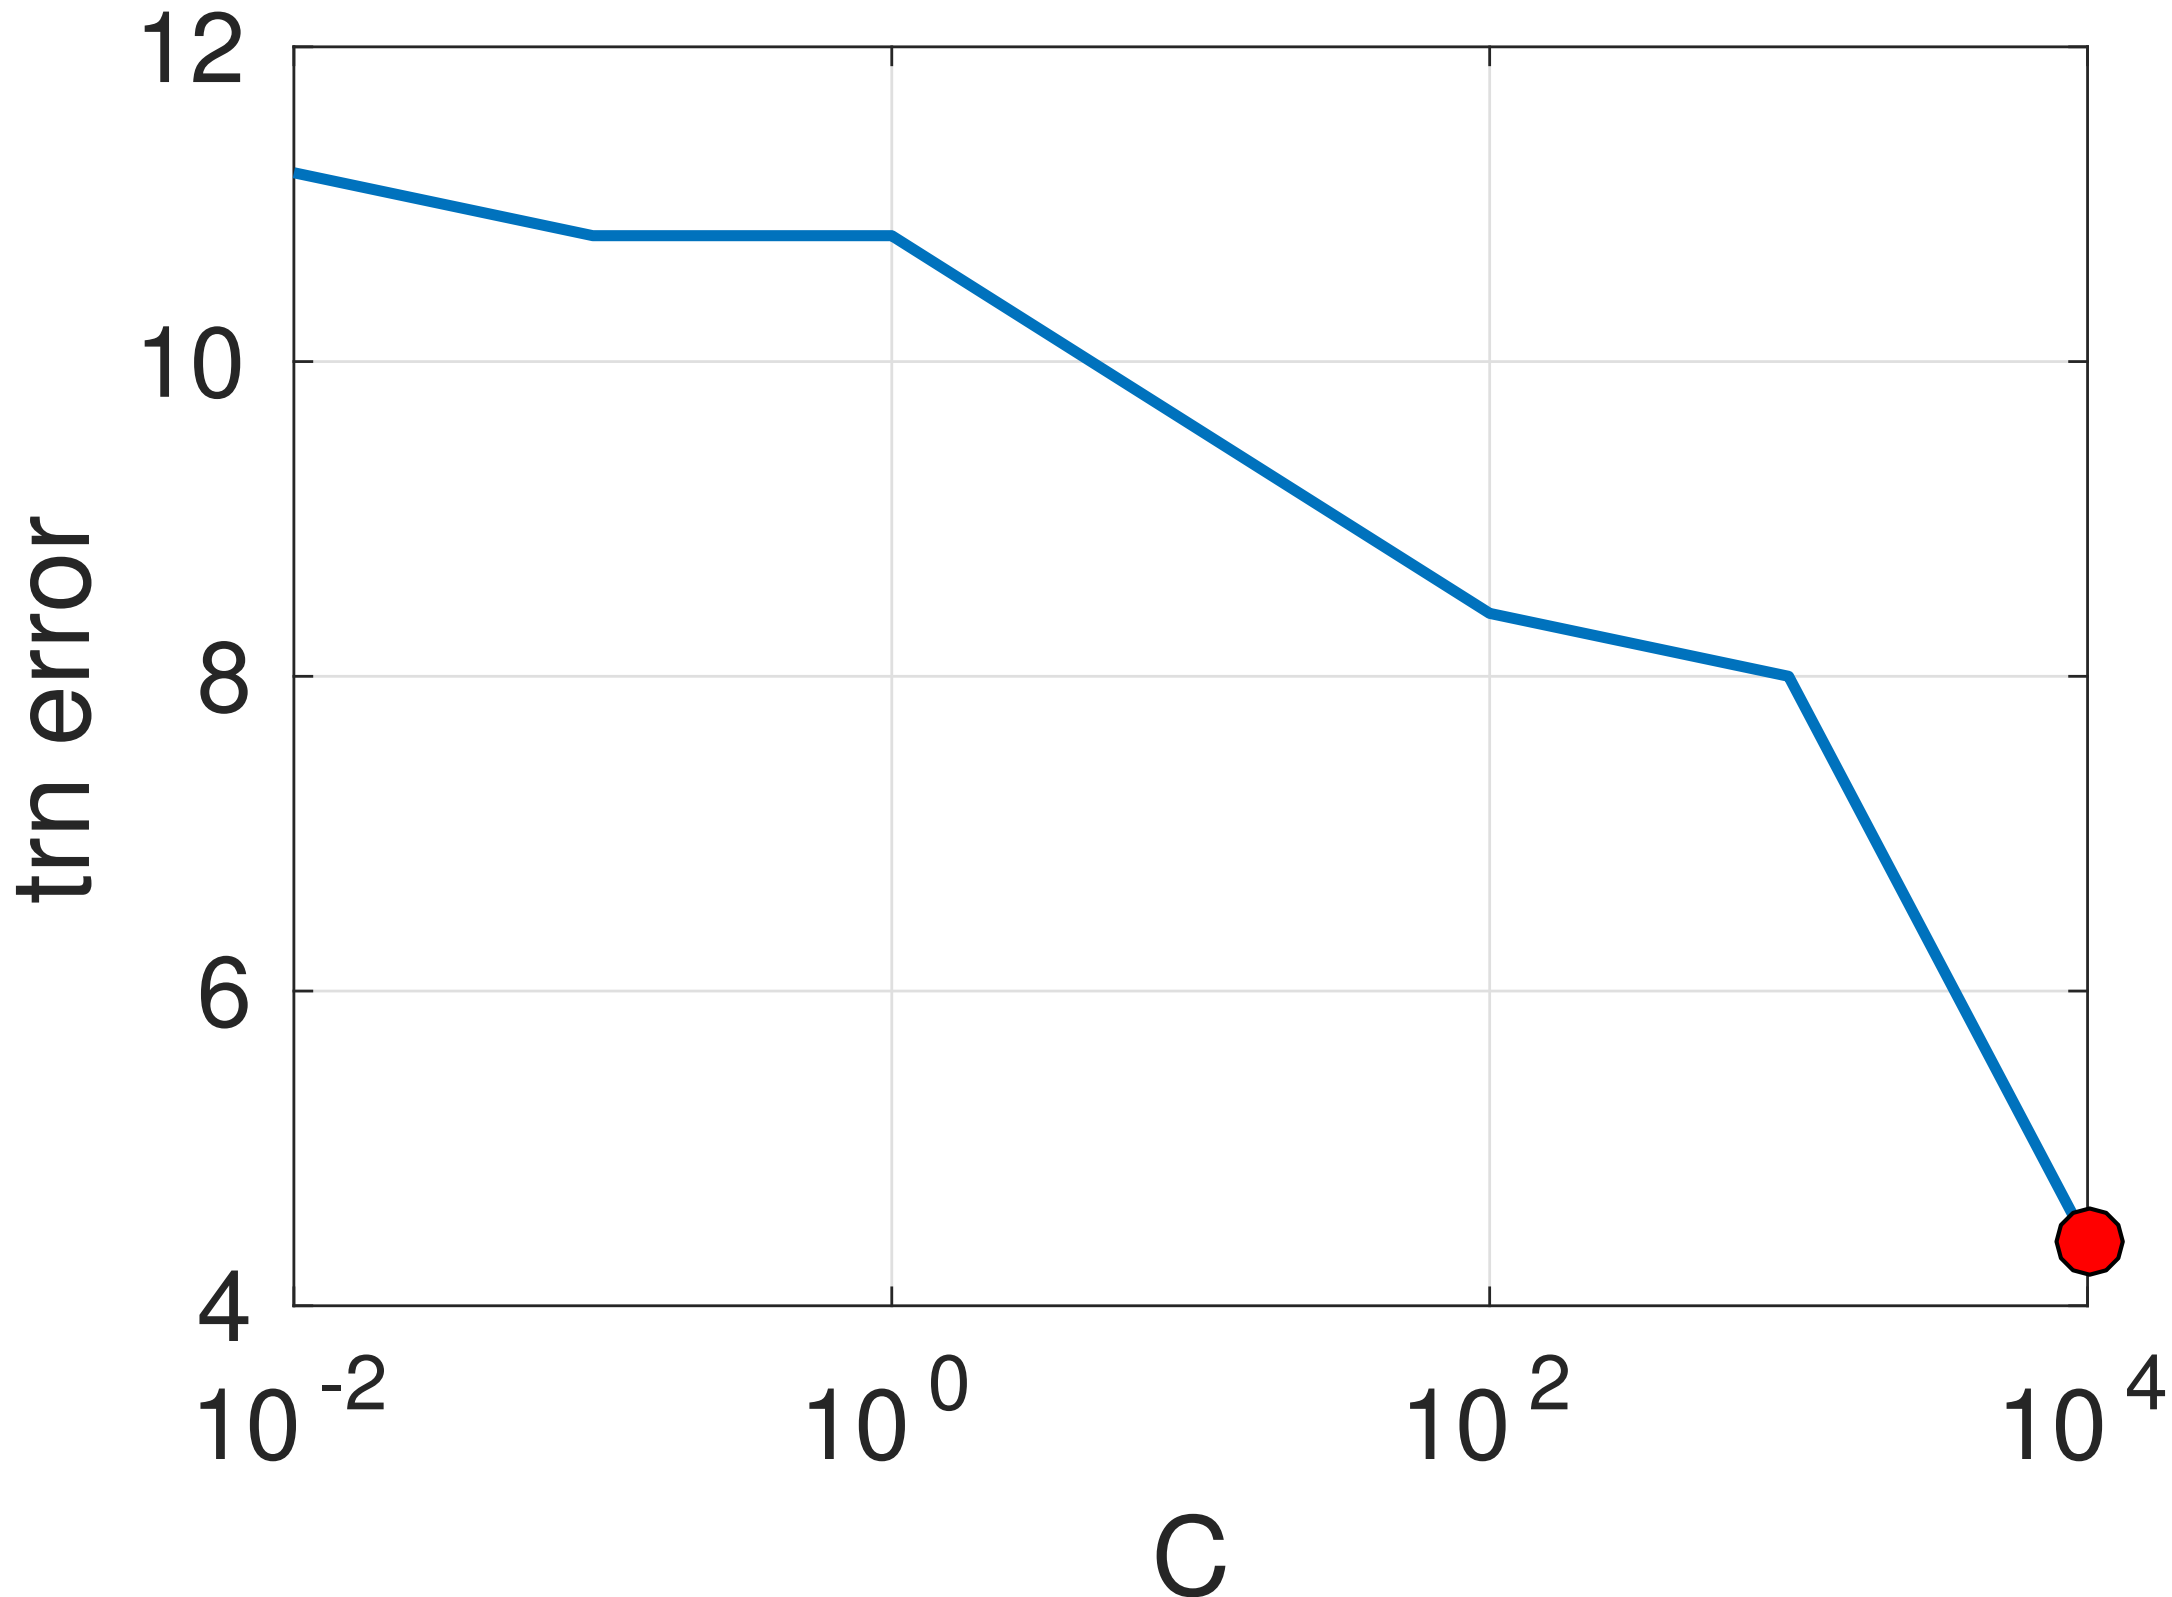


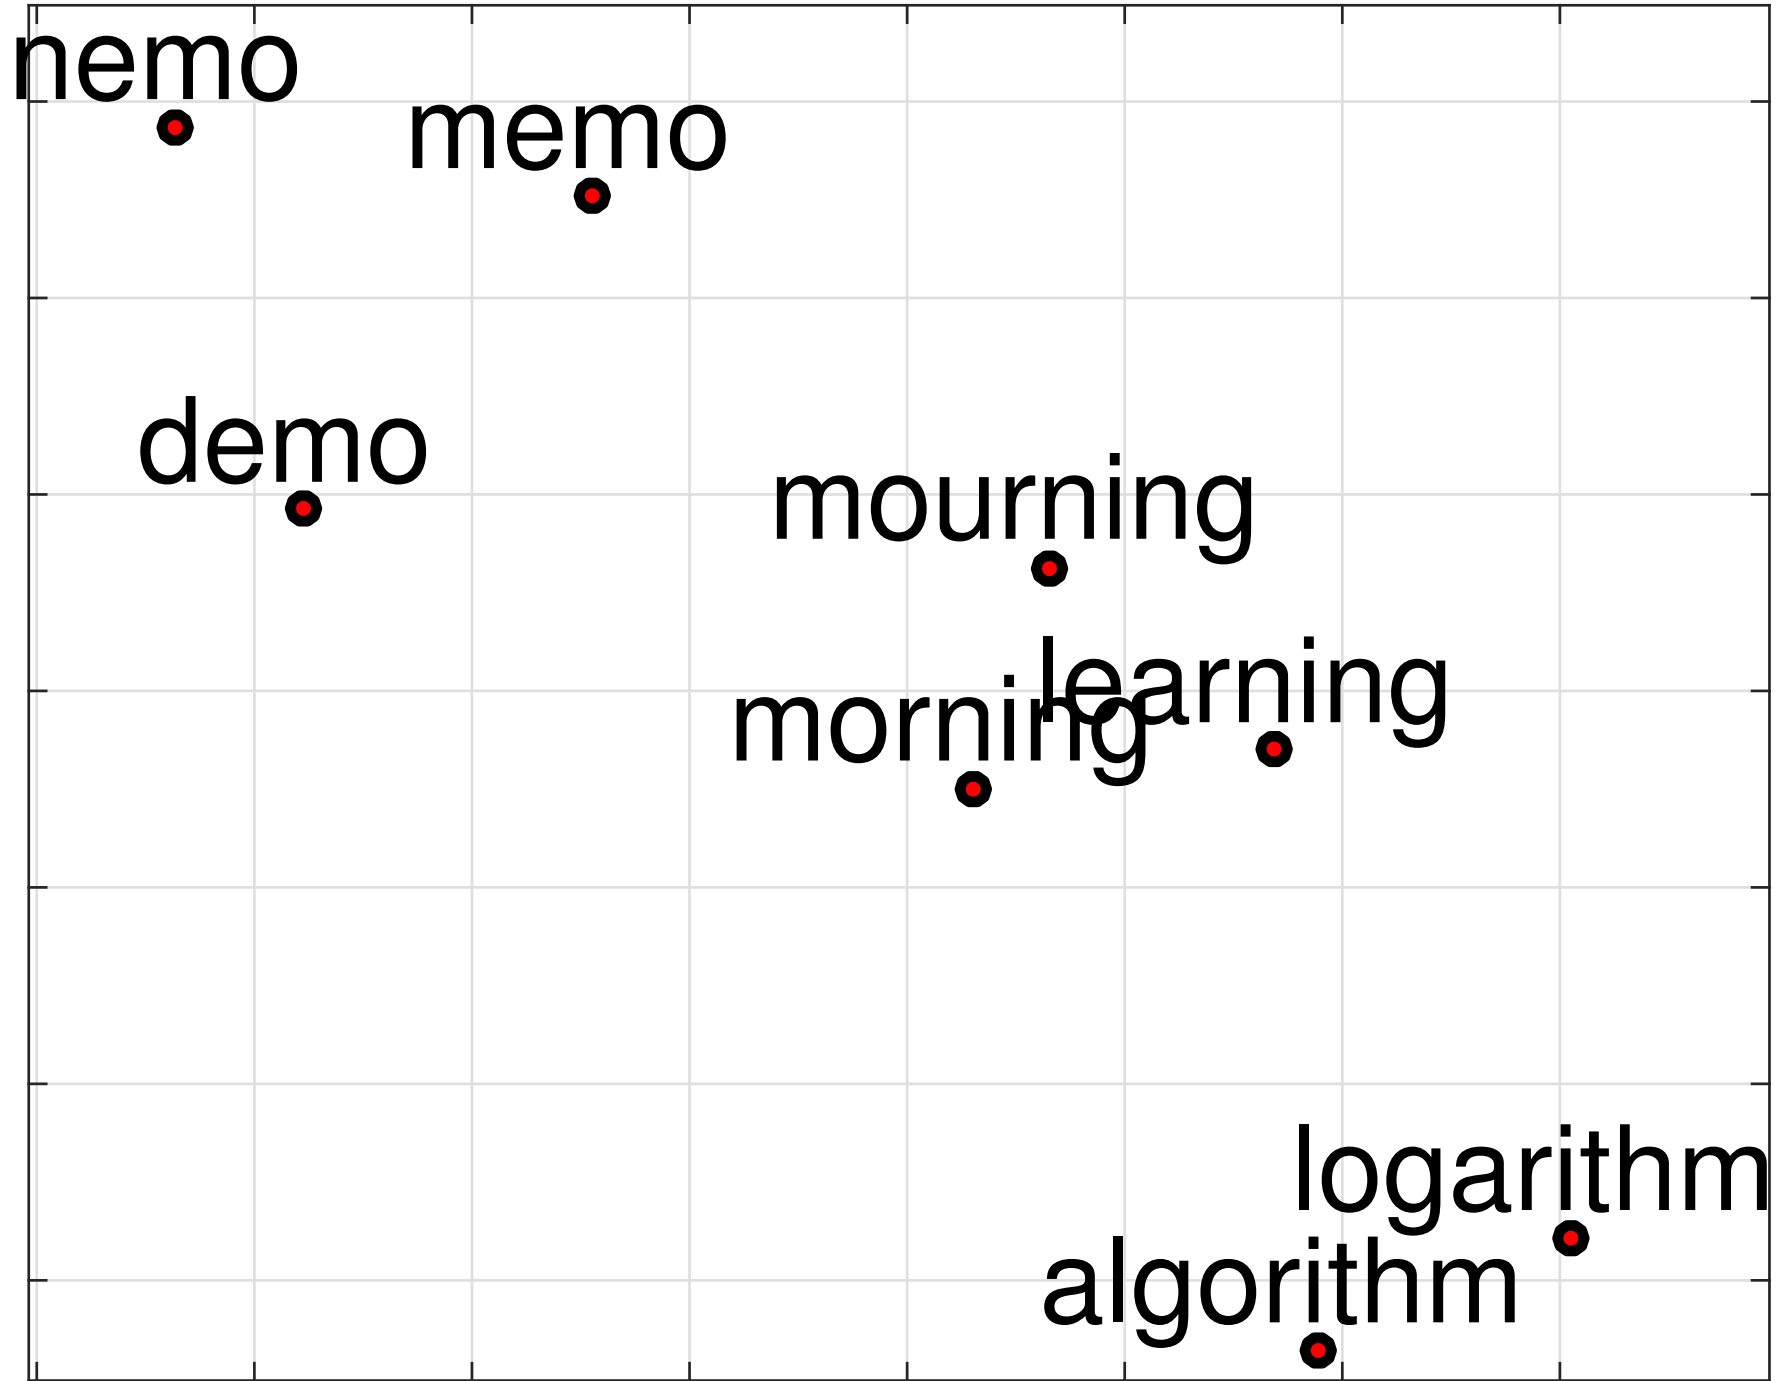
err=8.0%





err=4.4%





algorithm	0.24	0.15	0.01	0.04	0.02	0.00	0.00	0.00
logarithm	0.15	0.24	0.04	0.02	0.01	0.00	0.00	0.00
learning	0.01	0.04	0.23	0.13	0.13	0.00	0.00	0.00
morning	0.04	0.02	0.13	0.19	0.17	0.03	0.03	0.03
mourning	0.02	0.01	0.13	0.17	0.23	0.03	0.03	0.03
demo	0.00	0.00	0.00	0.03	0.03	0.09	0.06	0.06
memo	0.00	0.00	0.00	0.03	0.03	0.06	0.09	0.06
nemo	0.00	0.00	0.00	0.03	0.03	0.06	0.06	0.09

algorithm
 logarithm
 learning
 morning
 mourning
 demo
 memo
 nemo