Regularized Linear Regression

Mahdi Roozbahani
Georgia Tech

These slides are inspired based on slides from Andrew Zisserman, Jonathan Taylor, Chao Zhang and Yaser Abu-Mostafa.


EVERY GROUP PROJECT
DOES 99% OF THE WORK
HAS NO IDEA WHAT'S GOING ON THE WHOLE TIME
SAYS HE'S GOING TO HELP BUT HE'S NOT
DISAPPEAR AT THE VERY BEGINNING AND DOESN'T SHOW UP AGAIN TIL THE VERY END

Outline

- Overfitting and regularized learning
- Ridge regression
- Lasso regression
- Determining regularization strength

Regression: Recap

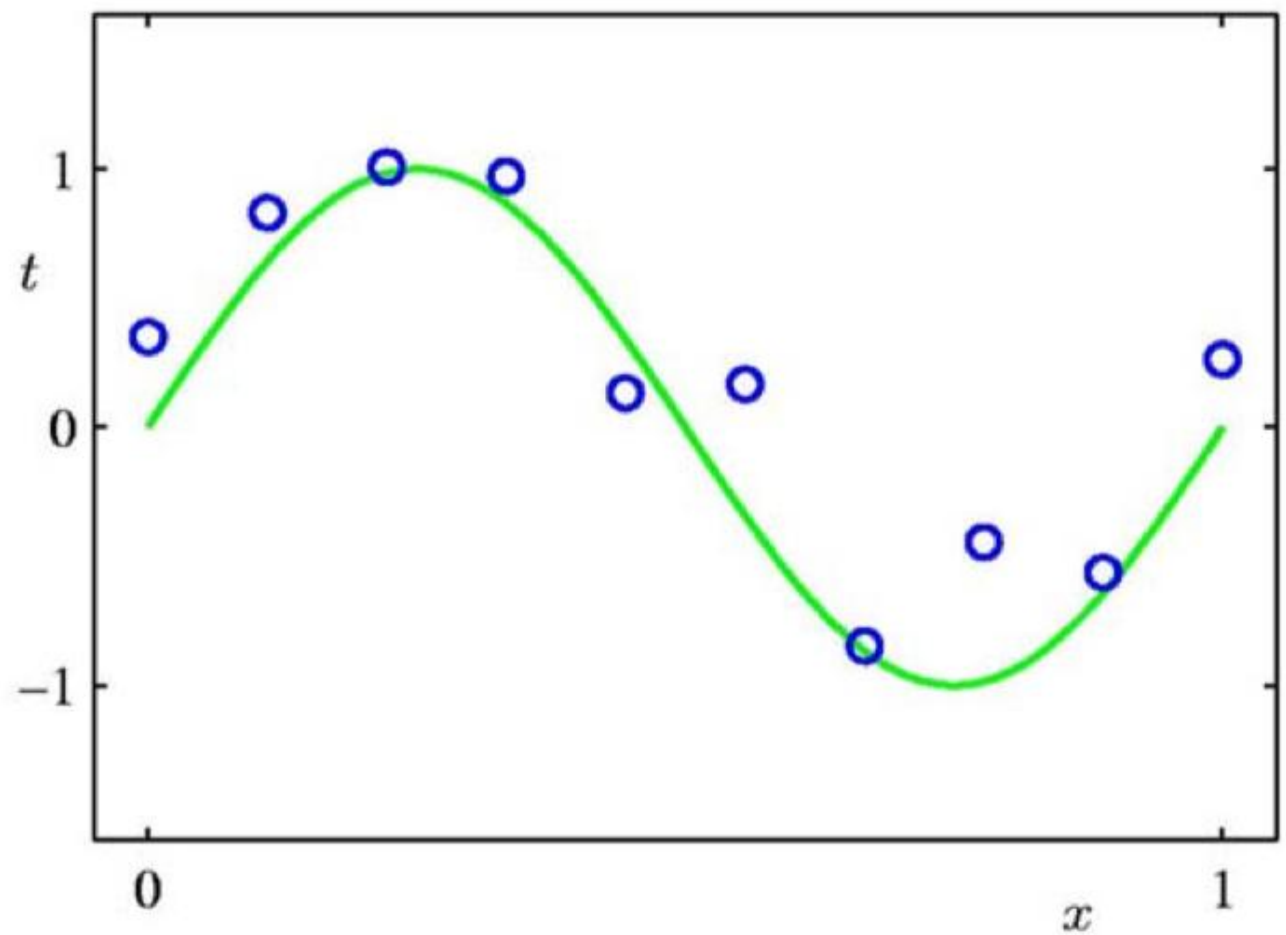

- Suppose we are given a training set of N observations

- Regression problem is to estimate y(x) from this data

Regression: Recap

$$X = \begin{array}{c} \begin{matrix} x_1 & x_2 \end{matrix} \\ \begin{bmatrix} \quad & \quad \end{bmatrix} \end{array}$$

$$Z = \begin{array}{c} \begin{matrix} x_1 & x_2 & x_1^2 & x_2^2 & x_1x_2^2 & x_1^2x_2^2 & x_1^2x_2^2 & \cdots \end{matrix} \\ \begin{bmatrix} \quad & \quad \end{bmatrix} \end{array}$$

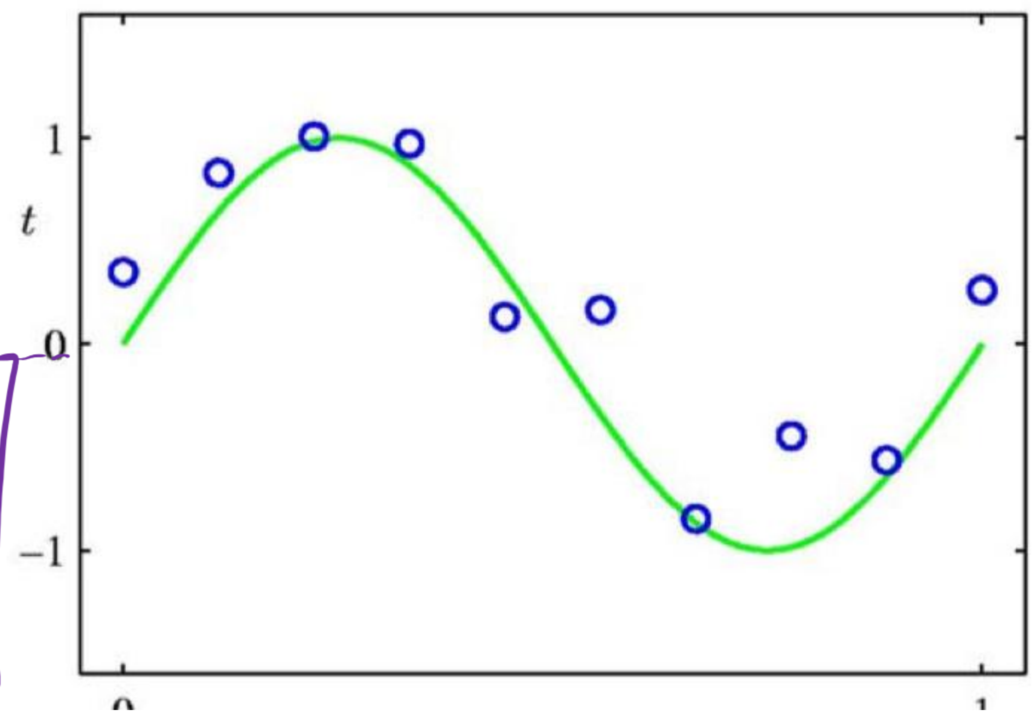

- Want to fit a polynomial regression model

$$y = \theta_0 + \theta_1 x + \theta_2 x^2 + \cdots + \theta_\mathrm{d} x^\mathrm{d} + \epsilon$$

- $z = \{1, x, x^2, \ldots, x^d\} \in R^d$ and $\theta = (\theta_0, \theta_1, \theta_2, \ldots, \theta_\mathrm{d})^\mathrm{T}$

$$y = z\theta$$

Which One is Better?

from Bishop

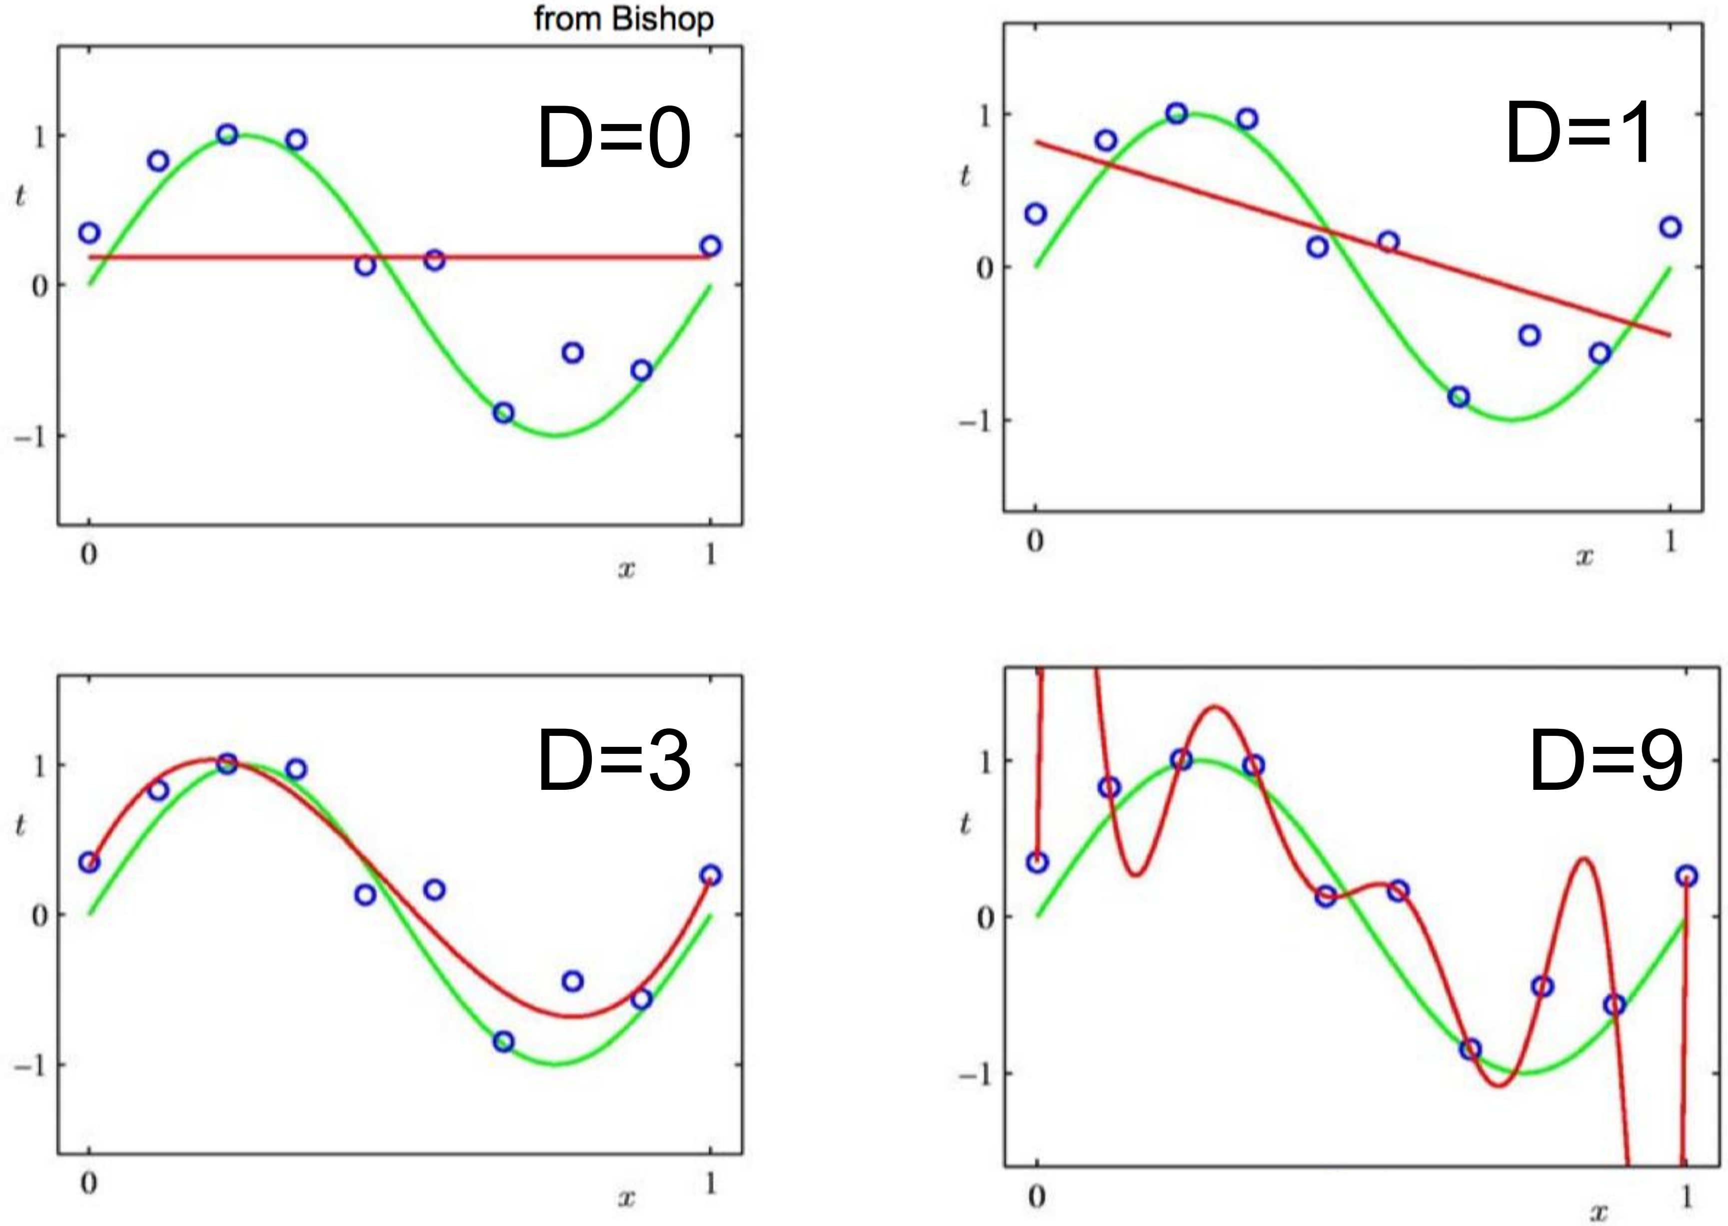

- Can we increase the maximal polynomial degree to very large, such that the curve passes through all training points?

No, this can lead to **overfitting**!

The Overfitting Problem

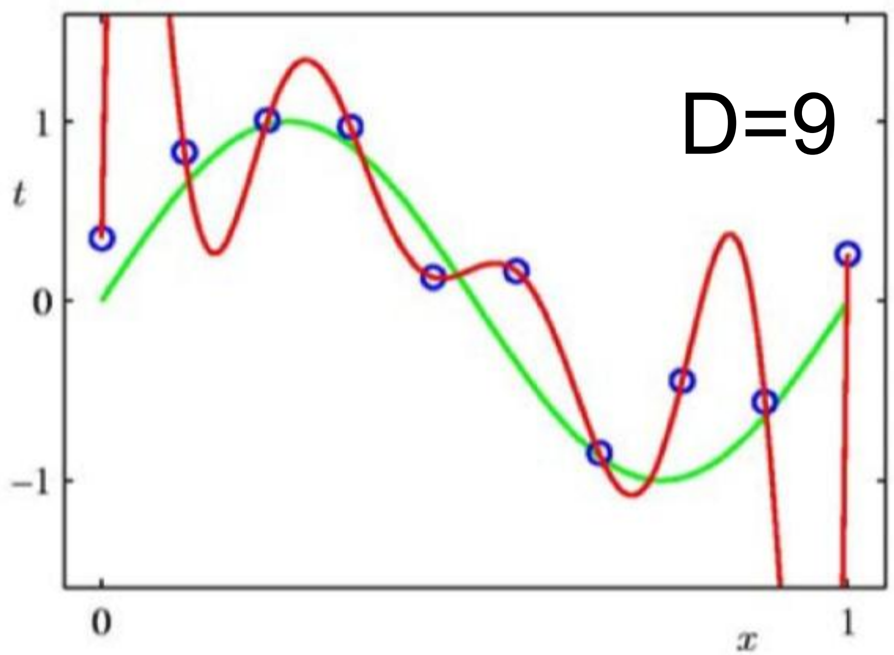

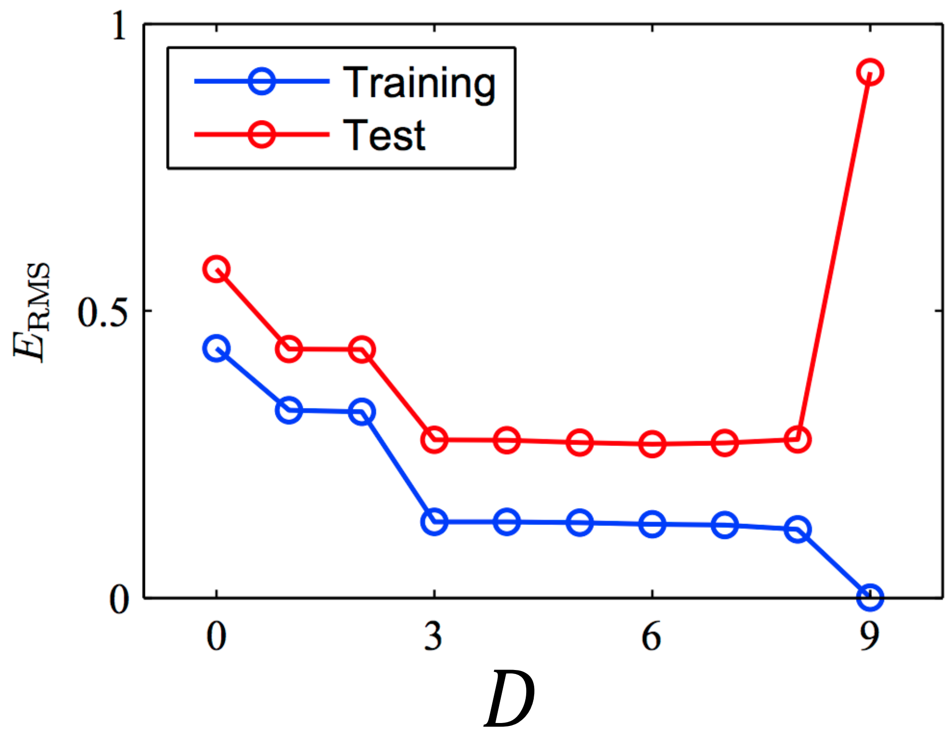

- The training error is very low, but the error on test set is large.
- The model captures not only patterns but also noisy nuisances in the training data.

The Overfitting Problem

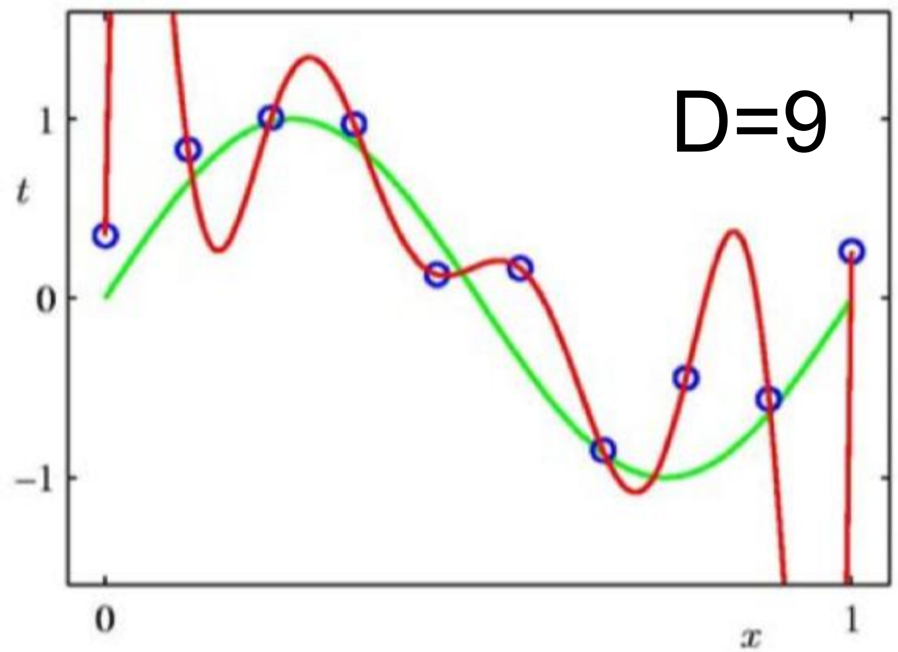

- In regression, overfitting is often associated with large Weights (severe oscillation)
- How can we address overfitting?

$$\hat{y}_p = \Theta_0 + \Theta_1 x + \Theta_2 x^2 + \Theta_2 x^3 + \cdots$$

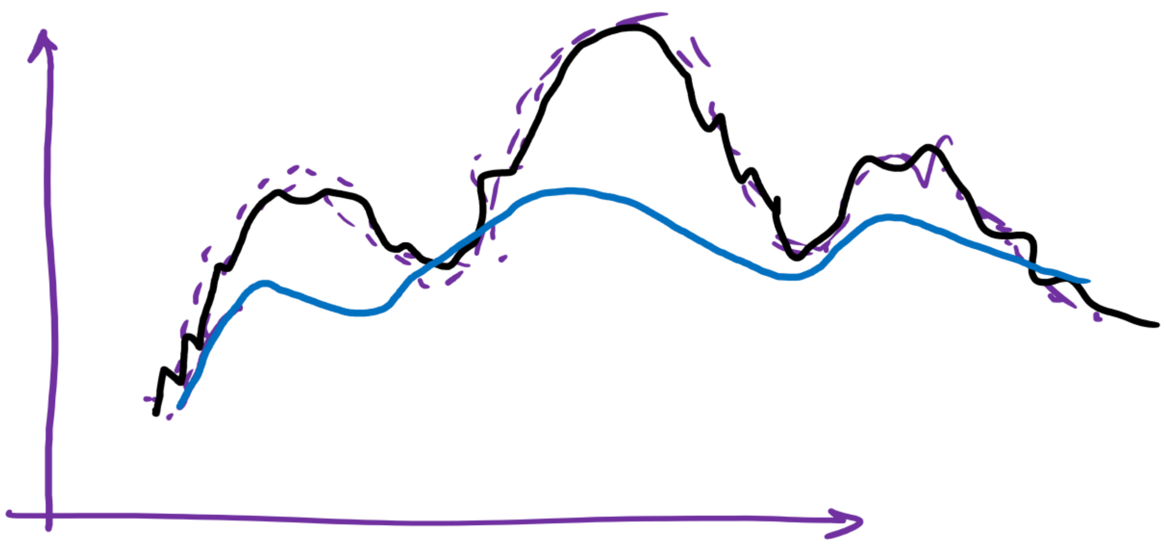

Regularization

(smart way to cure overfitting disease)

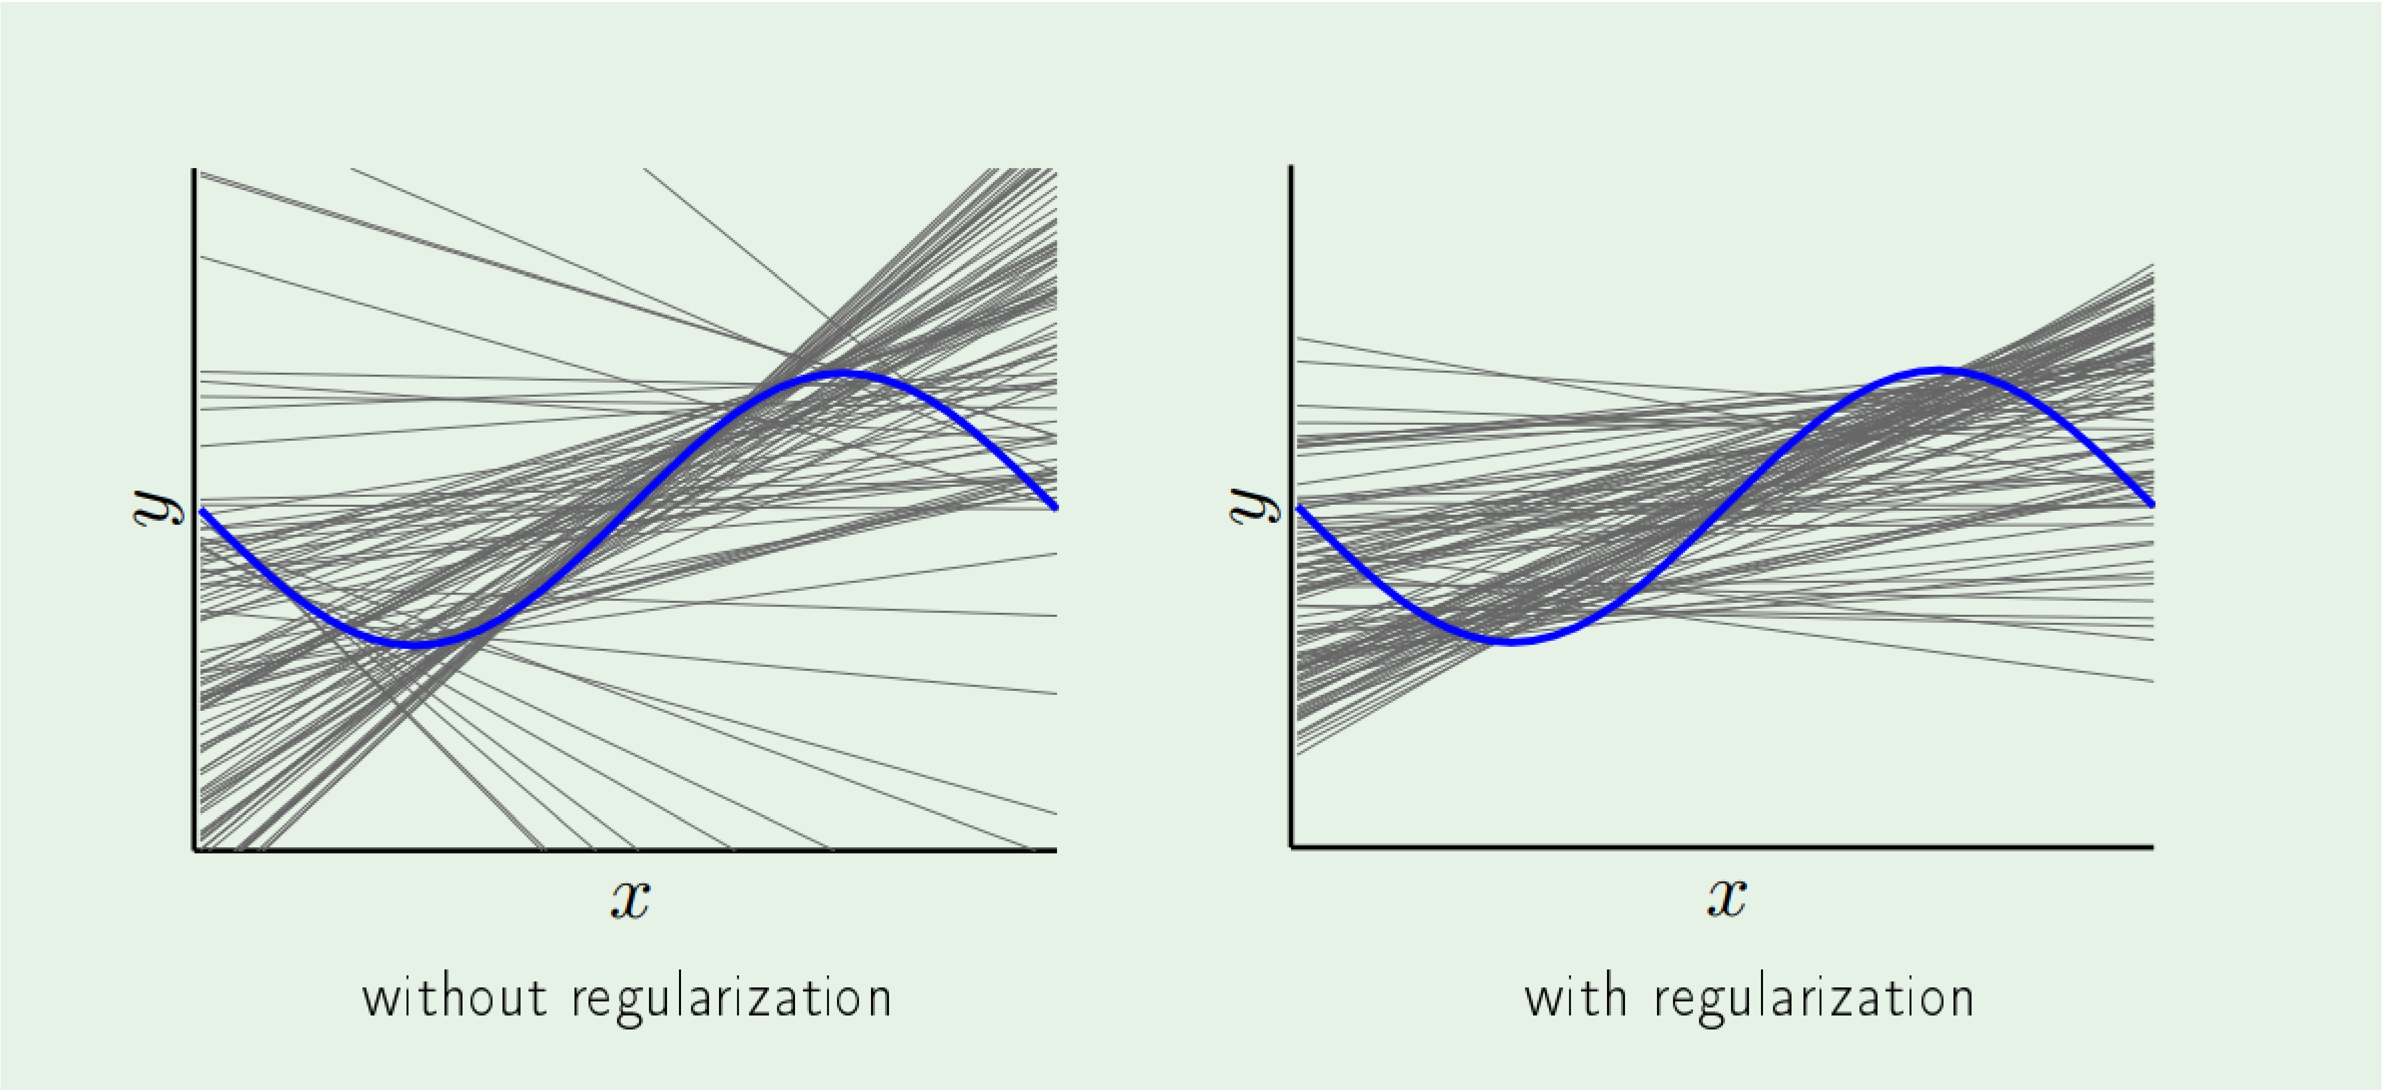

Put a brake on fitting

Fit a linear line on sinusoidal with just two data points

Who is the winner?

$E(f) = \text{bias}^2 + \text{Variance}$

$\bar{g}(x)$: average over all lines

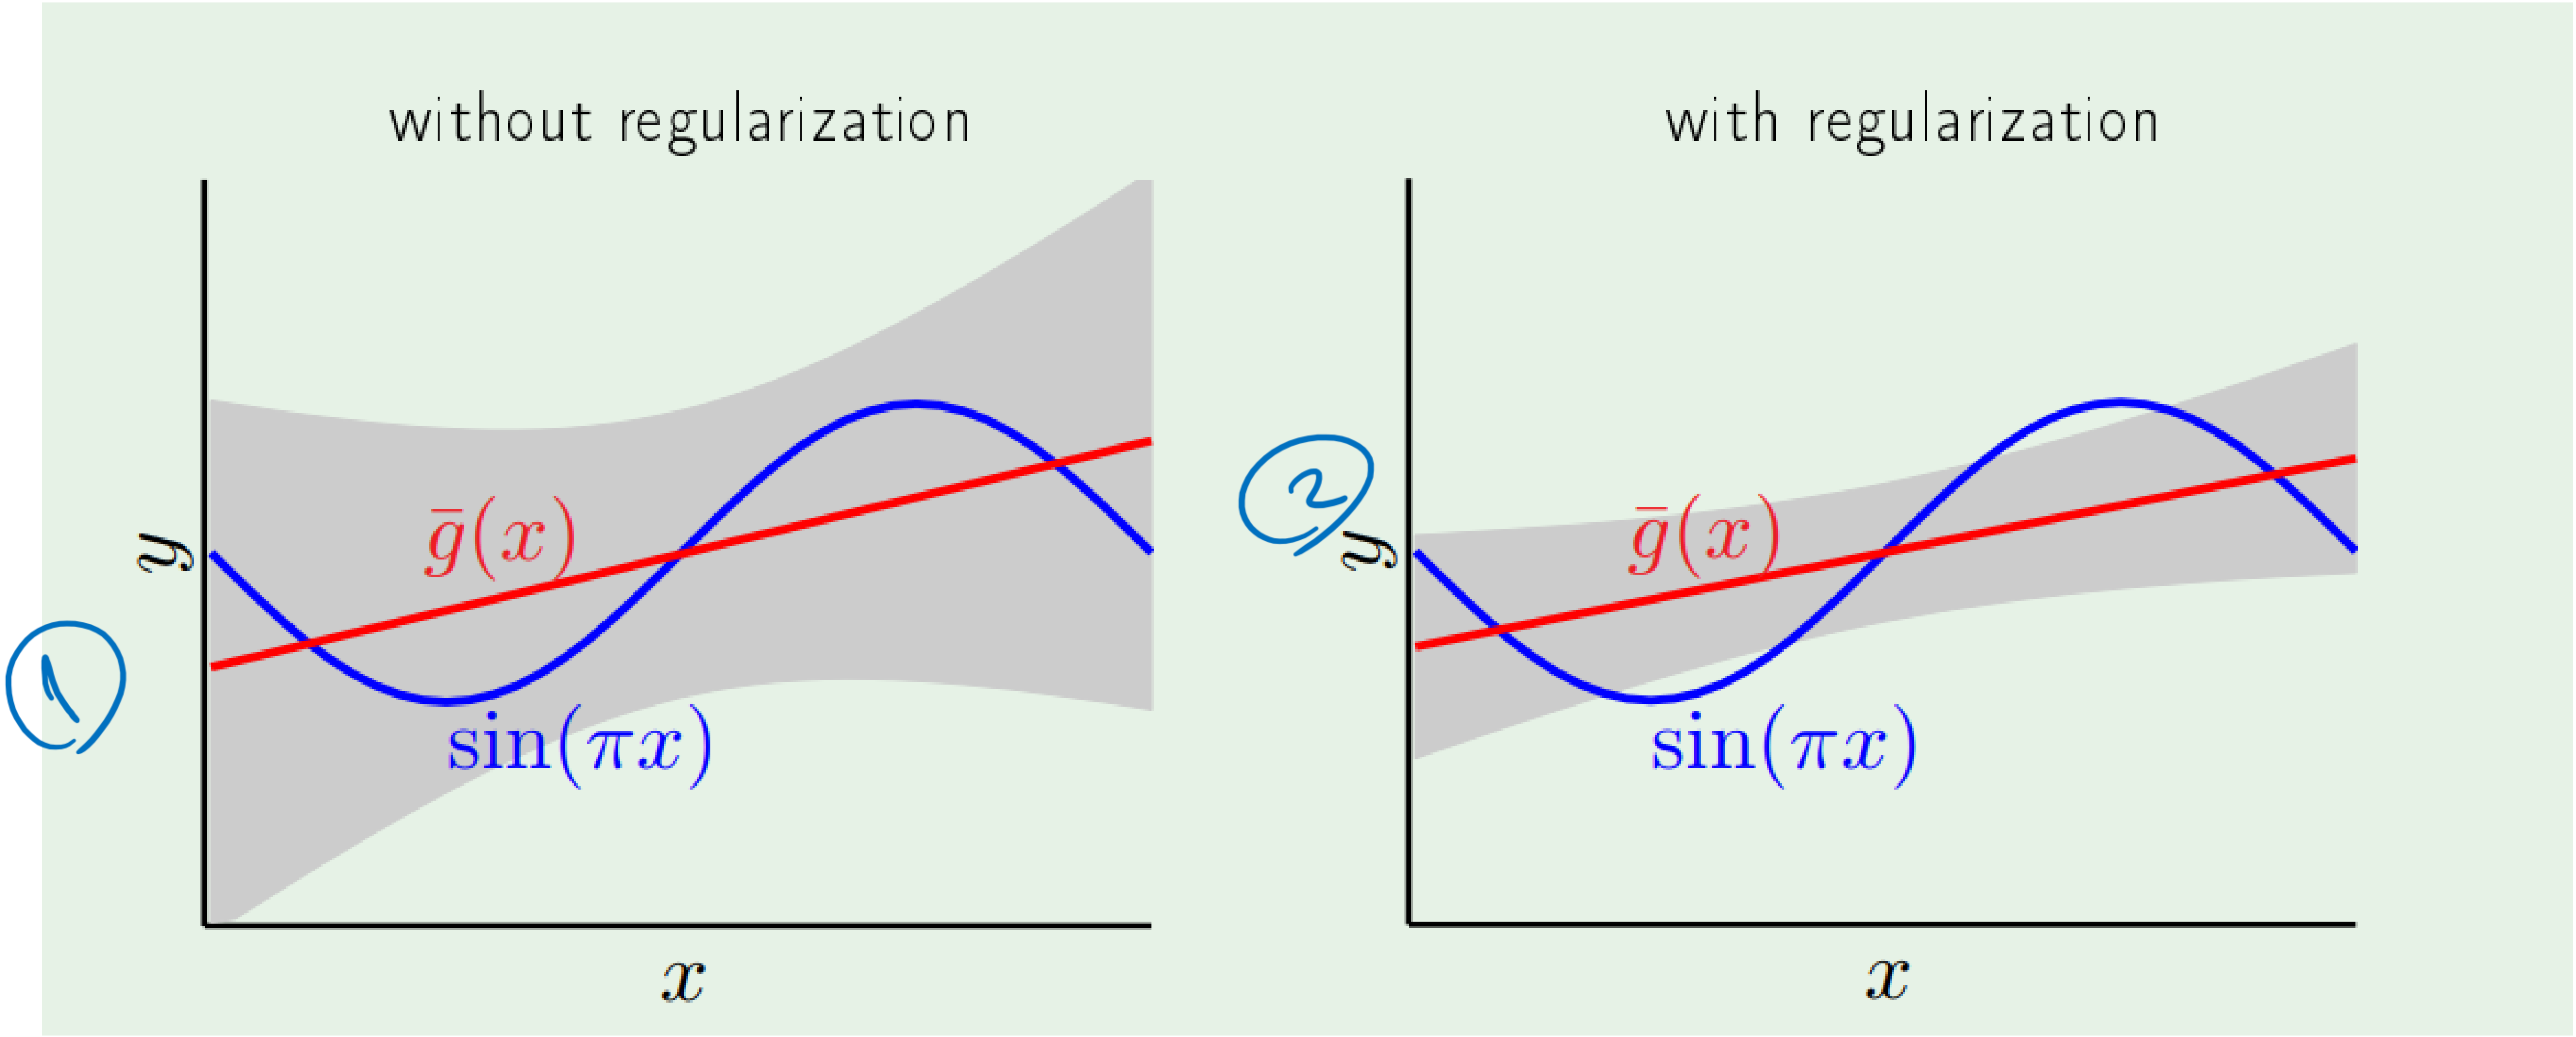

bias=0.21; var=1.69

bias=0.23; var=0.33

Polynomial Model

Want to fit a polynomial regression model

$$y = \theta_0 + \theta_1 x + \theta_2 x^2 + \cdots + \theta_d x^d + \epsilon$$

Let's rewrite it as:

$$y = \theta_0 + \theta_1 z_1 + \theta_2 z_2 + \cdots + \theta_d z_{\mathrm{d}} + \epsilon = \boldsymbol{z\theta}$$

Regularizing is just constraining the weights (*θ*)

For example: let's do a **hard** constraining

$$y = \theta_0 + \theta_1 z_1 + \theta_2 z_2 + \cdots + \theta_d z_{\mathrm{d}}$$

subject to

$$\theta_d = 0 \; for \; d > 2$$

$$\Downarrow$$

$$y = \theta_0 + \theta_1 z_1 + \theta_2 z_2 + 0 + \cdots + 0$$

$$\underbrace{E(\theta) = \frac{1}{N}\sum_{i=1}^{n}\left(y^{\{i\}} - z^{\{i\}}\theta\right)^2}$$

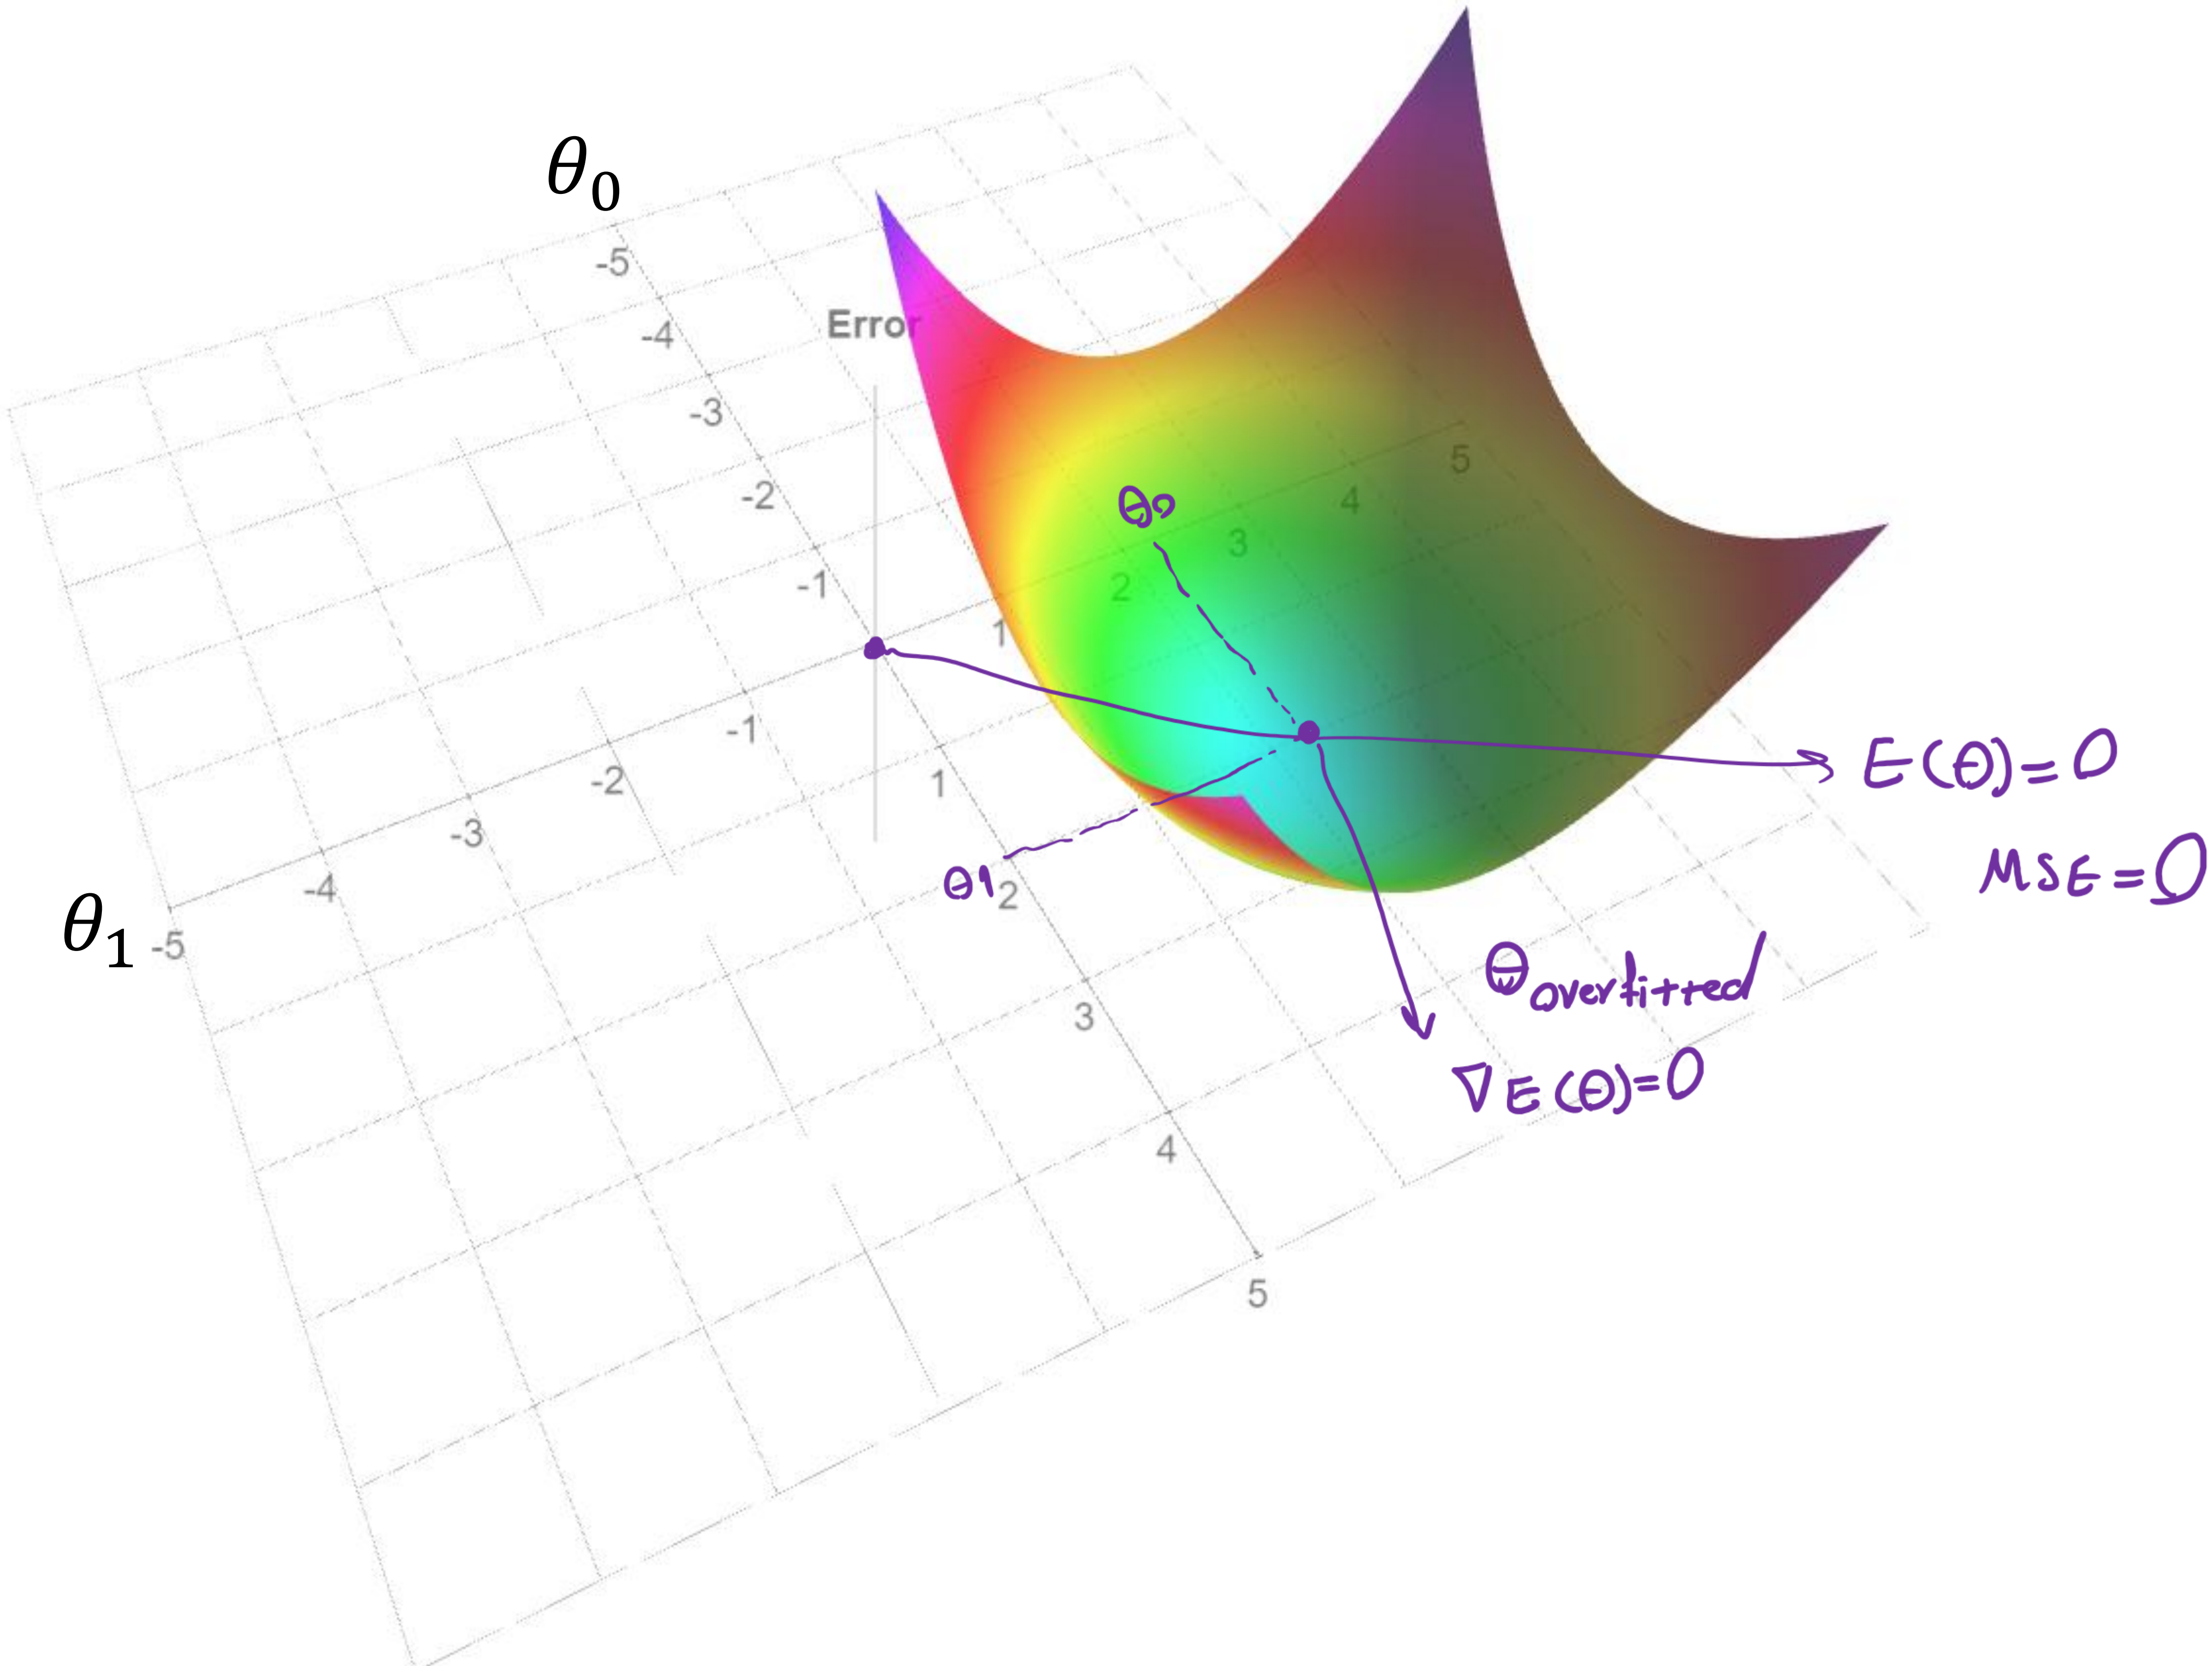

Project the same graph on x-y using contour plot

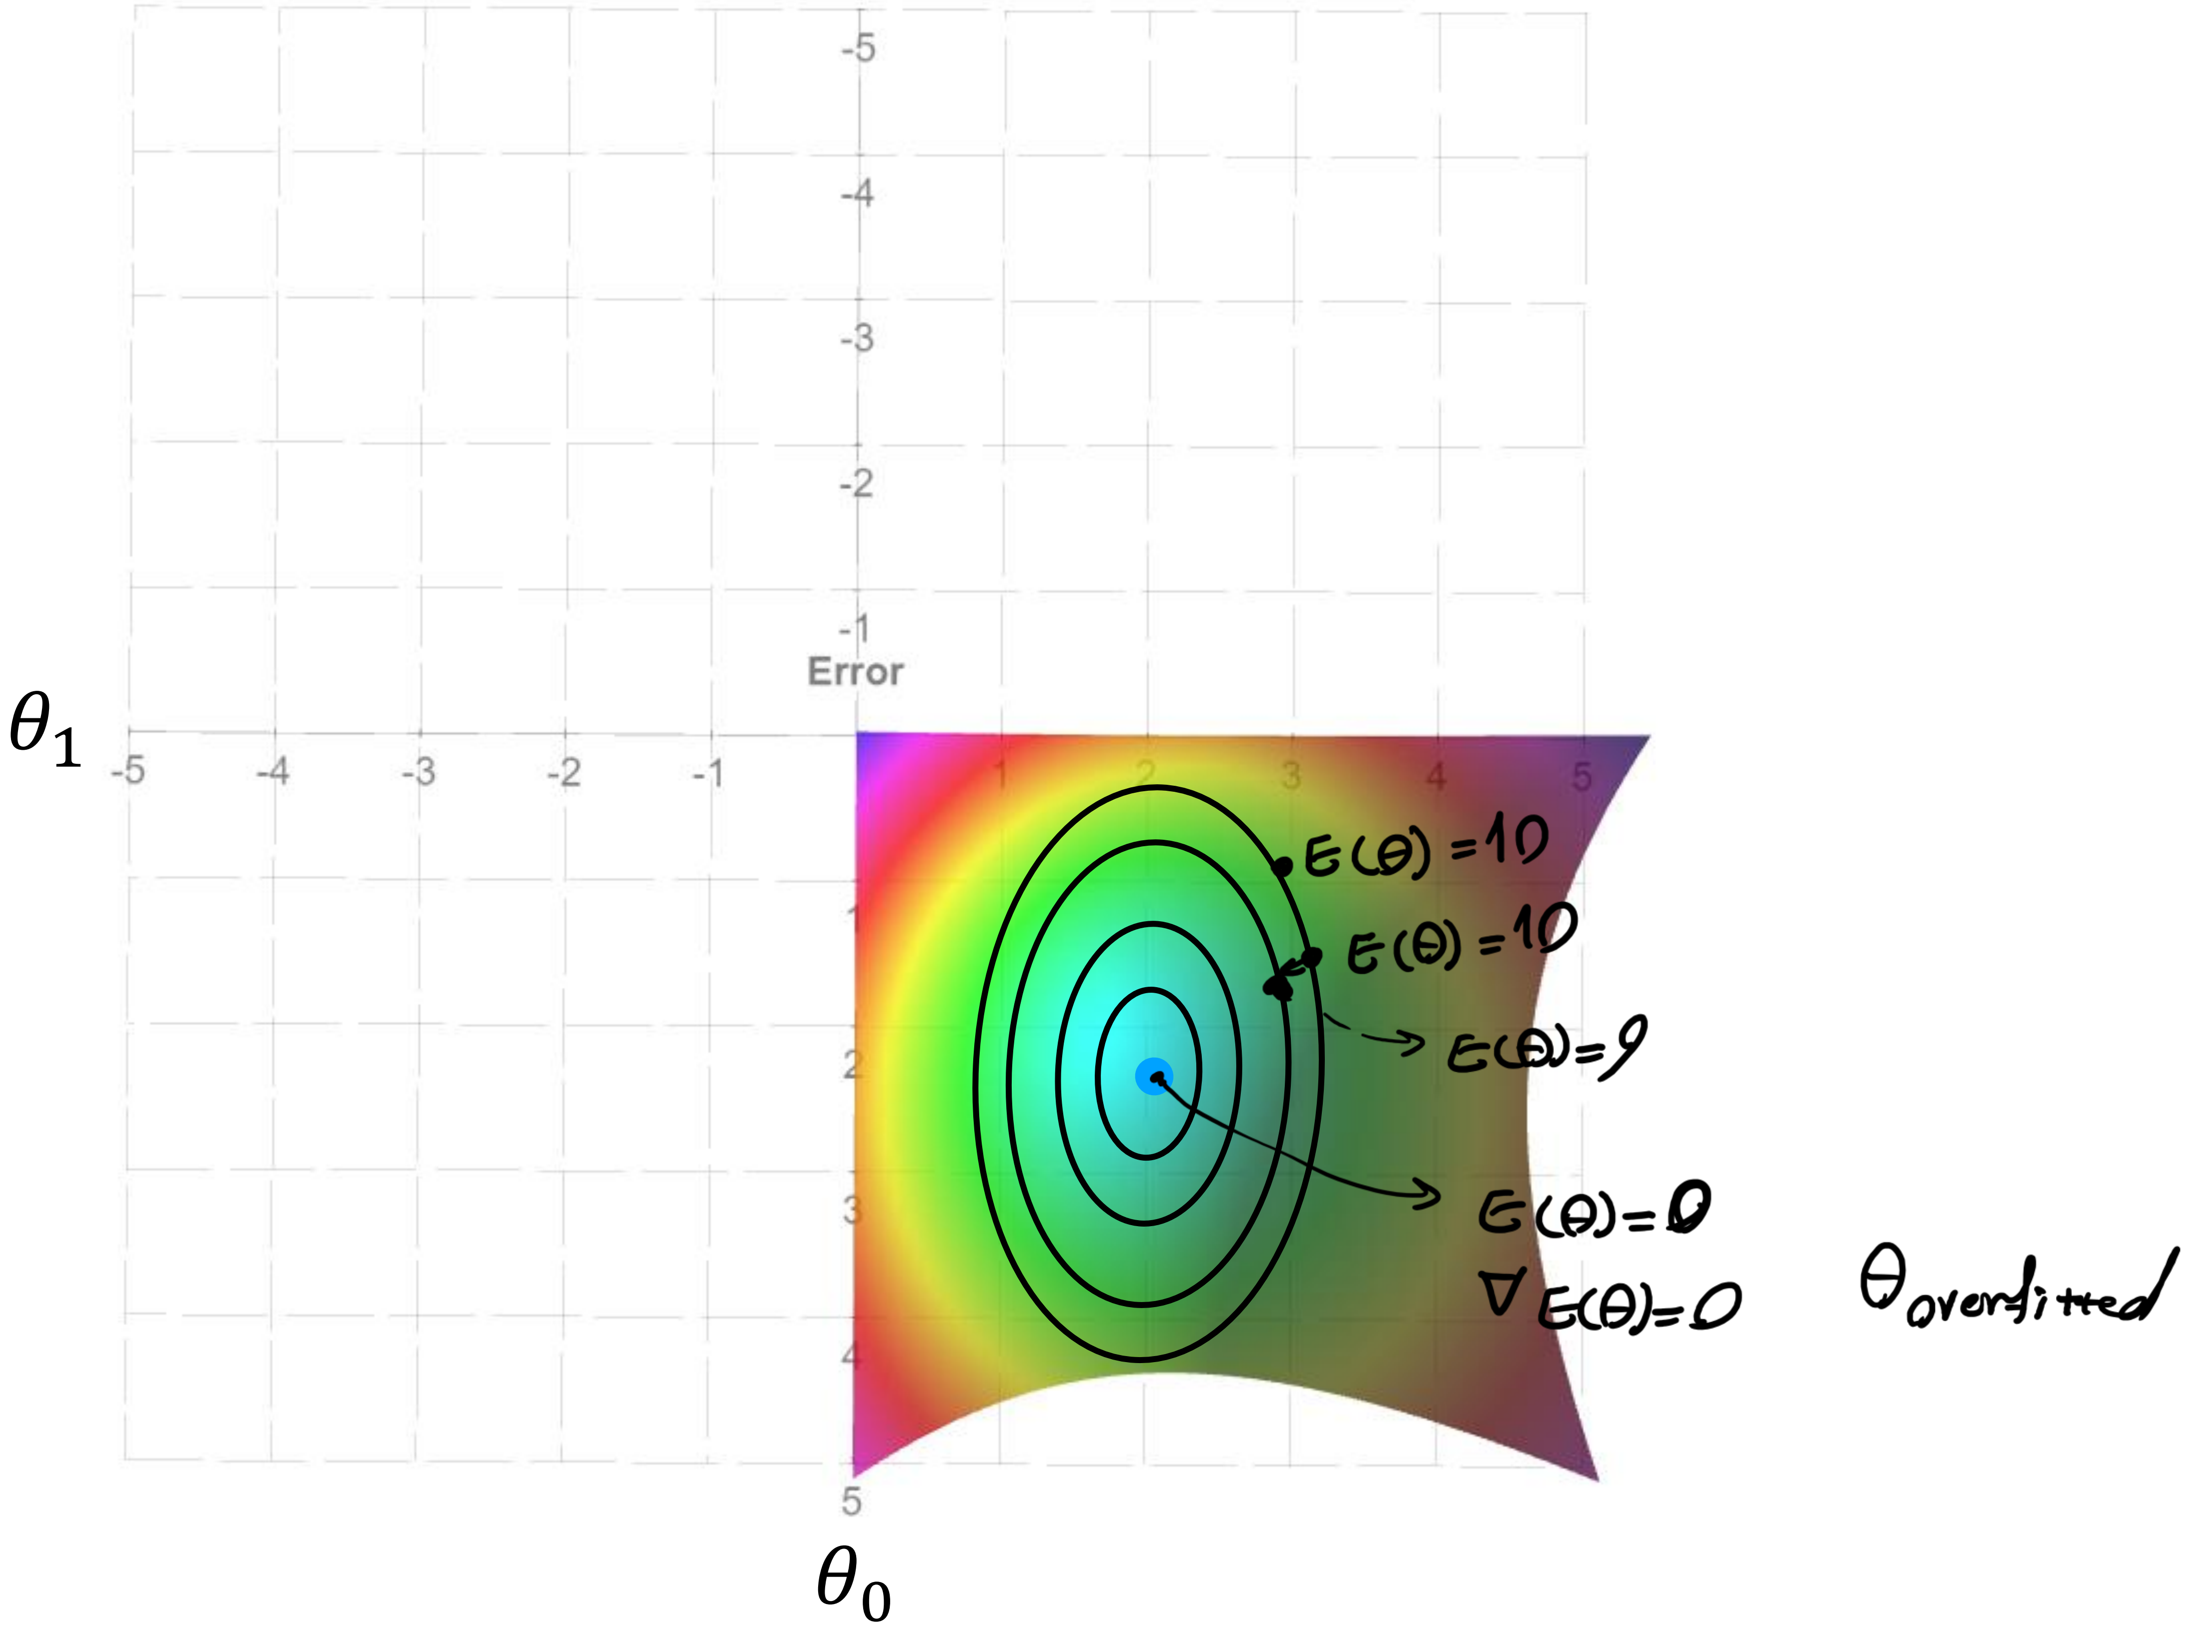

Linear regression with a very high polynomial degree solution

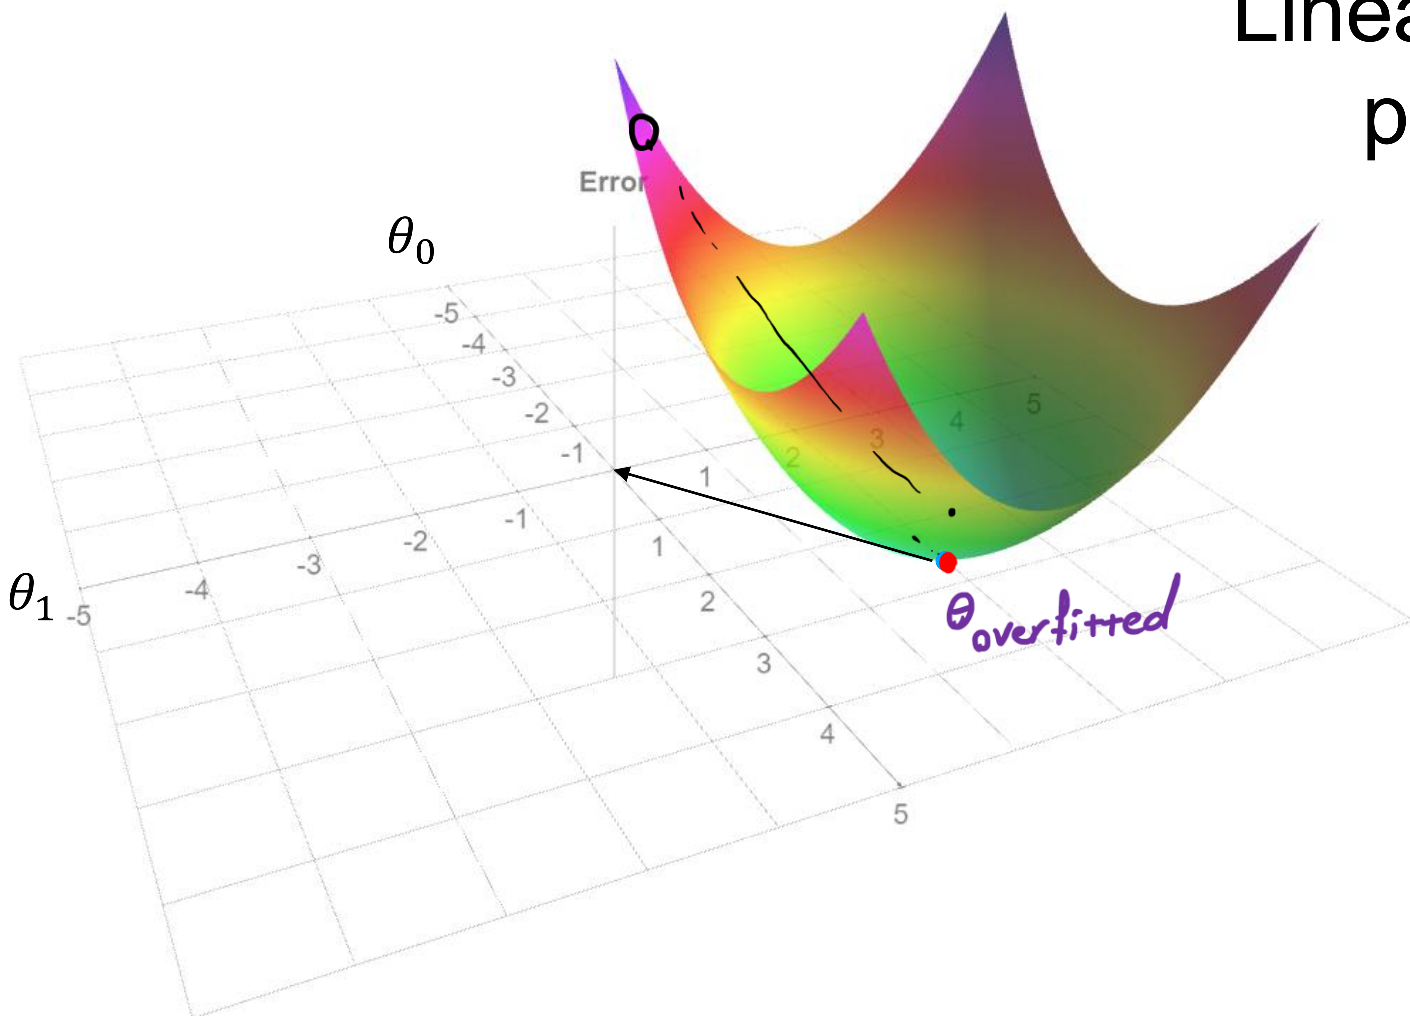

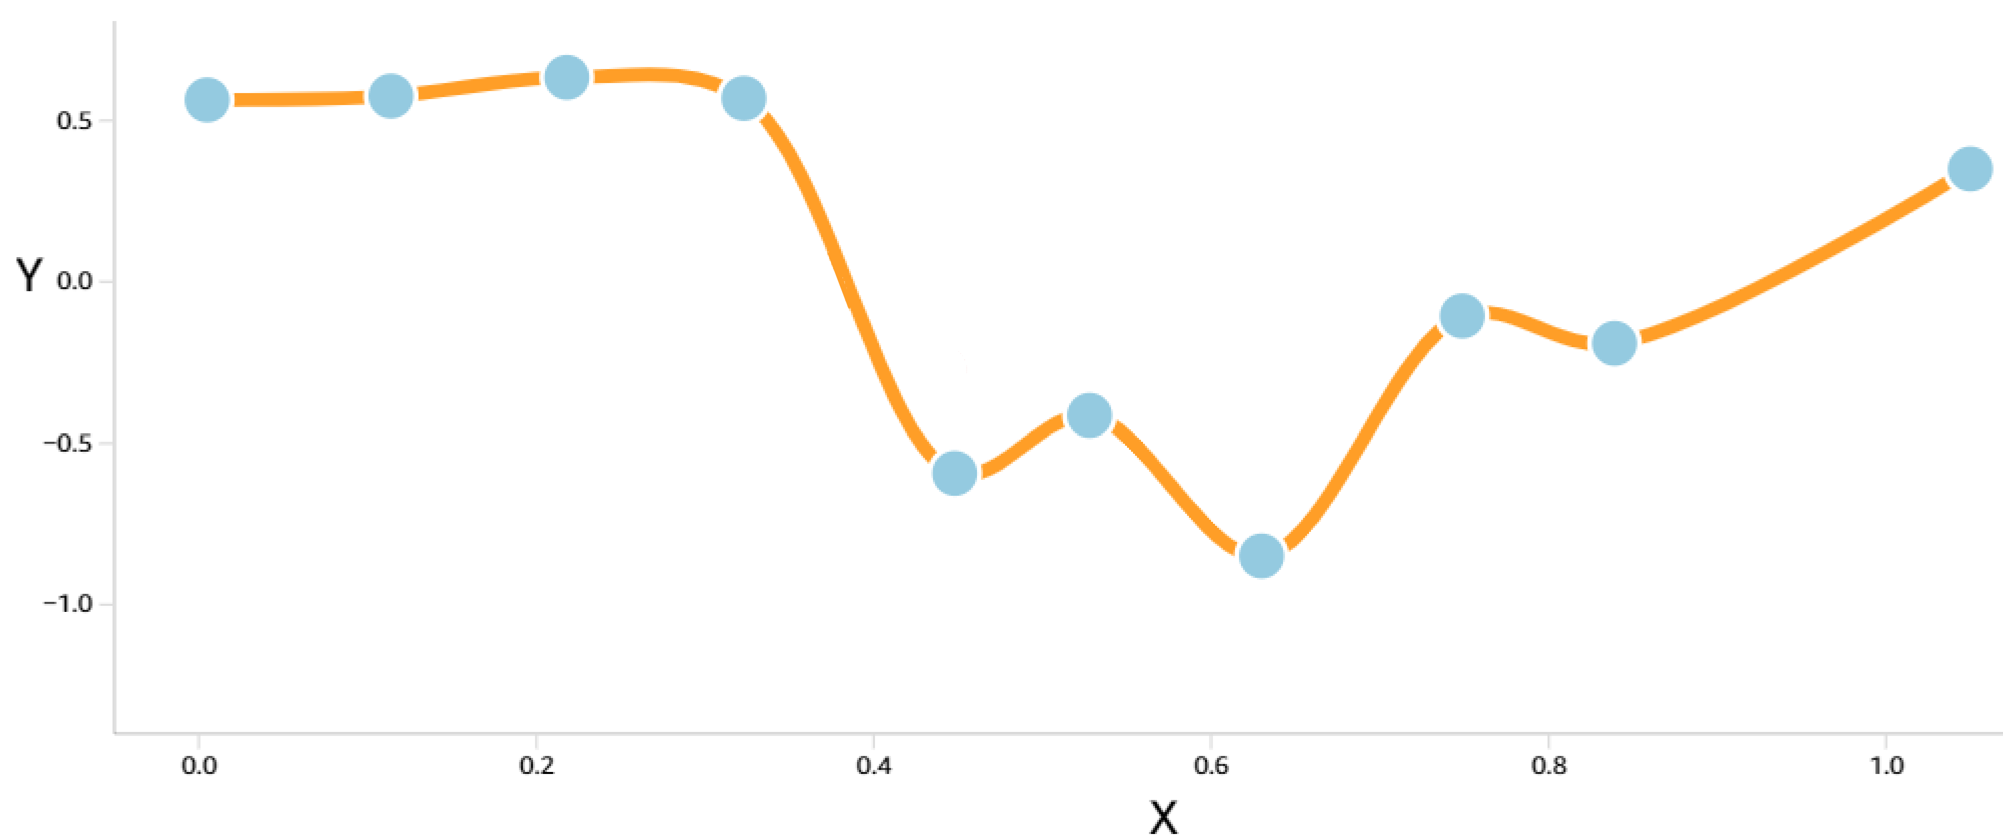

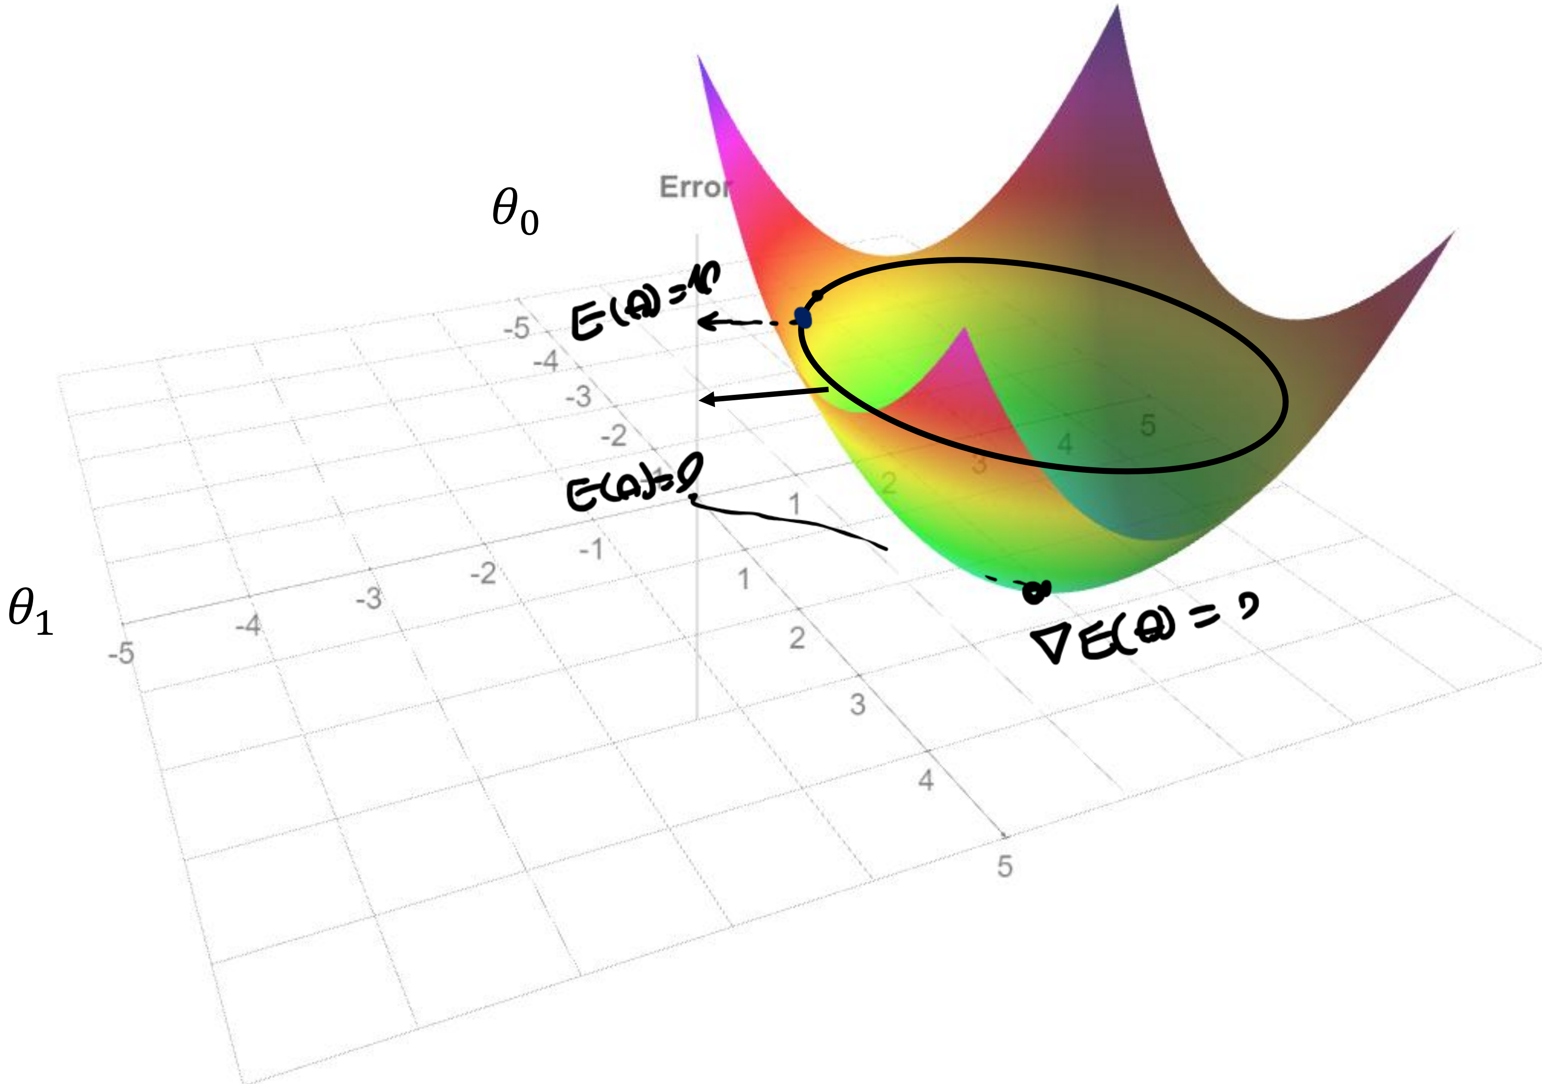

Error
θ_0
θ_1
E(θ)=10
E(θ)=0
∇E(θ) = 0

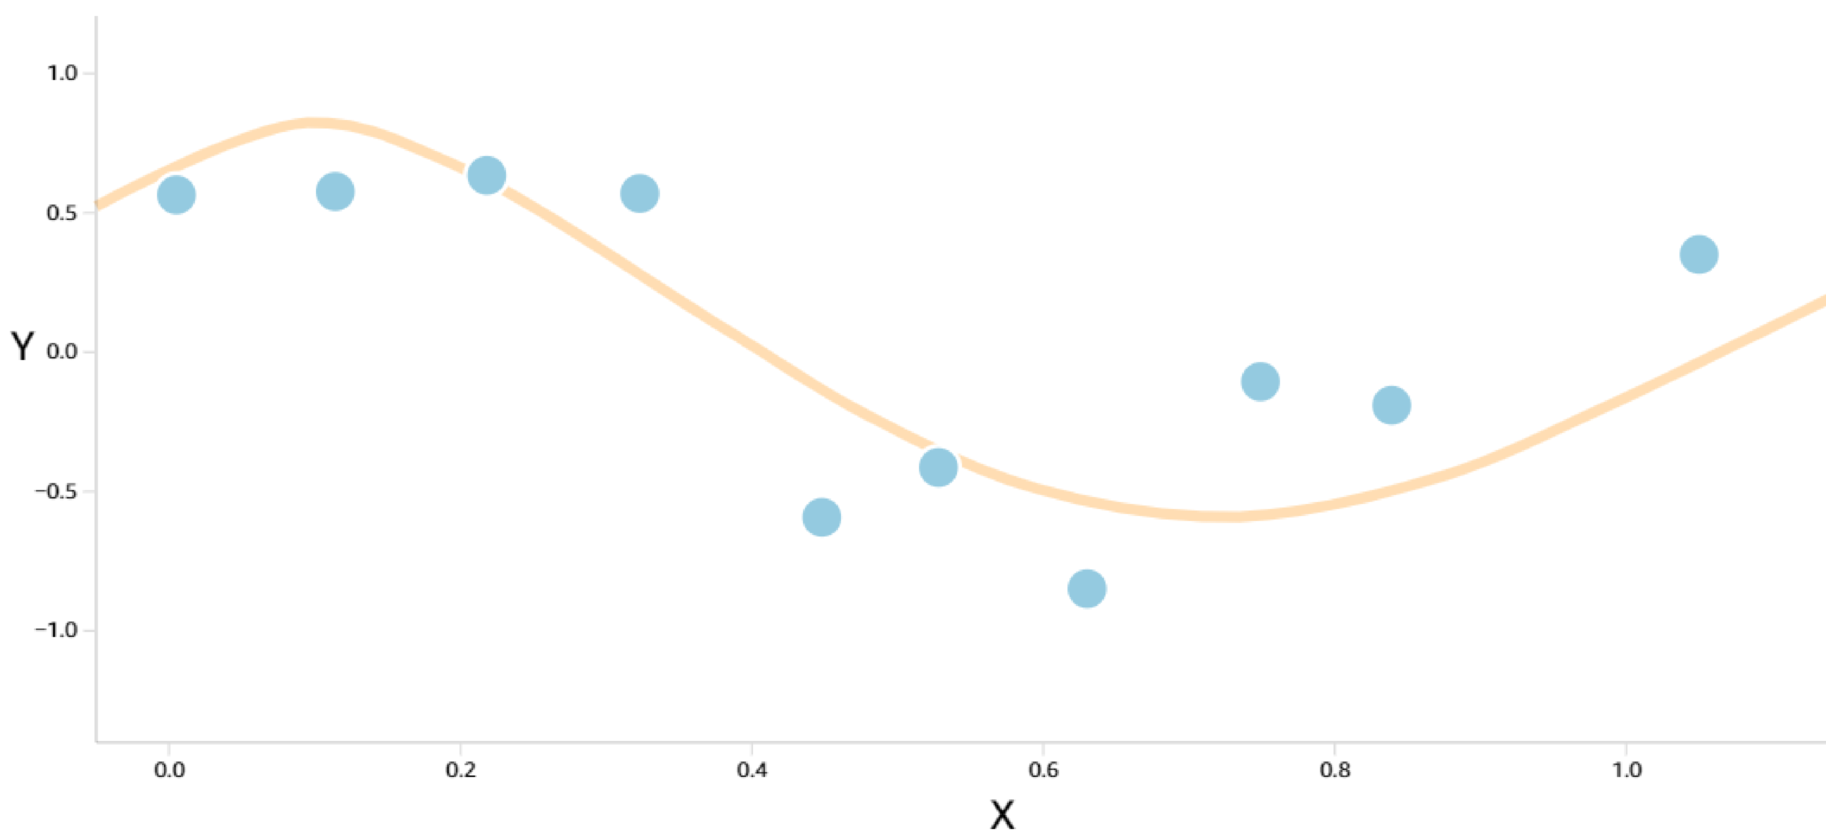

Y
X

How can we get an optimal solution with a positive error for a model that overfits?

We need to introduce a constraint

$$g(\theta) = \theta_0^2 + \theta_1^2 = \theta^T\theta = C$$

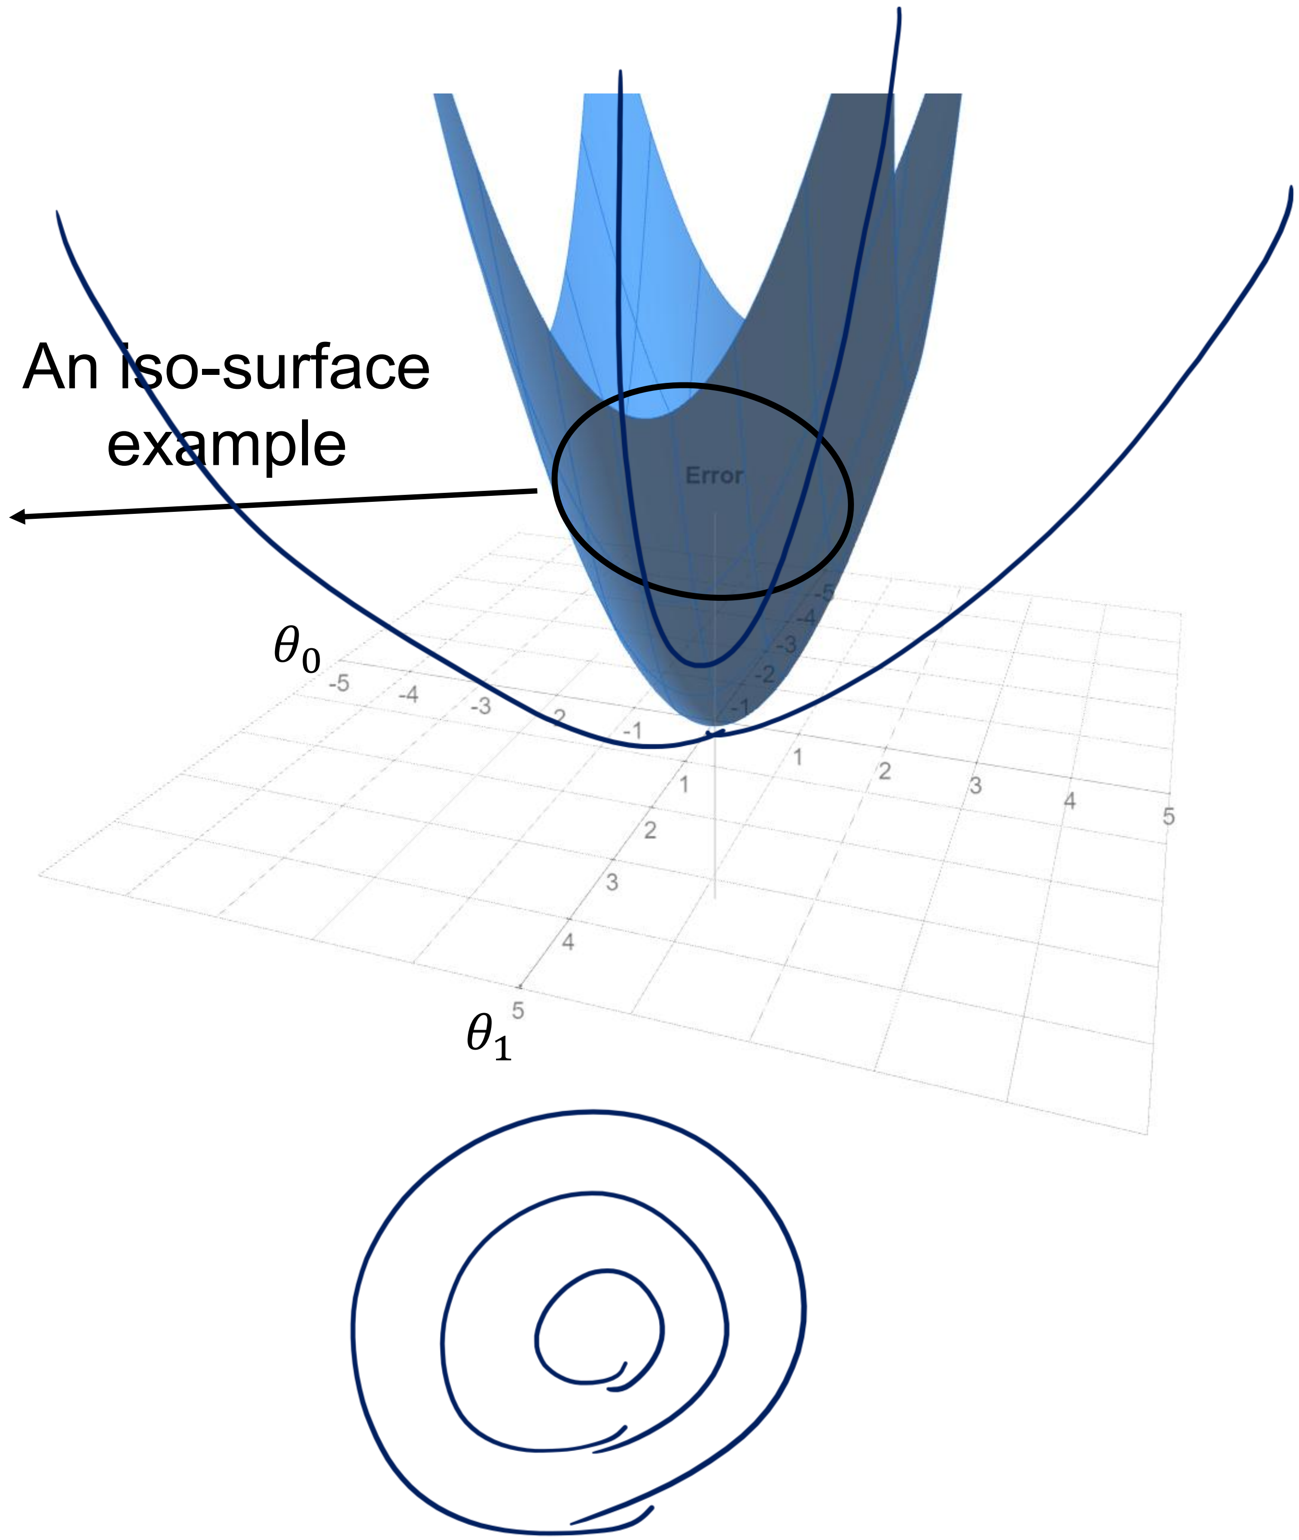

Error function together with a new introduced constraint

$$g(\theta) = \theta_0^2 + \theta_1^2 = \theta^T\theta$$

$$E(\theta) = \frac{1}{n}\sum_{i=1}^{n}\left(y^i - z_i\theta\right)^2$$

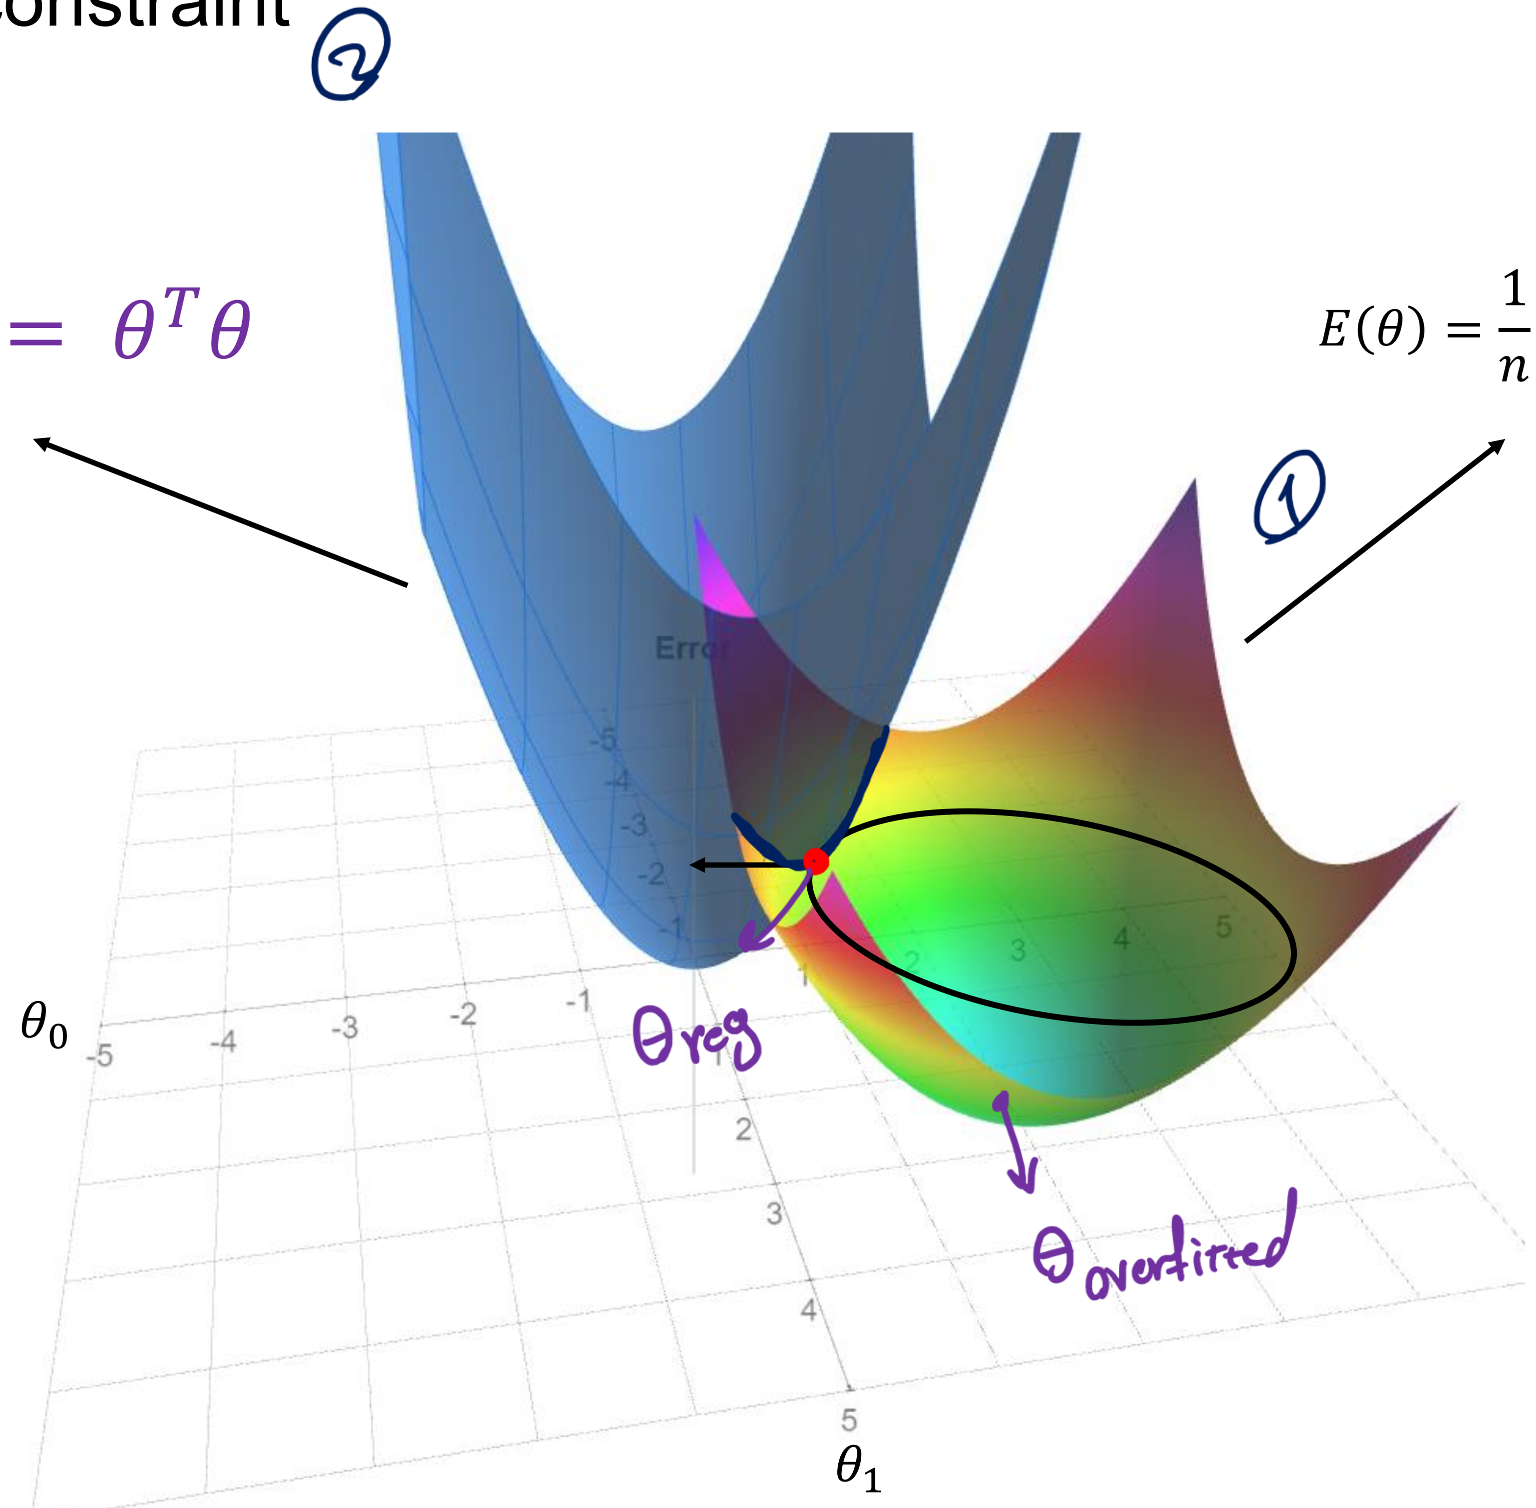

Let’s define the Lagrange function

$$L(\theta, \lambda) = \widetilde{E(\theta)} + \lambda \widetilde{g(\theta)}$$

$$L(\theta, \lambda) = E(\theta) + \lambda(\theta^T \theta - c)$$

$$\nabla L(\theta, \lambda) = 0 \qquad \nabla[E(\theta) + \lambda \theta^T \theta] = 0$$

$$\nabla[E(\theta)] + \lambda \nabla[\theta^T \theta] = 0$$

How to enforce the gradient of Lagrange function to be zero

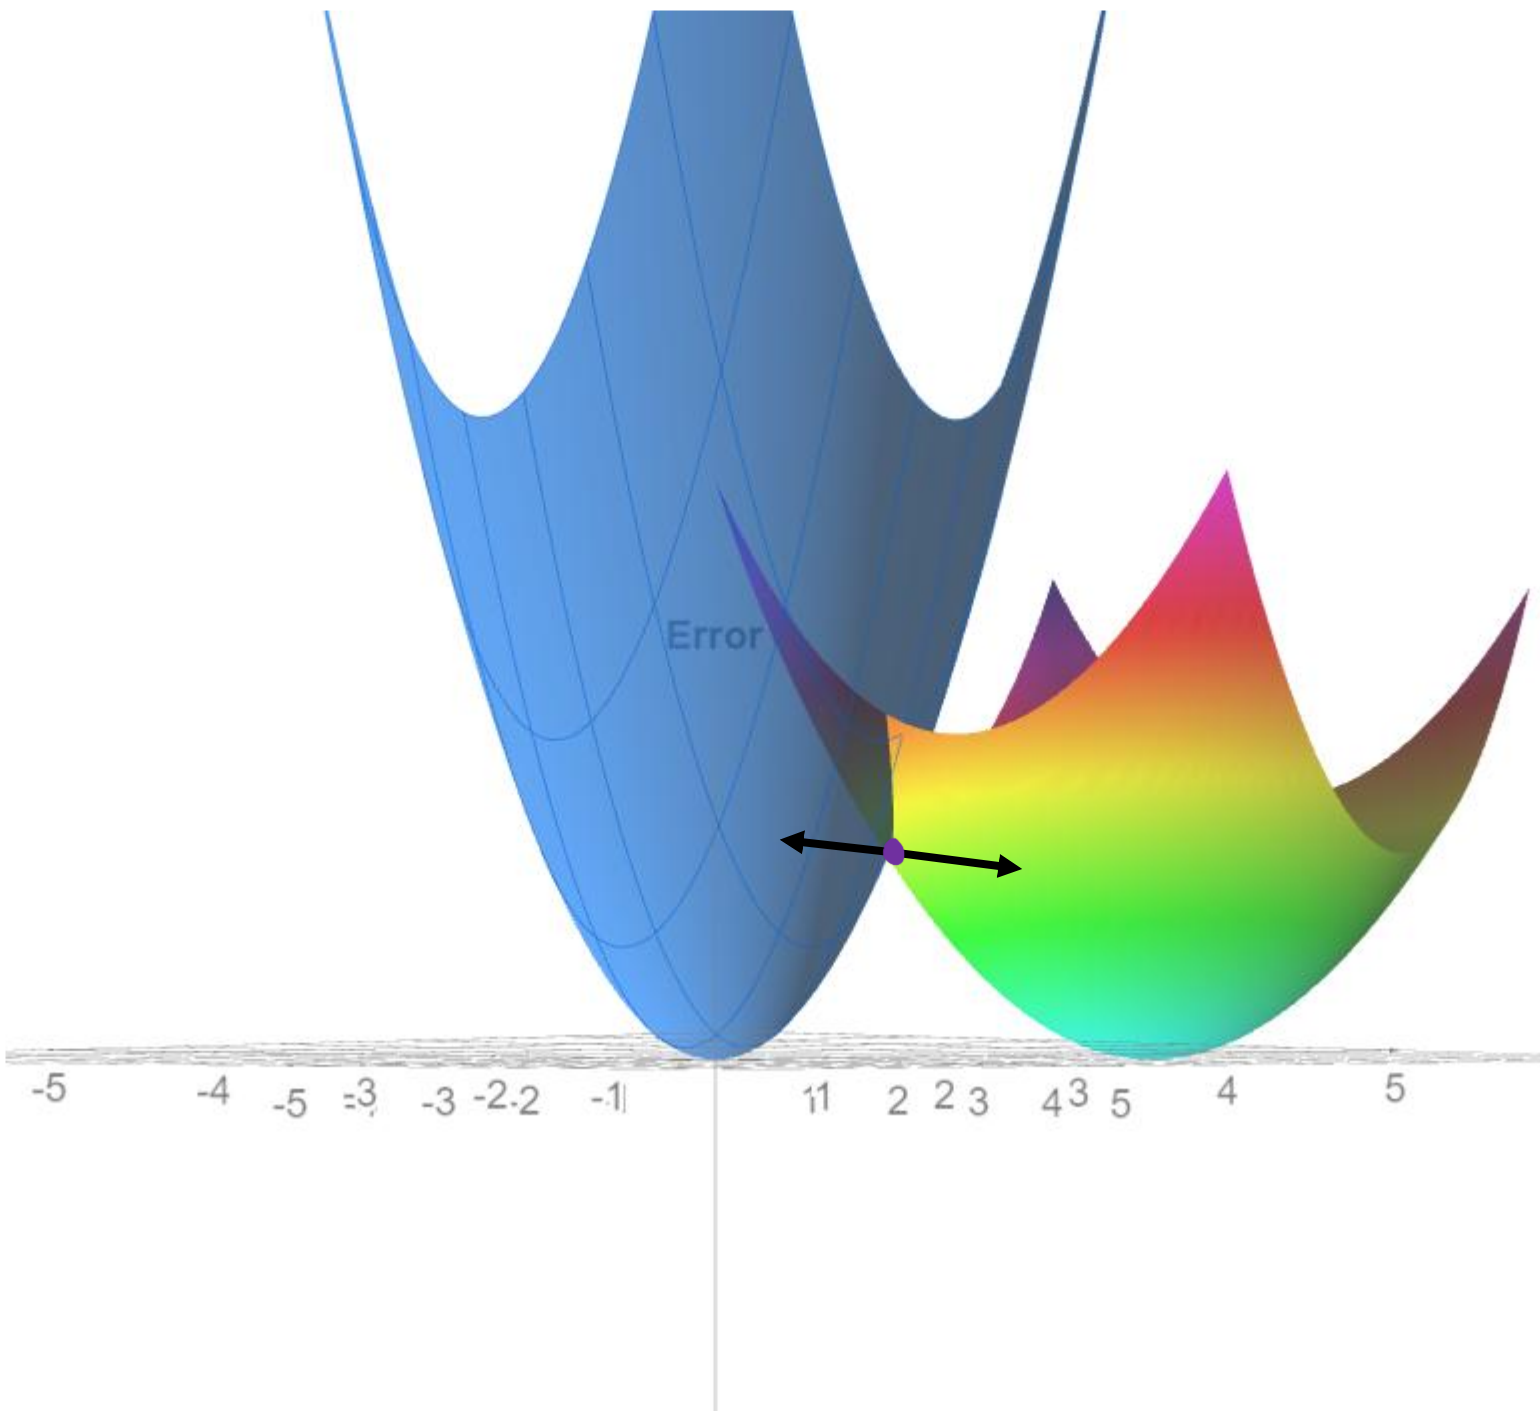

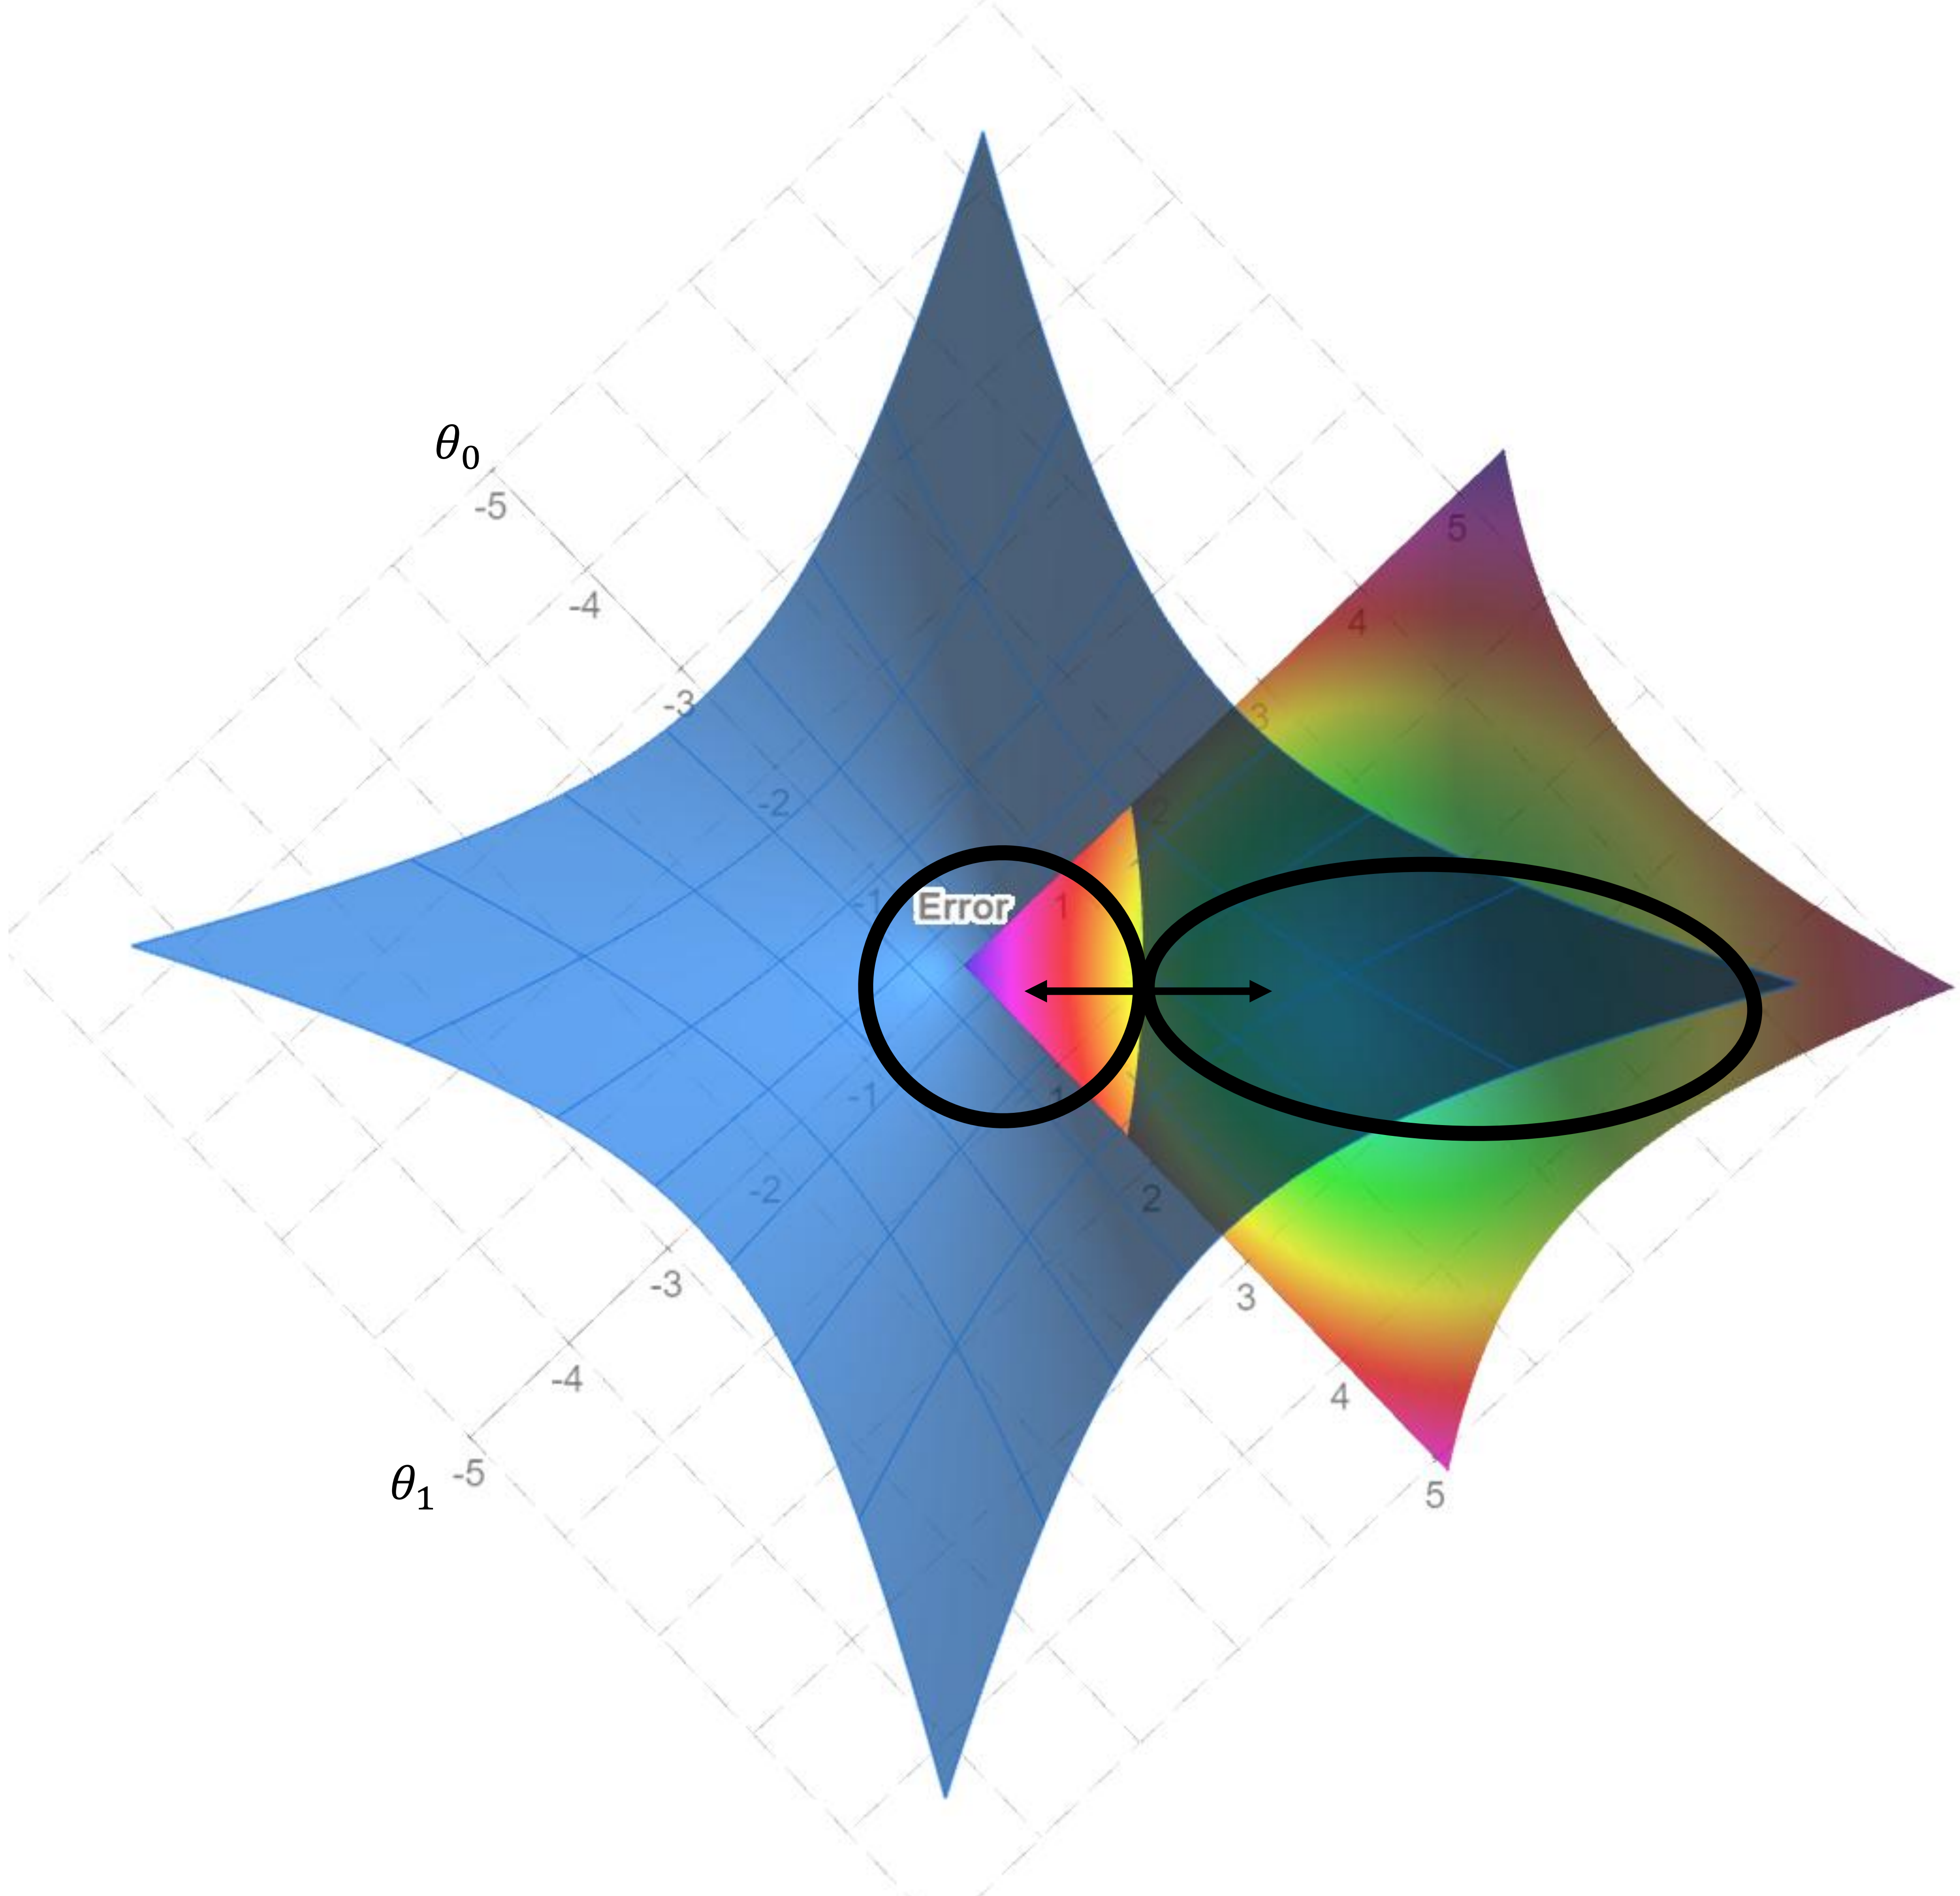

θ_0
θ_1
Error
-5
-4
-3
-2
-1
1
2
3
4
5

Let's calculate the gradients

Gradient of constraint $g(\theta)$ $\nabla[\theta^T\theta] = 2\theta$

$$\nabla[E(\theta)] + \lambda\nabla[\theta^T\theta] = 0$$

$$\nabla[E(\theta)] = -\lambda\nabla[\theta^T\theta]$$

$$\nabla E(\theta) = -2\lambda\theta$$

$$\nabla E(\theta) + 2\lambda\theta = 0$$

Let's do integration $E(\theta) + \lambda\theta^T\theta$

$$L(\theta, \delta) = E(\theta) + \delta(\theta^T\theta - c)$$

$$\frac{\partial L(\theta, \delta)}{\partial \theta} = 0 \quad ; \quad \frac{\partial L(\theta, \delta)}{\partial \delta} = 0 \qquad \text{implicit equation}$$

What if we know δ in advance $\rightsquigarrow$ δ becomes constant

regularized error

$$L(\theta) = \tilde{E}(\theta) = E(\theta) + \delta\theta^T\theta \cancel{- \delta c}$$

$\rightarrow$ Penalty term

minimize $\rightsquigarrow$ $\underbrace{\tilde{E}(\theta)}_{\text{penalized error}} = E(\theta) + \boxed{\delta\,\underbrace{\theta^T\theta}_{c}}$

$\rightarrow$ Penalty term

$\rightarrow$ Regularization term

The effect of low Lambda

$$E(\theta) + \frac{\lambda}{N}\theta^T\theta$$

c

0

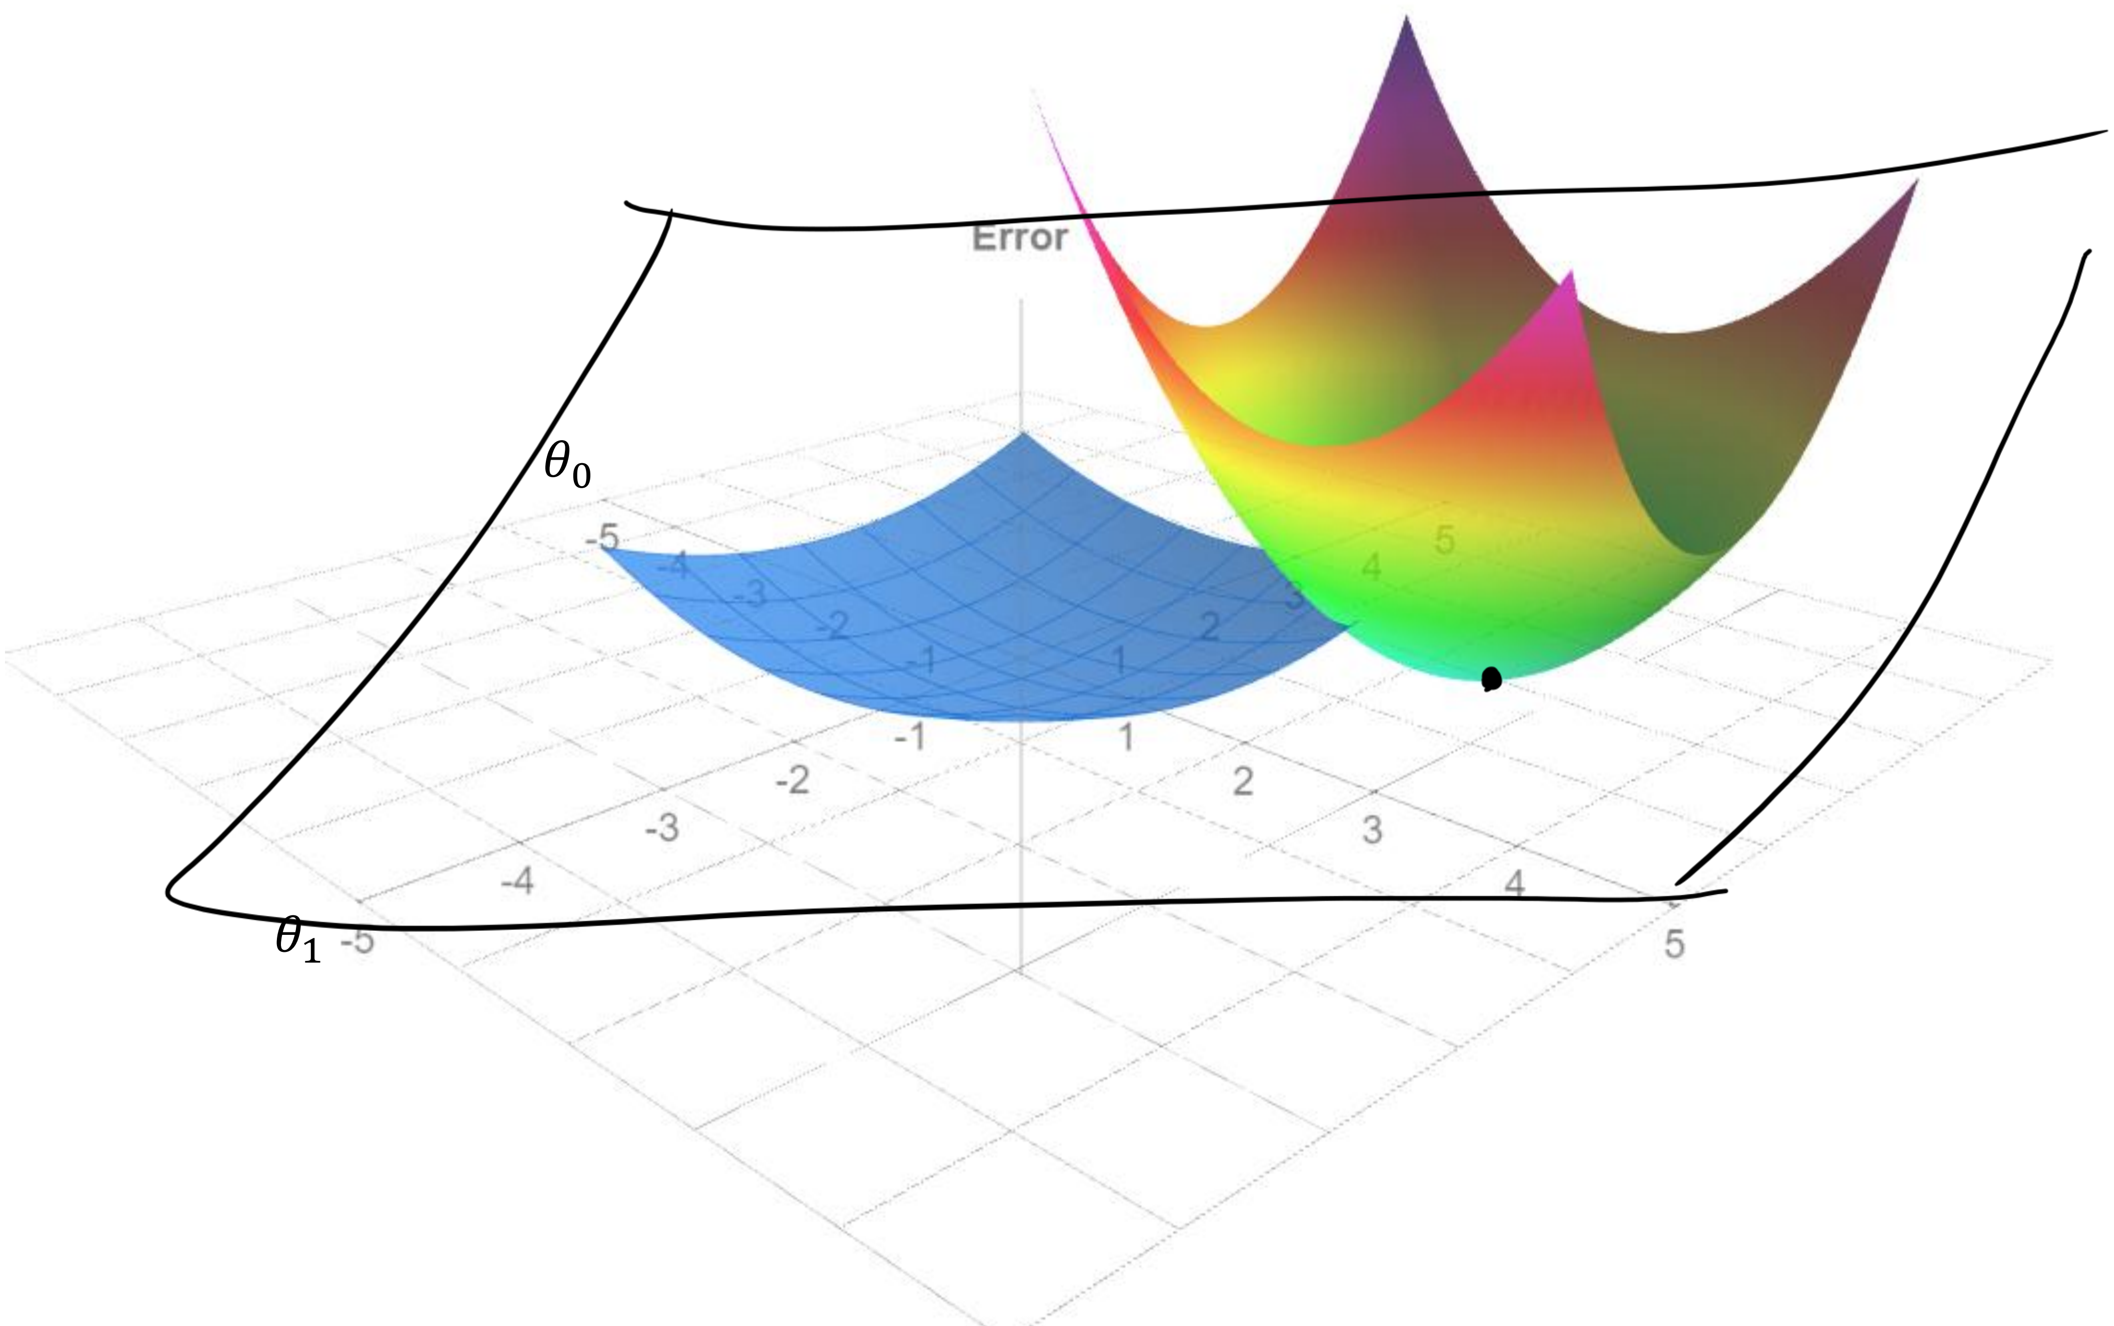

The effect of high Lambda

$$E(\theta) + \frac{\uparrow\lambda}{N}\theta^T\theta \downarrow$$

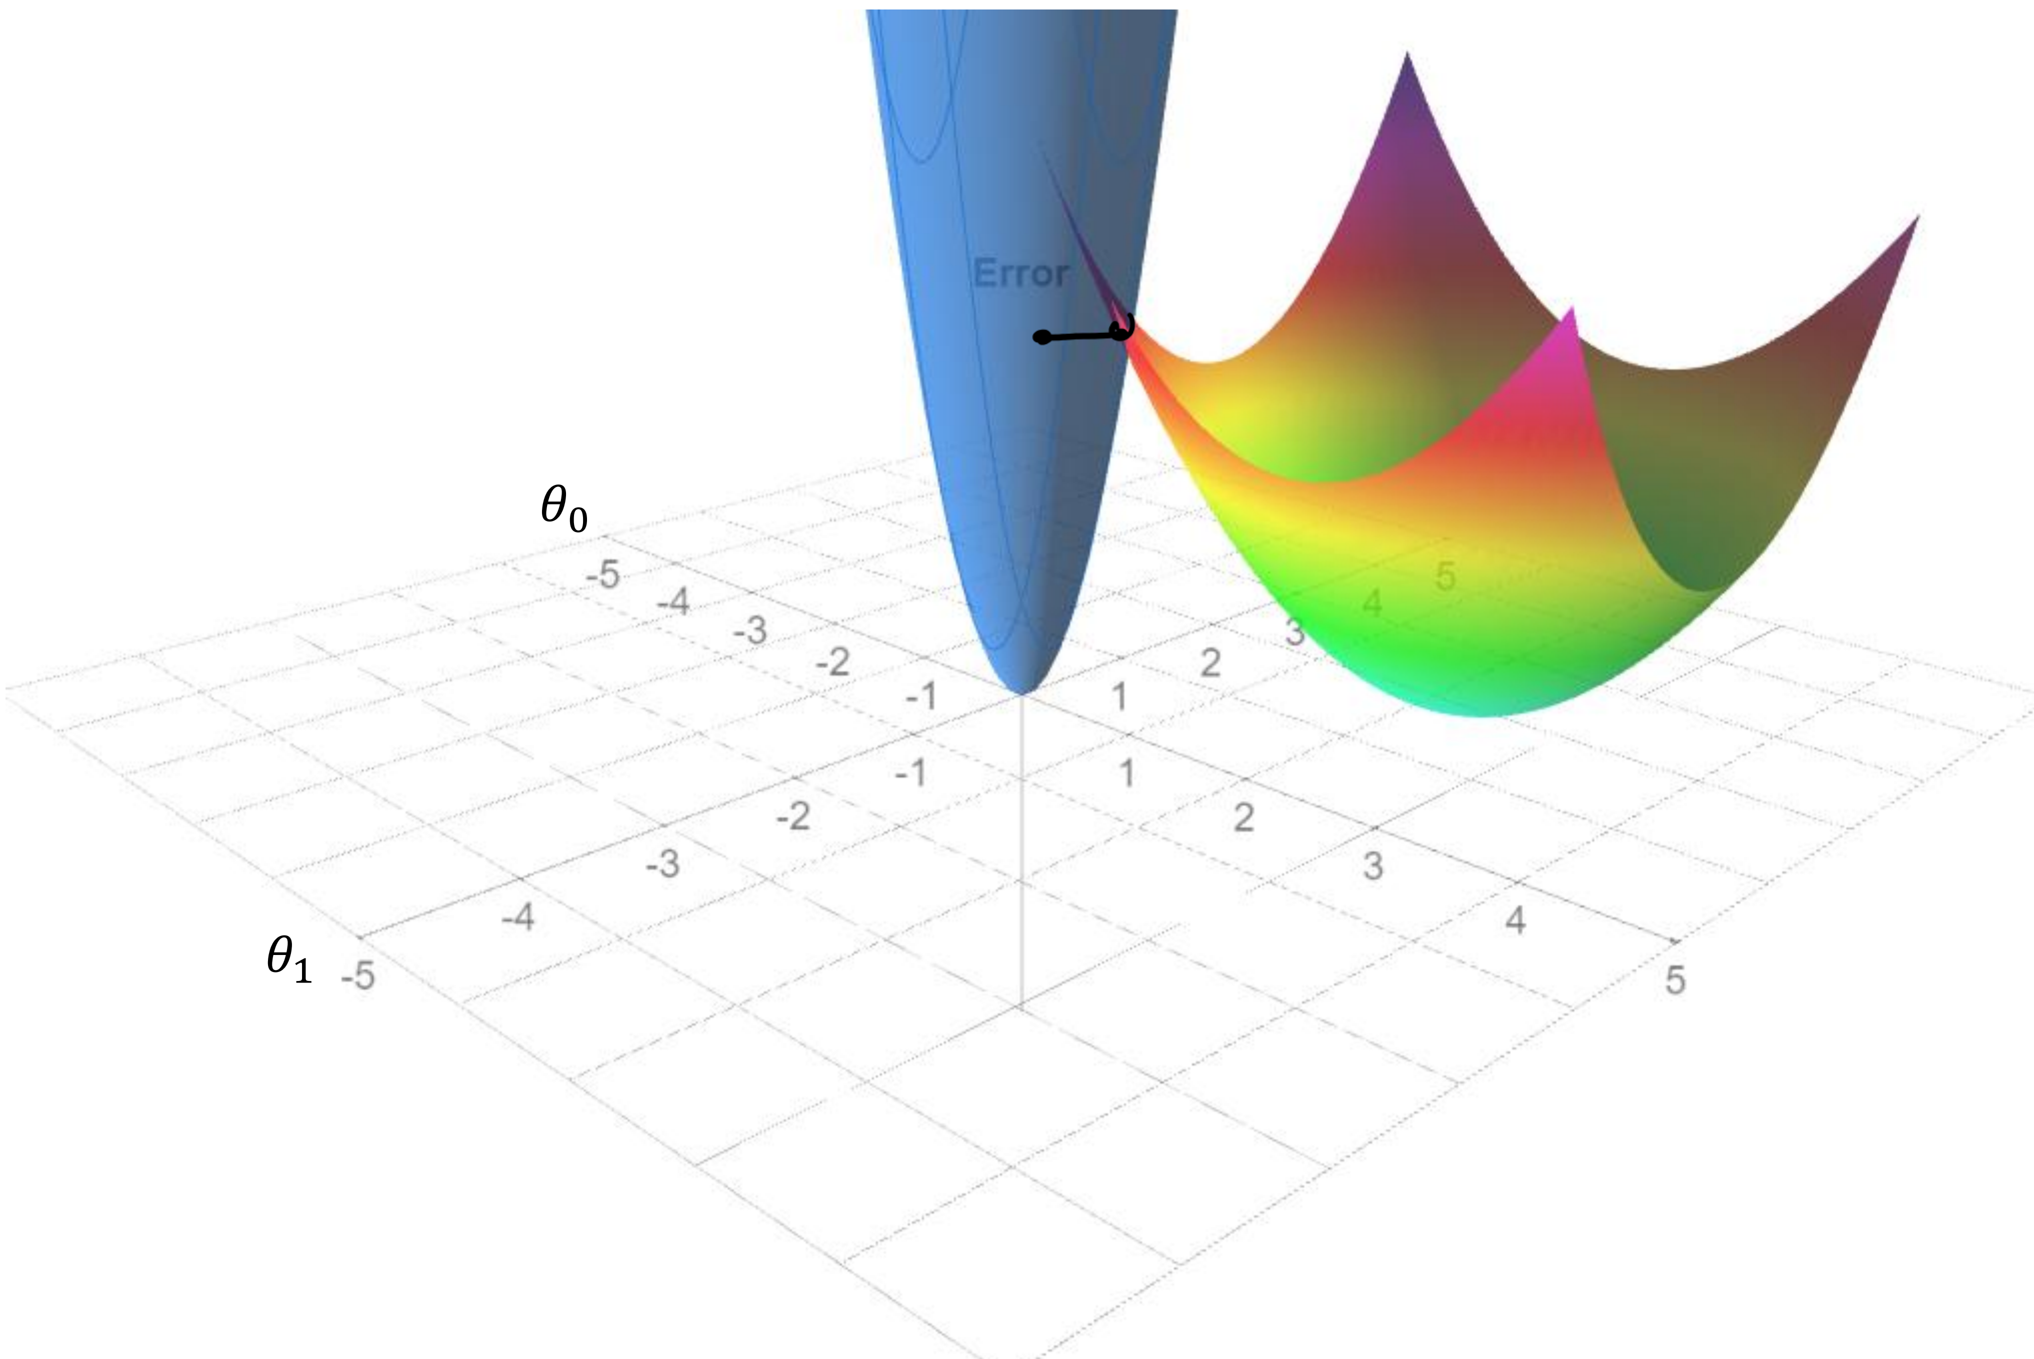

Regularized Learning

Now we know Why this term leads to the regularization of parameters

Minimize $\boxed{E(\theta) + \lambda\theta^T\theta}$

$\tilde{E}(\theta)$

Regularized Error $= \dfrac{1}{N}\sum_{i=1}^{n}\left(y^{\{i\}} - z^{\{i\}}\theta\right)^2 + \dfrac{\lambda}{2N}\|\theta\|_2^2$

L2 Regularization term

Outline

- Overfitting and regularized learning
- Ridge regression
- Lasso regression
- Determining regularization strength

Ridge Regression

$$\tilde{E}(\theta)$$
$$= \frac{1}{N}\sum_{i=1}^{n}\left(y^{\{i\}} - z^{\{i\}}\theta\right)^2 + \boxed{\lambda\|\theta\|_2^2}$$

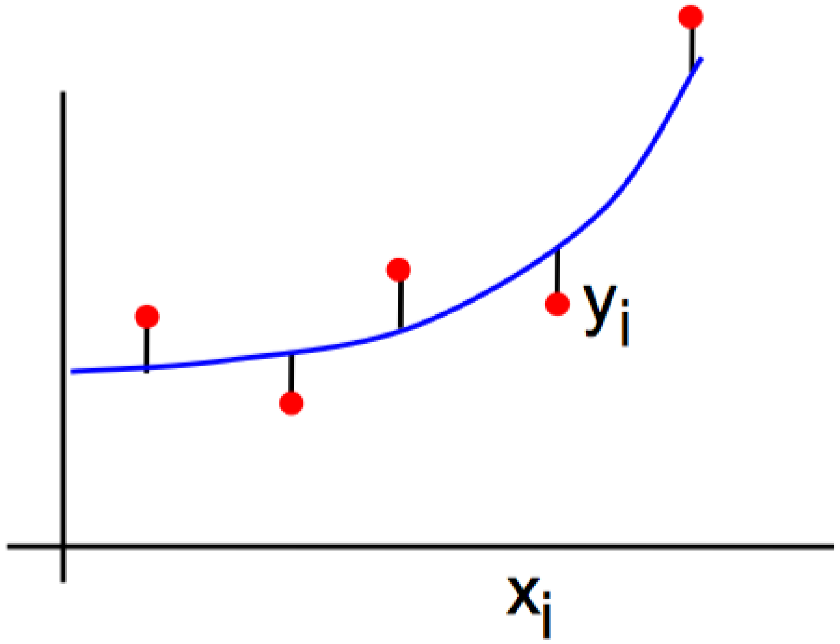

$$\theta_0 + \theta_1 z_1 + \theta_2 z_2 + \cdots + \theta_d z_\mathrm{d} + \epsilon = \boldsymbol{z\theta}$$

General form $$\tilde{E}(\theta) = \frac{1}{N}\sum_{i=1}^{n}\left(y^{\{i\}} - z^{\{i\}}\theta\right)^2 + \frac{\lambda}{2}\|\theta\|_2^2$$

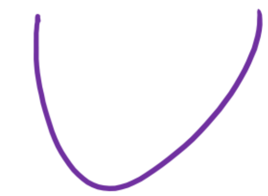

Matrix form $$\tilde{E}(\theta) = \frac{1}{N}(y - z\theta)^{\mathrm{T}}(\mathrm{y} - z\theta) + \frac{\lambda}{2}\|\theta\|_2^2$$

$$\frac{\partial\tilde{E}(\theta)}{\partial\theta} = -z^T(y - z\theta) + \lambda\theta$$

$$\theta_{overfitted} = (z^T z)^{-1} z^T y$$

$$(z^T z + \lambda I)\theta = z^T y$$

$$\boxed{\theta_{reg} = (z^T z + \lambda I)^{-1} z^T y}$$

$$\frac{1}{z^T z + \lambda I}$$

Ridge Regression Example

- The red curve is the true function (which is not a polynomial)
- The data points are samples from the curve with added noise in y.
- There is a choice in both the degree, D, of the basis functions used, and in the strength of the regularization

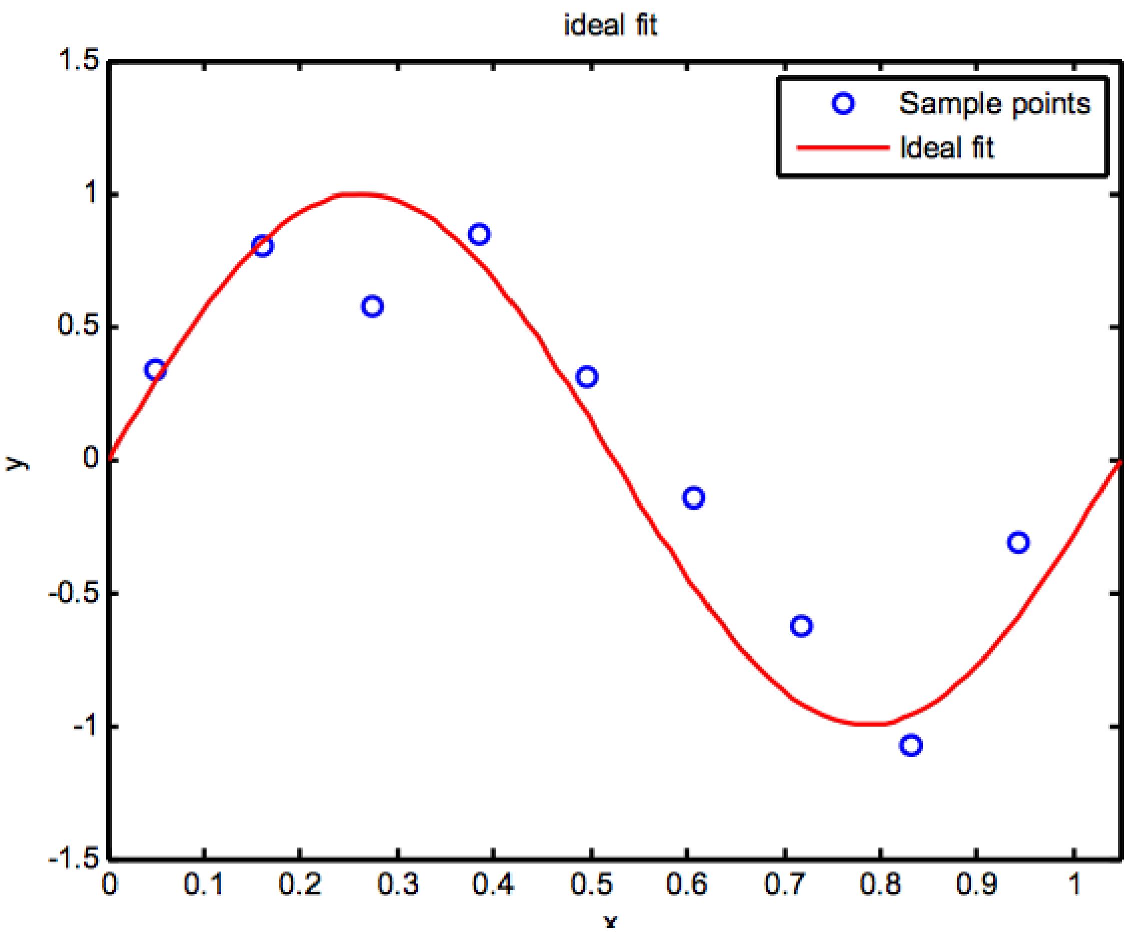

$$f(x, \theta) = z\theta \qquad z: x \to z$$

$$\tilde{E}(\theta) = \frac{1}{N} \sum_{i=1}^{n} \left(y^{\{i\}} - z^{\{i\}}\theta\right)^2 + \lambda \|\theta\|_2^2 \quad \theta \in \mathbb{R}^{D+1}$$

N = 9 samples, D = 7

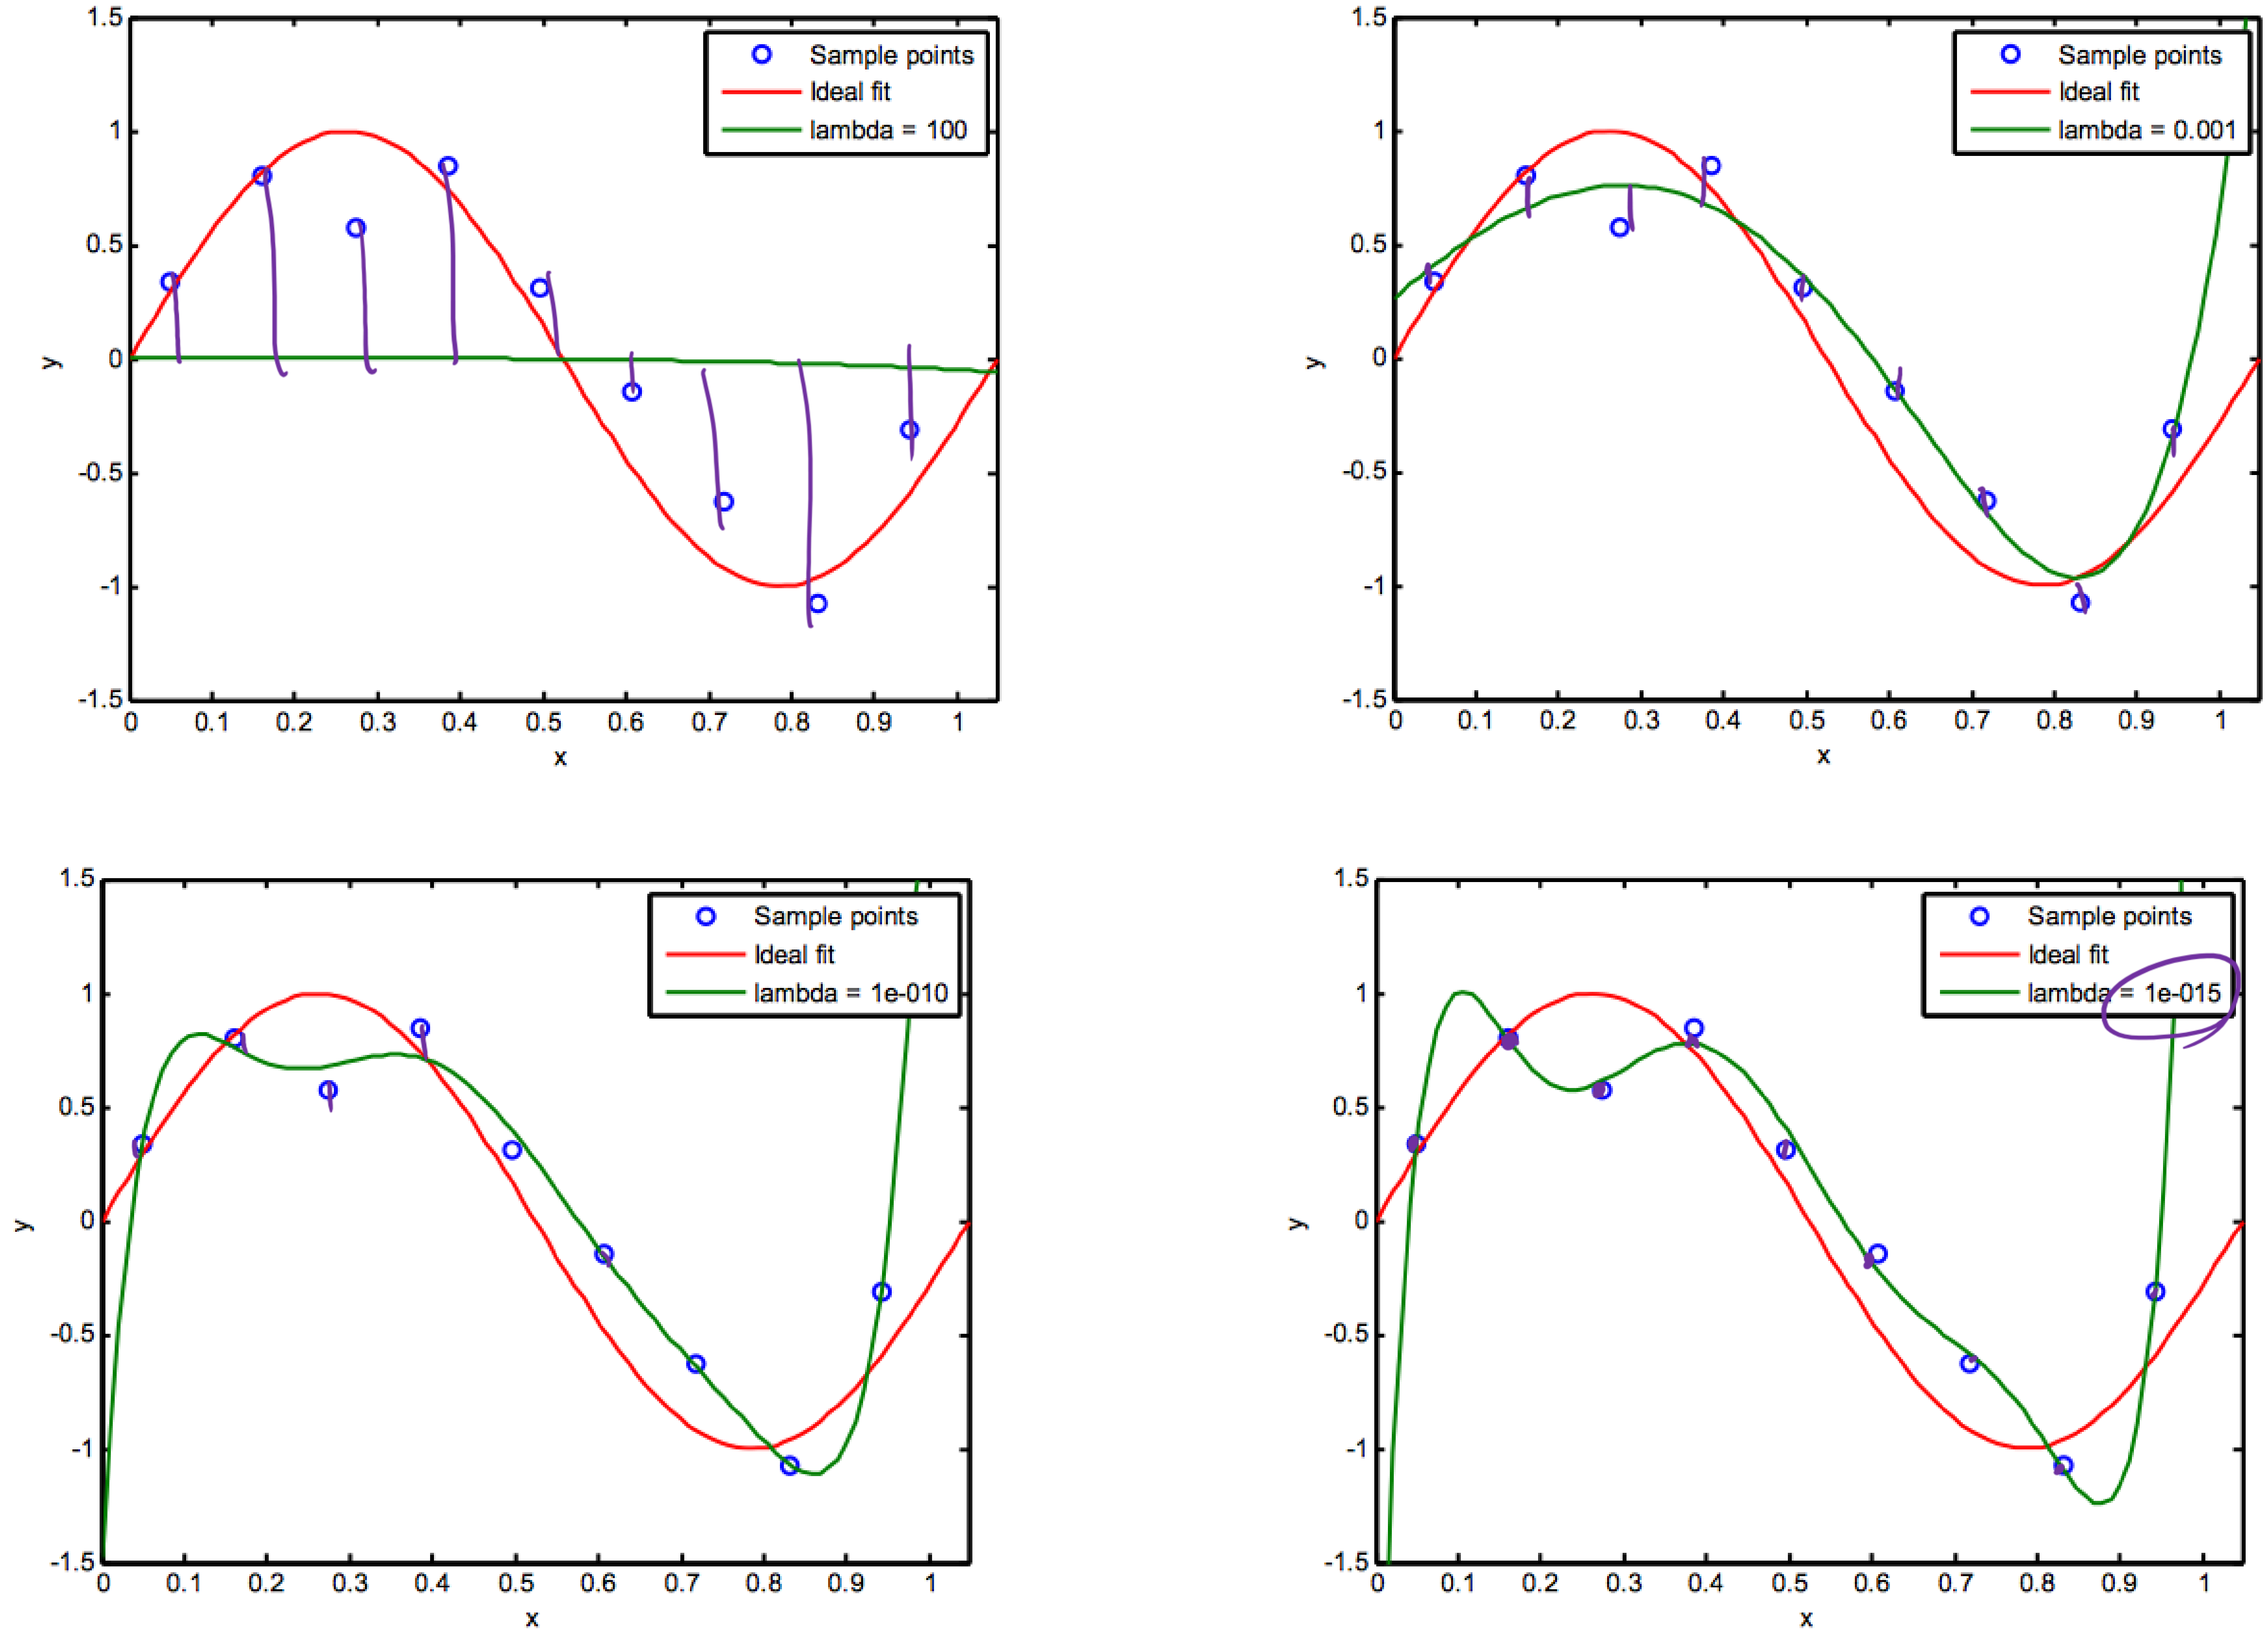

D = 3

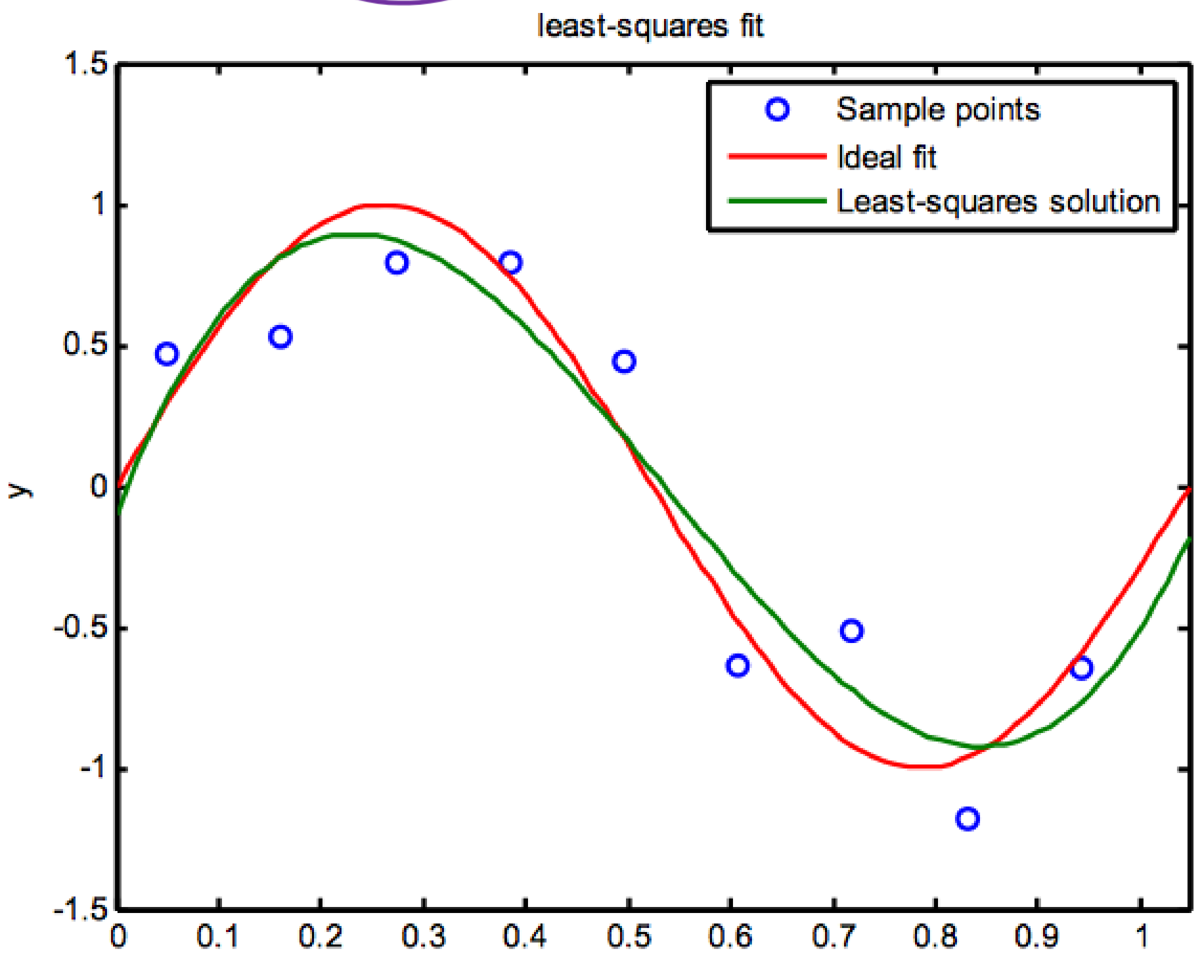

least-squares fit
Sample points
Ideal fit
Least-squares solution
y
1.5
1
0.5
0
-0.5
-1
-1.5
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

D = 5

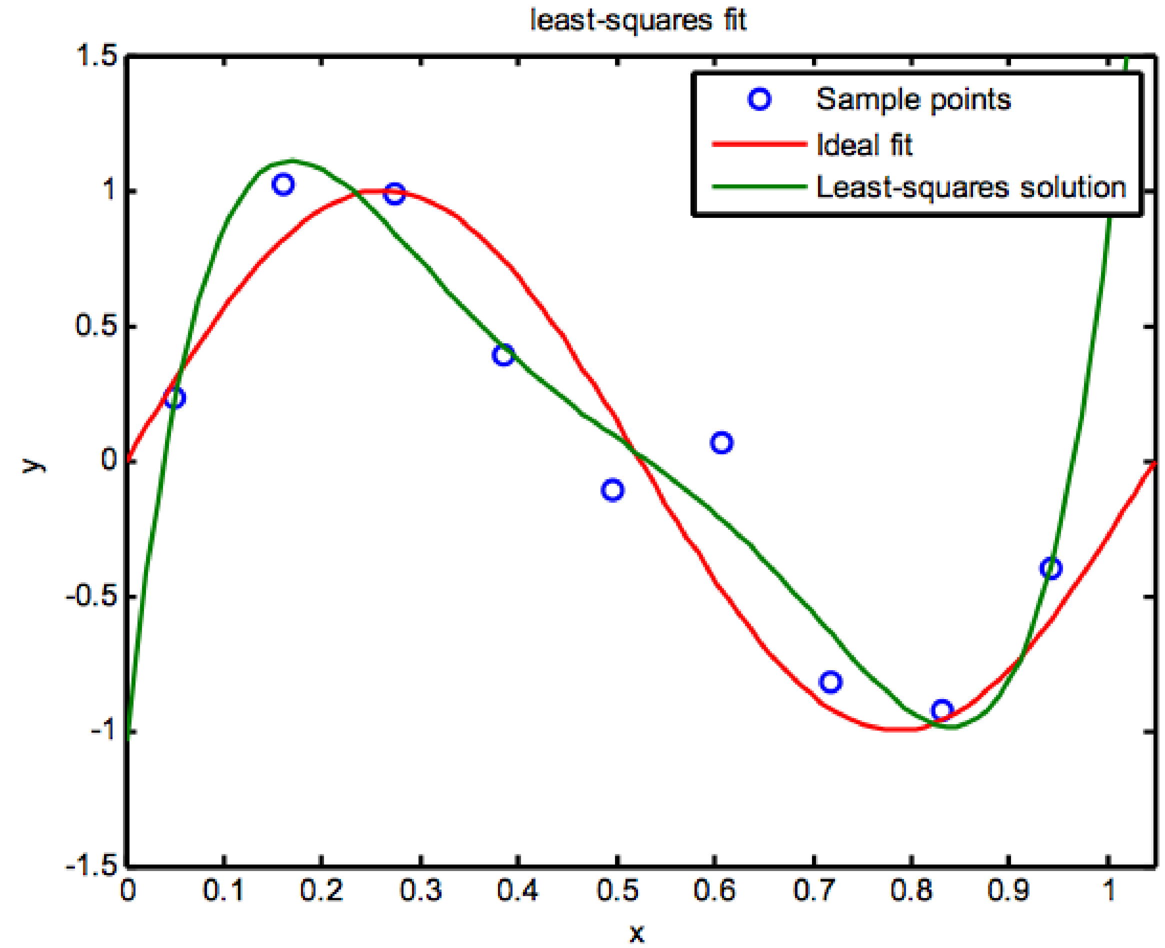

least-squares fit
Sample points
Ideal fit
Least-squares solution
y
x
1.5
1
0.5
0
-0.5
-1
-1.5
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

Outline

- Overfitting and regularized learning
- Ridge regression
- Lasso regression ⬅ Dimensionality reduction
- Determining regularization strength

- PCA ⇝ max variance ⇝ unsupervised   PCA ⇒ np.svd(X)
- Forward Feature selection → supervised
- Backward " " → "
- Lasso → "

Regularized Regression

$$\tilde{E}(\theta) = \frac{1}{N}\sum_{i=1}^{n}\left(y^{\{i\}} - z^{\{i\}}\theta\right)^2 + \lambda\|\theta\|_2^2$$

Squared loss\Error $\frac{1}{N}\sum_{i=1}^{n}\left(y^{\{i\}} - z^{\{i\}}\theta\right)^2$

L2 Regularizer $\lambda\|\theta\|_2^2$

Now let's look at another regularization choice.

The Lasso Regularization (L1 norm) and sparsity

Lasso = **L**east Absolute **S**hrinkage and **S**election **O**perator

$y = |x|$

$$\tilde{E}(\theta) = \frac{1}{N}\sum_{i=1}^{n}\left(y^{\{i\}} - z^{\{i\}}\theta\right)^2 + \lambda\|\theta\|_1$$

L1 norm induces sparsity. This means that some of the weights become zero, and the feature contribution will be completely removed. L1 Regularizer could be used for feature selection

Ridge Regularizer

$$g(\theta) = \theta_0^2 + \theta_1^2 = \theta^T\theta$$

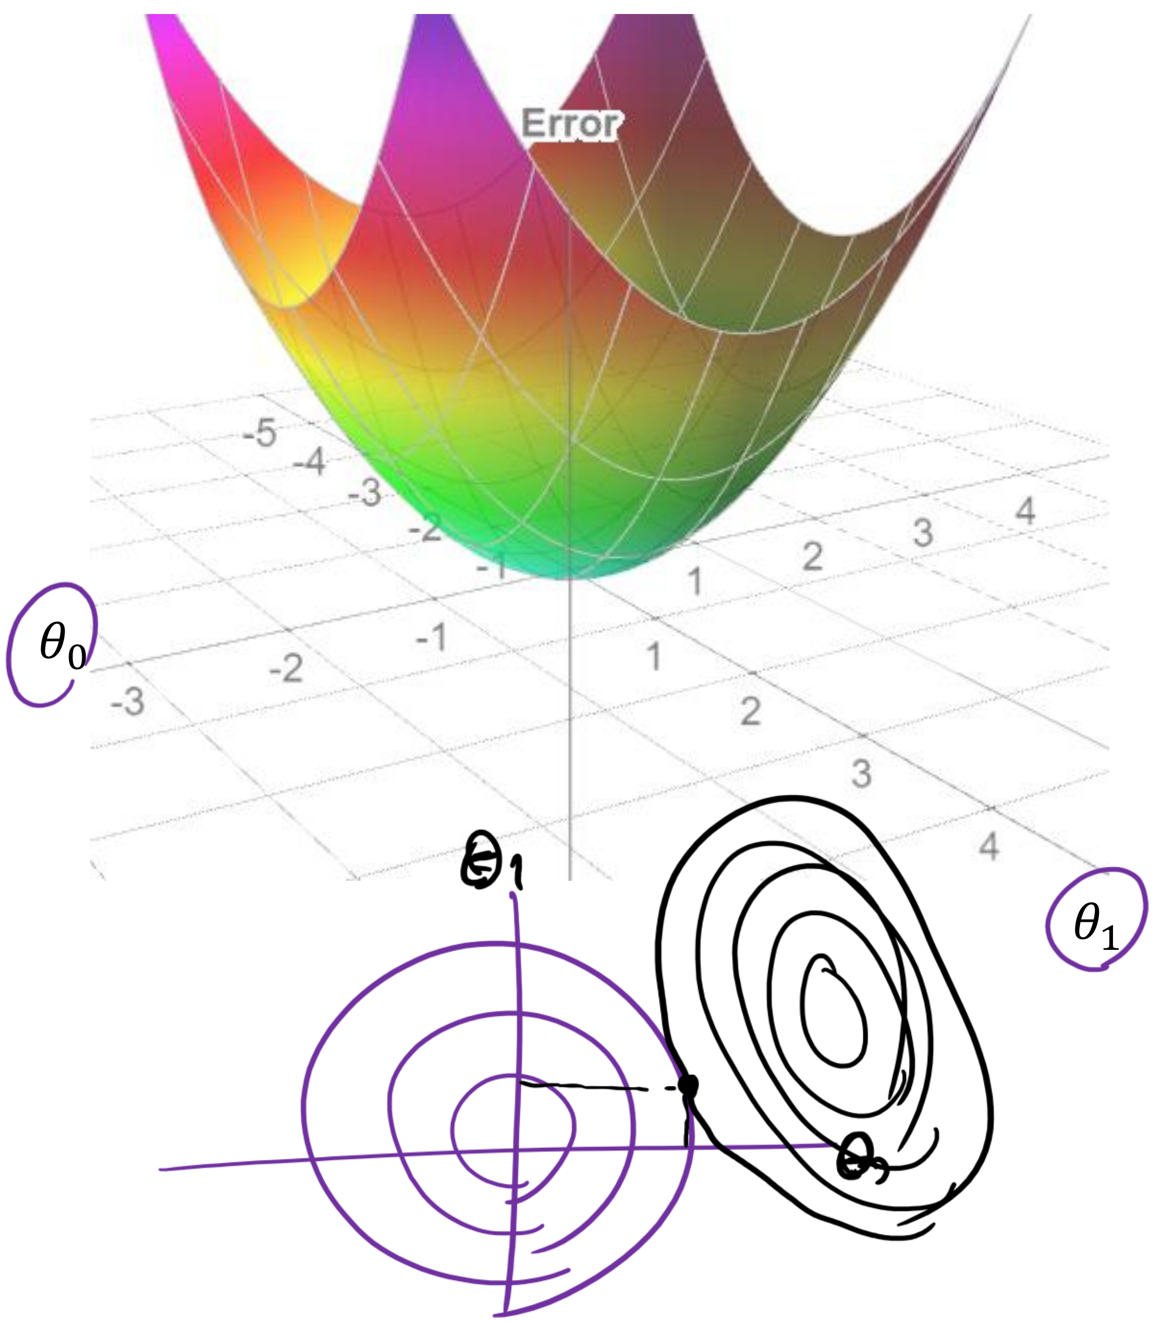

Lasso Regularizer

$$g(\theta) = \theta_0 + \theta_1 = \theta$$

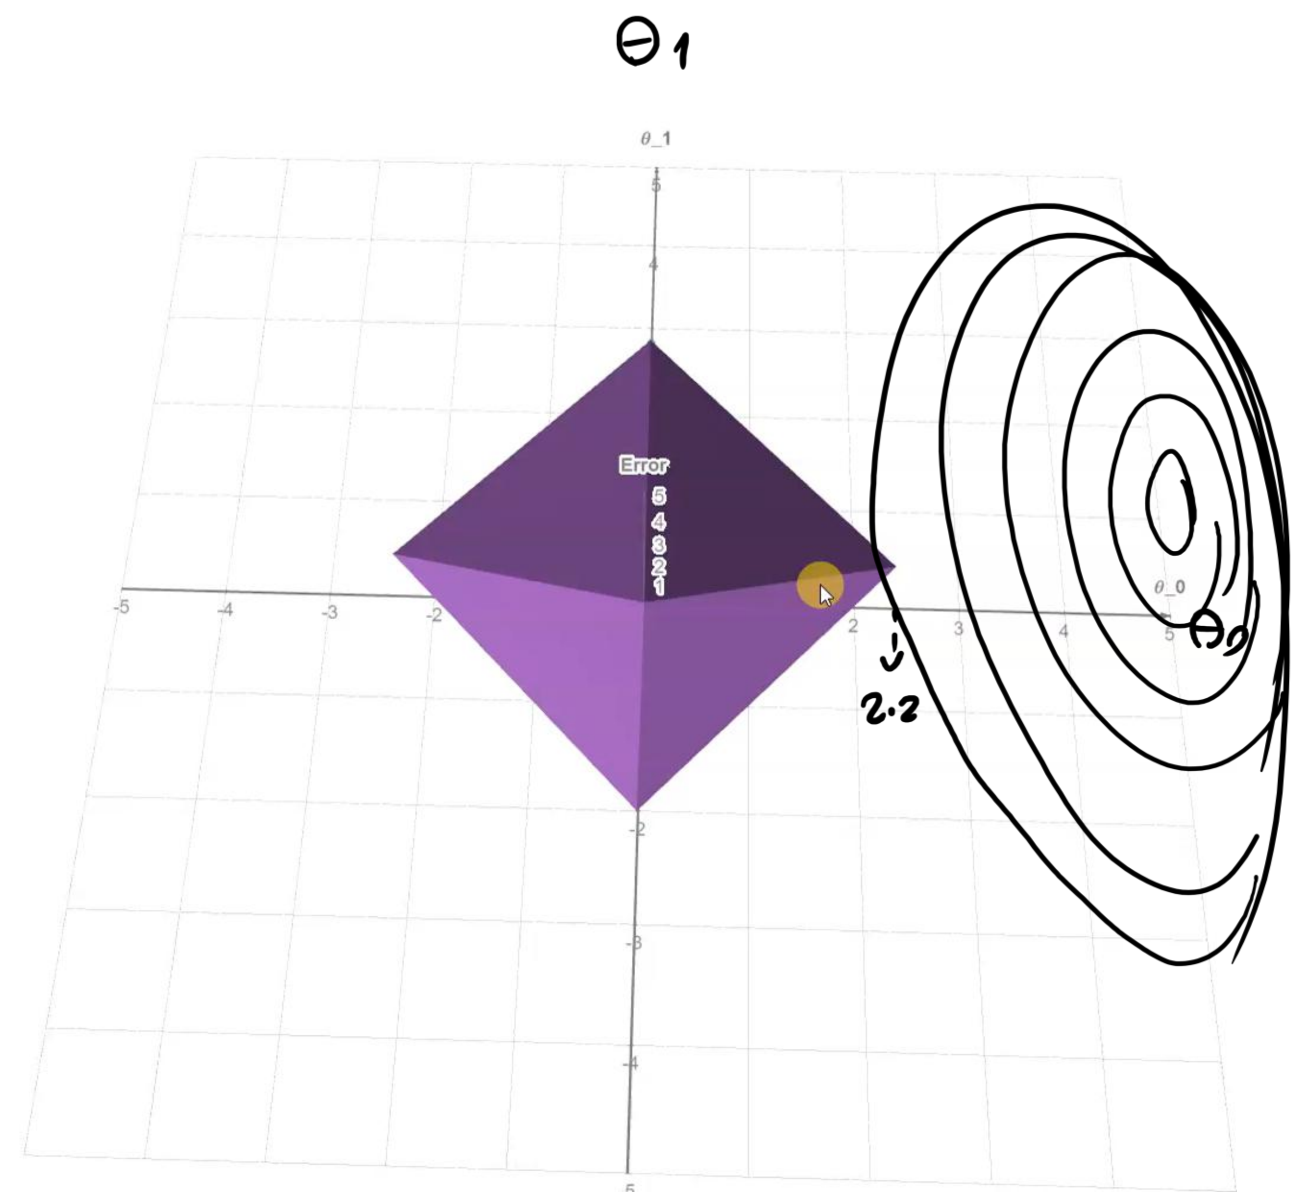

[Animation]()

Let's say we have two parameters (θ_0 and θ_1)

$$Min\ E(\theta) = \frac{1}{N}(z\theta - y)^T(z\theta - y) + \lambda\|\theta\|_1$$

$$\theta = \begin{bmatrix} \theta_0 \\ \mathbf{0} \end{bmatrix}$$

Interesting way for feature selection

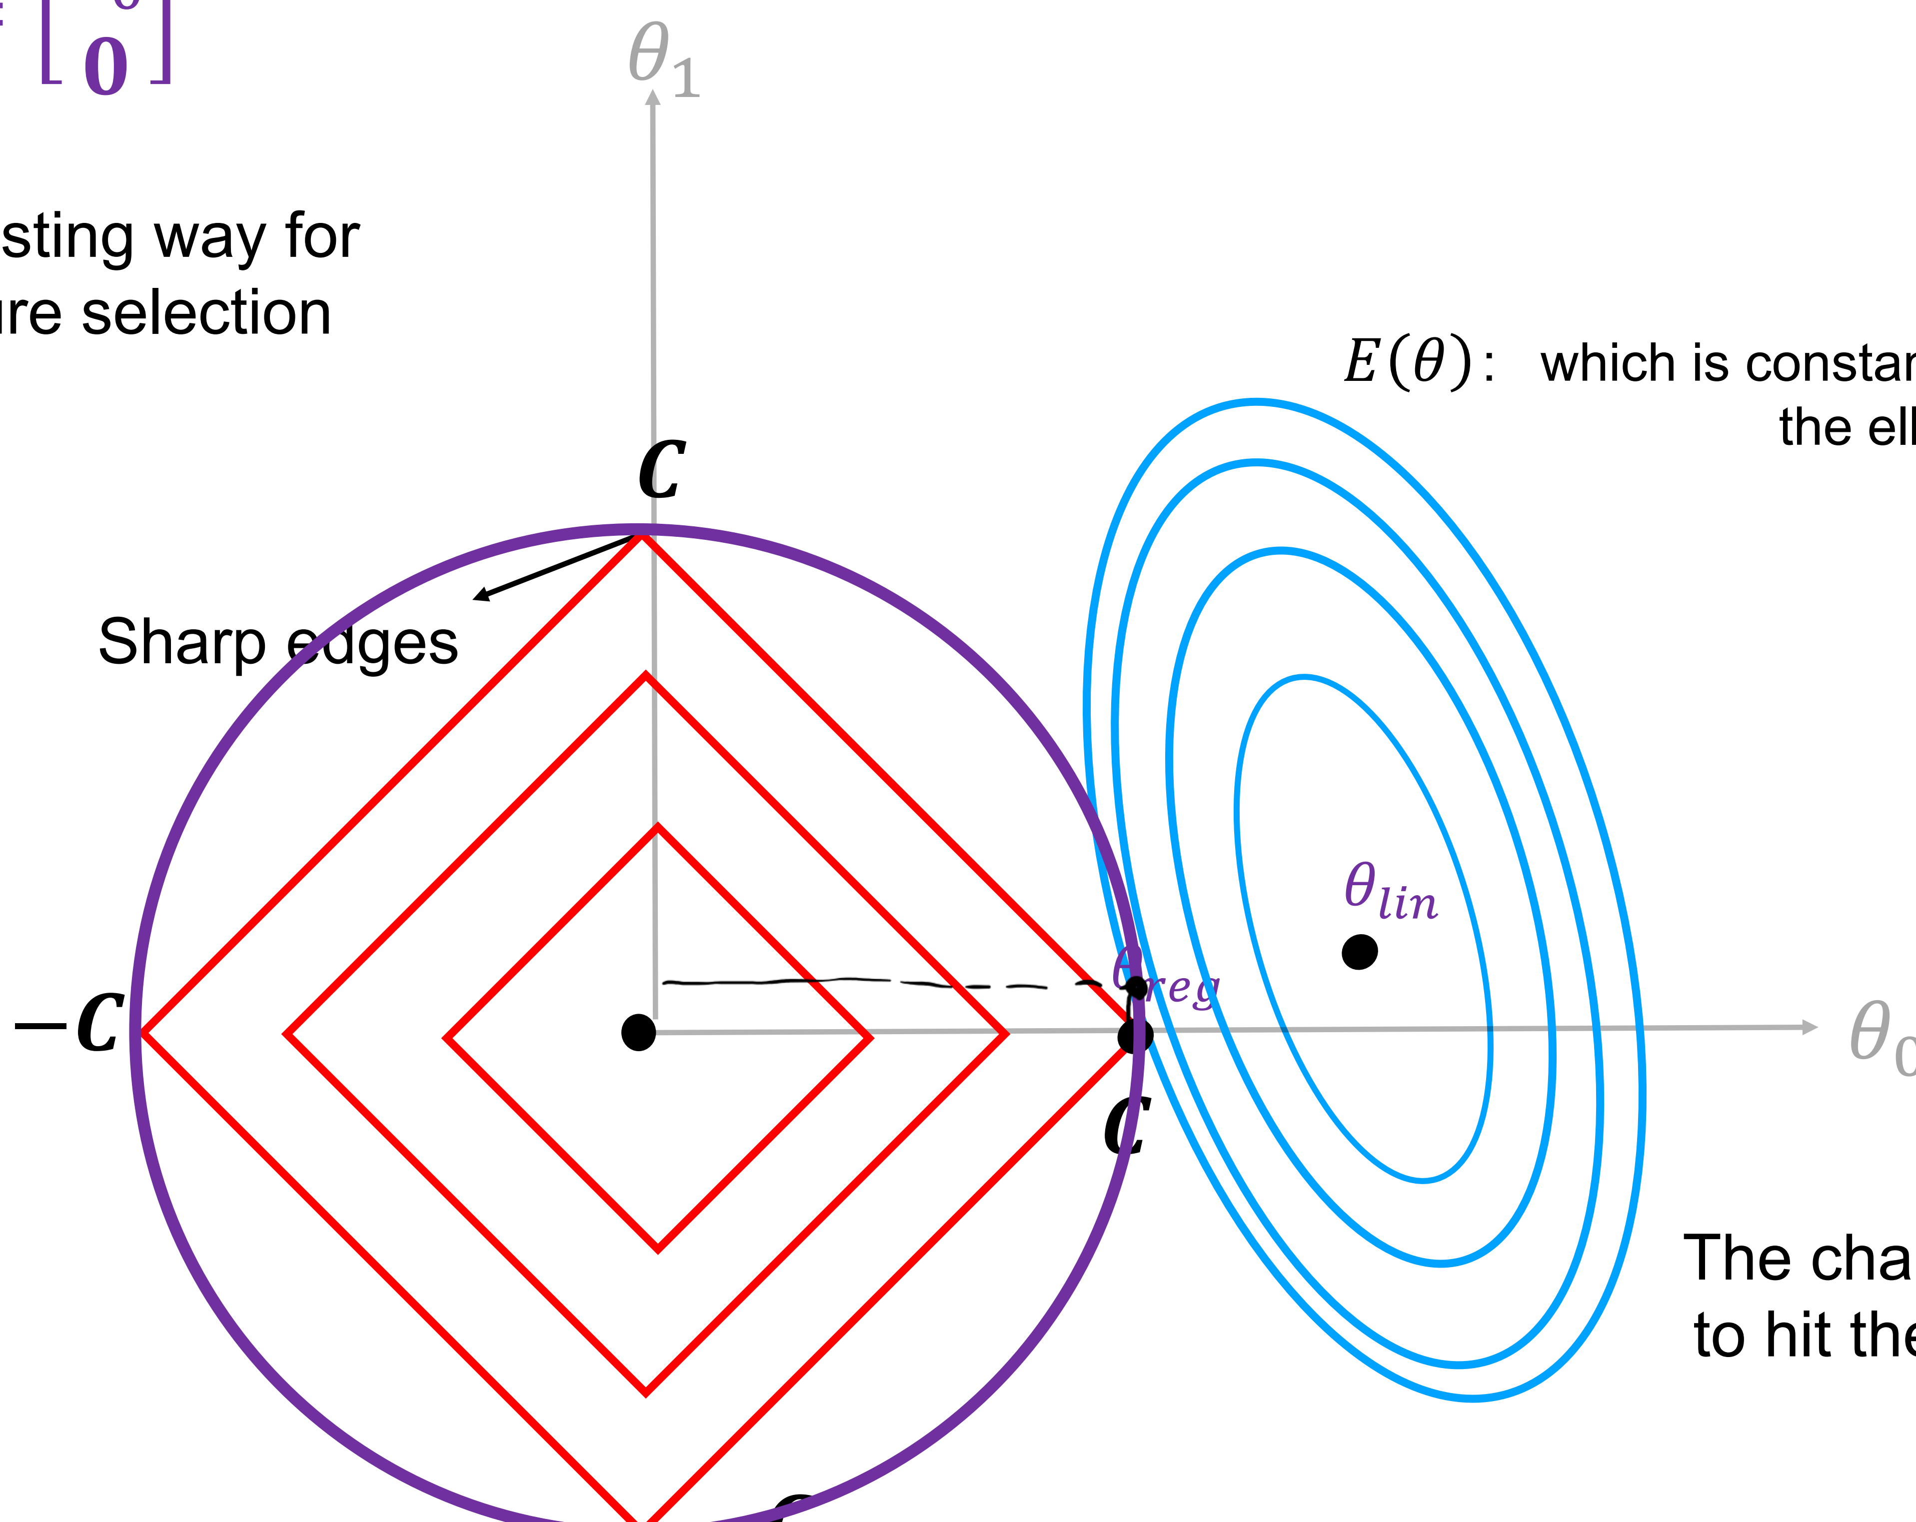

The chance is very high to hit the sharp corners first

Graph

Ridge versus Lasso

Ridge

$$\tilde{E}(\theta) = \frac{1}{N}(y - z\theta)^{\mathrm{T}}(\mathrm{y} - z\theta) + \lambda\|\theta\|_2^2$$

It is a convex model

Both mean squared error and L2 regularizer are differentiable.

We can get a closed form solution

Lasso

$$\tilde{E}(\theta) = \frac{1}{N}(y - z\theta)^{\mathrm{T}}(\mathrm{y} - z\theta) + \lambda\|\theta\|_1$$

It is a convex model

L1 regularizer is NOT differentiable.

We can **NOT** get a closed form solution

Sub-gradient Descend in Lasso

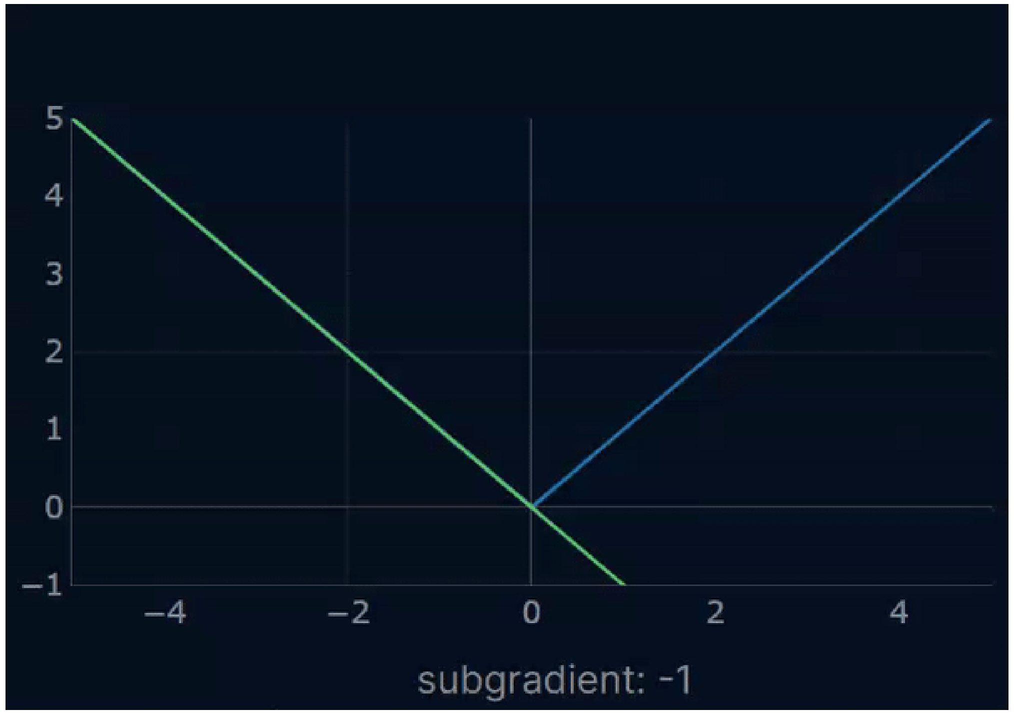

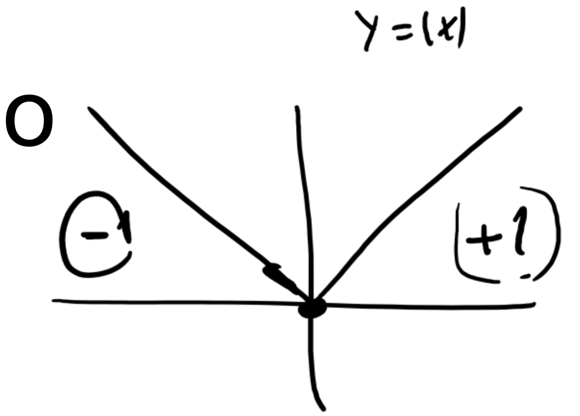

$$\tilde{E}(\theta) = \frac{1}{N}(y - z\theta)^{\mathrm{T}}(y - z\theta) + \lambda\|\theta\|_1$$

$$\frac{\partial \tilde{E}(\theta)}{\partial \theta} = -z^T(y - z\theta) + \frac{\partial(\lambda\|\theta\|_1)}{\partial \theta}$$

Sign + → +1, 0 → 0, −1 → −1

$$\theta^{\{t+1\}} \leftarrow \theta^{\{t\}} - \alpha \frac{\partial \tilde{E}(\theta)}{\partial \theta}$$

Using Sub-gradient

$$\frac{\partial \tilde{E}(\theta)}{\partial \theta} = -z^T(y - z\theta) + \lambda sign(\theta)$$

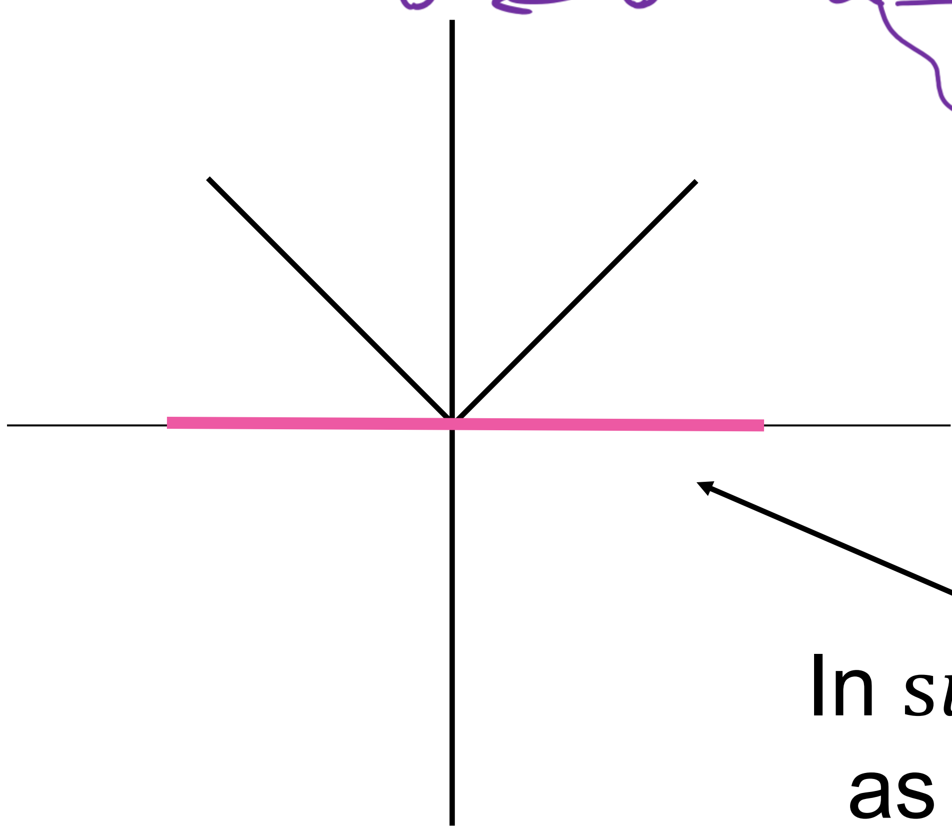

In $sign$ function, we use this sub-gradient line as our under-estimator (below our function)

Outline

- Overfitting and regularized learning
- Ridge regression
- Lasso regression
- Determining regularization strength

Hyper Parameter

Leave-One-Out Cross Validation

$\delta_1 = 0.01$ $\delta_2 = 0.02$ . . . $\delta_5 = 0.05$

For every $i = 1, \ldots, n$:

$\underset{d+1 \times 1}{\theta_{reg}} = (Z^T Z + \delta I)^{-1}_{d+1 \times d+1} Z^T \underset{n \times 1}{y}$

- train the model on every point except i,
- compute the test error on the held out point.

LCF

$\hat{y}_p = Z\theta$

Average the test errors. $$CV_{(n)} = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i^{(-i)})^2$$

$E_1(\theta) = \sum_{i=1}^{N} (y_a - \hat{y}_p)^2$

Validation datapoint

$E_1(\theta) = y_a - Z\theta$

$\{i\}$ Validation

Training datapoints

$E_2(\theta)$

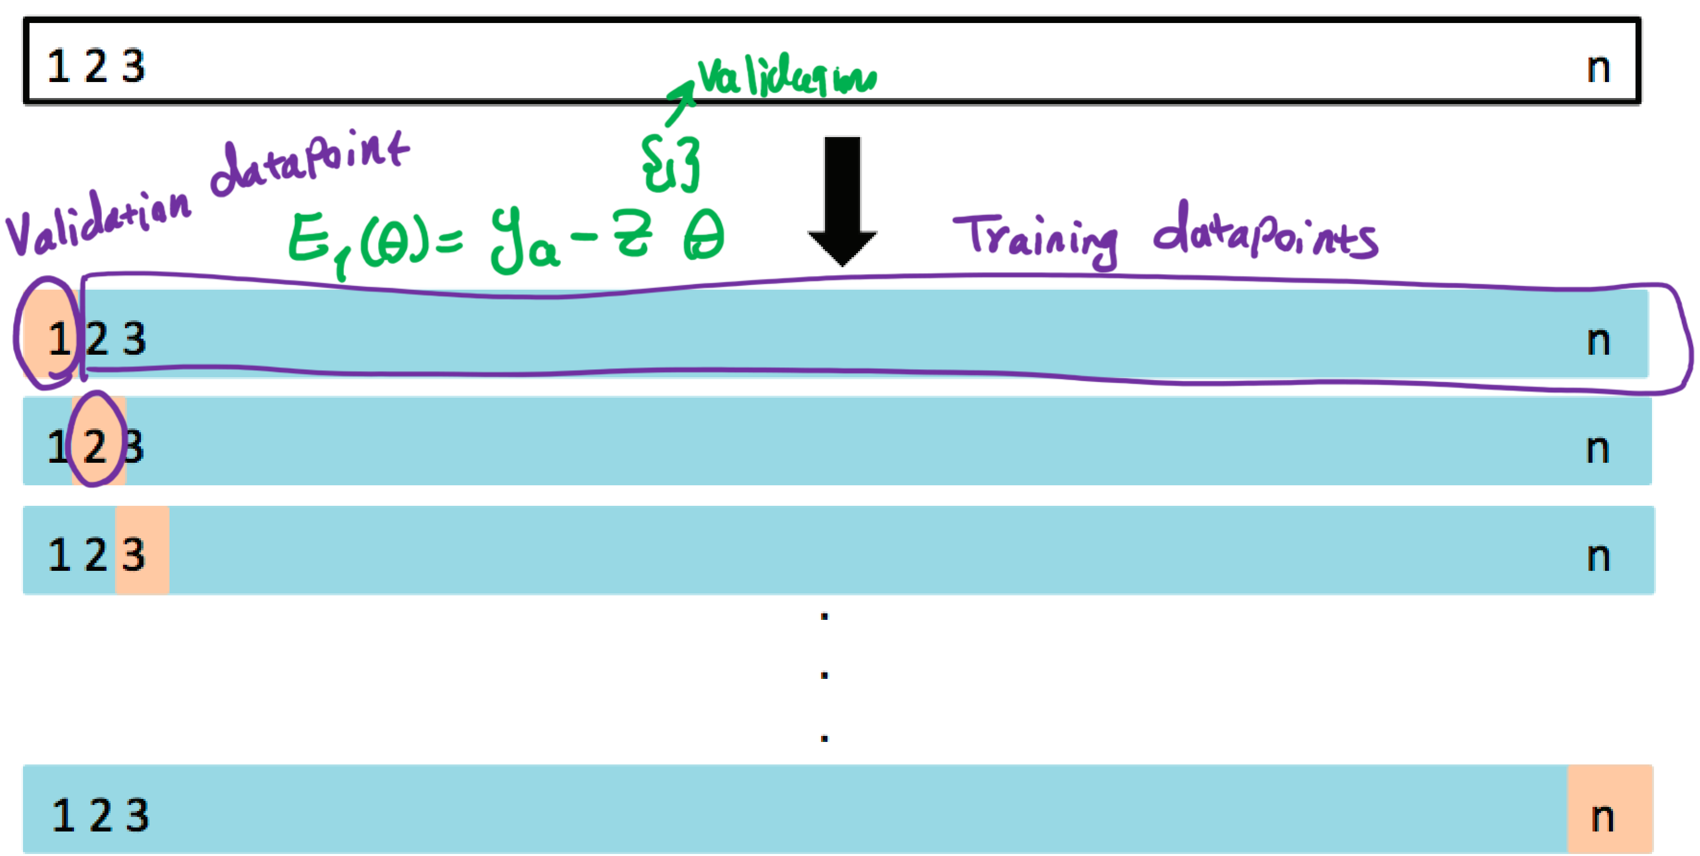

$E_{100}(\theta) = ?$

avg$(E_1(\theta) + \cdots E_{100}(\theta))$

K-Fold Cross Validation

Split the data into k subsets or *folds*.

For every $i = 1, \ldots, k$:

- train the model on every fold except the ith fold,
- compute the test error on the ith fold.

Average the test errors.

$\boxed{k = 10}$

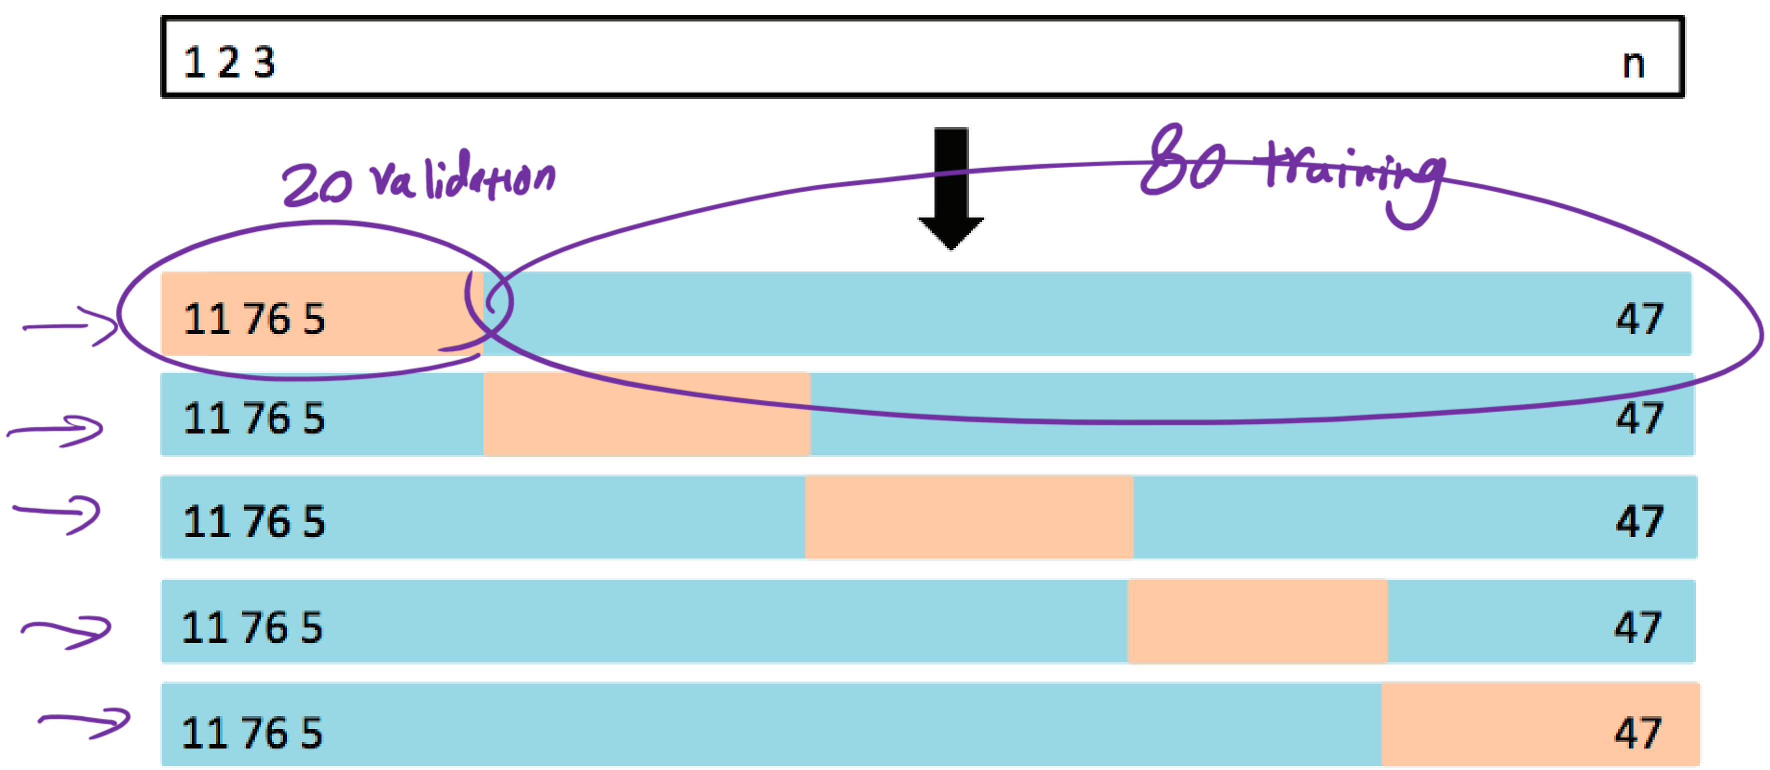

Choosing λ Using Validation Dataset

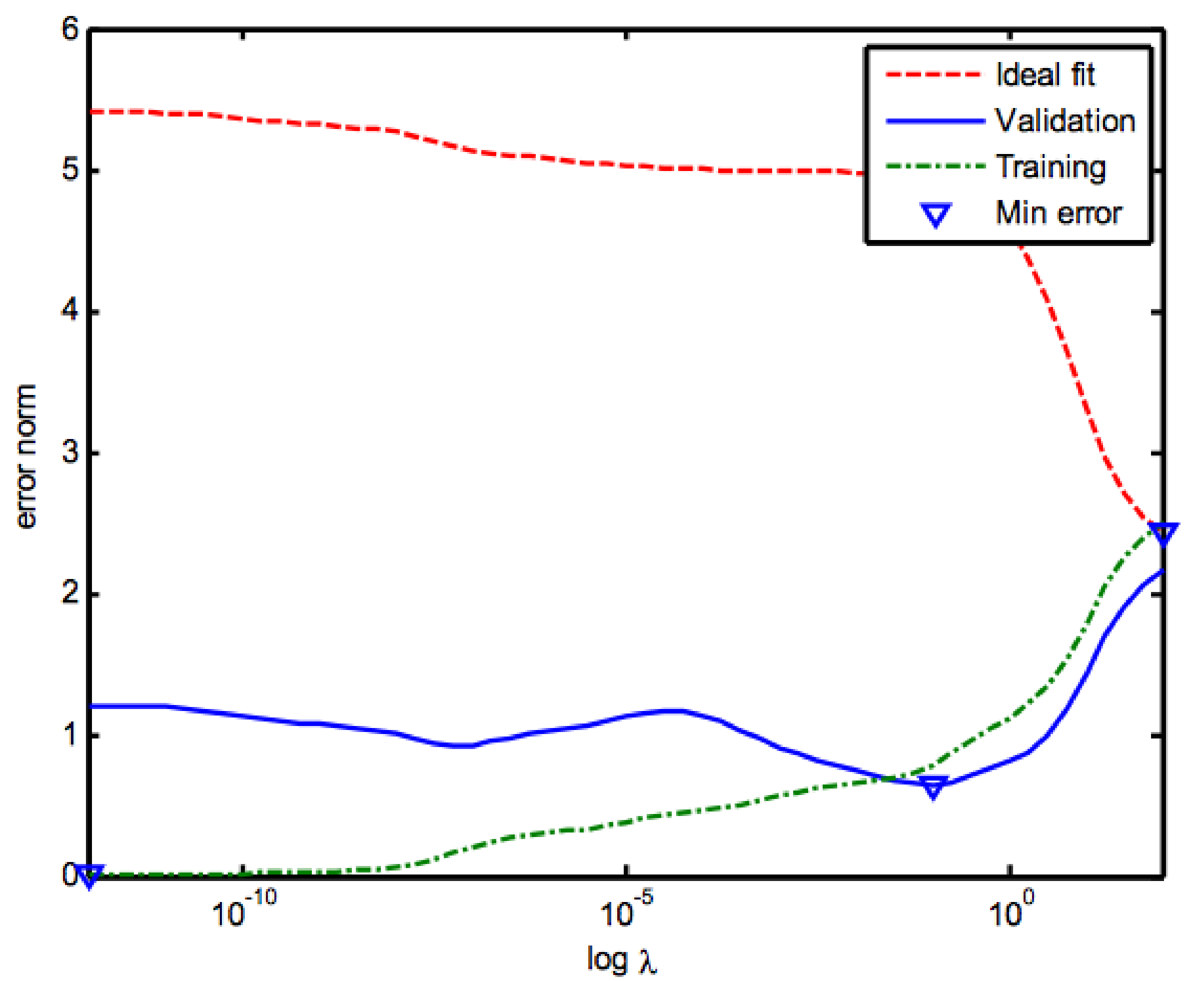

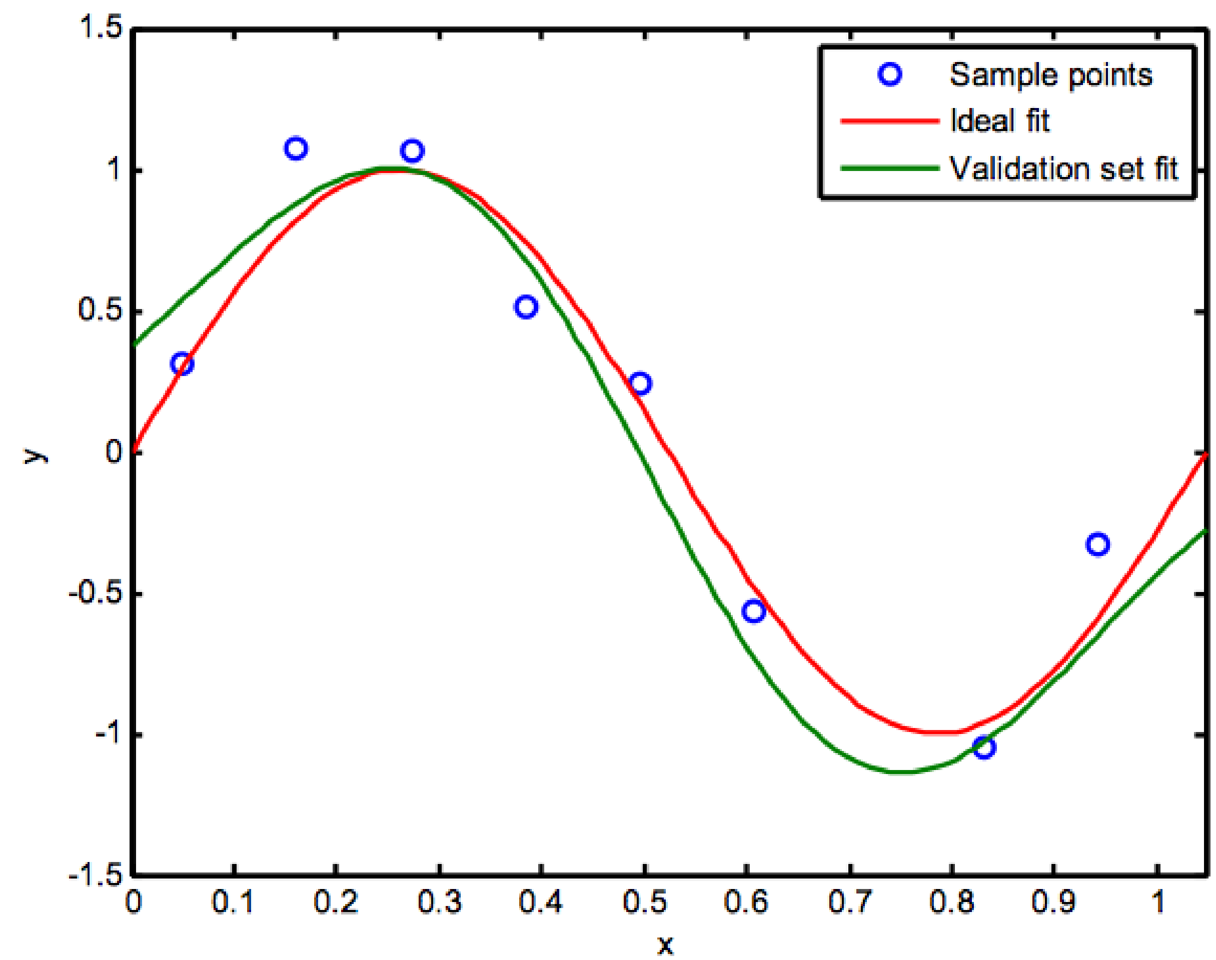

Pick up the lambda with the lowest mean value of rmse calculated by Cross Validation approach

Take-Home Messages

- What is overfitting
- What is regularization
- How does Ridge regression work
- Sparsity properties of Lasso regression
- How to choose the regularization coefficient λ