Generalization

CE417: Introduction to Artificial Intelligence
Sharif University of Technology
Fall 2023

Soleymani

Topics

- Beyond linear models
- Evaluation & model selection
- Regularization

Recall: Linear regression (squared loss)

- Linear regression functions

$\boldsymbol{w} = [w_0, w_1, \ldots, w_d]^T$ are the parameters we need to set.

$$g : \mathbb{R} \to \mathbb{R} \quad g(x; \boldsymbol{w}) = w_0 + w_1 x$$

$$g : \mathbb{R}^{\mathrm{d}} \to \mathbb{R} \quad g(\boldsymbol{x}; \boldsymbol{w}) = w_0 + w_1 x_1 + \ldots w_d x_d$$

- $J(\boldsymbol{w})$: Sum of squares error

$$J(\boldsymbol{w}) = \sum_{i=1}^{n} \left(y^{(i)} - g\left(\boldsymbol{x}^{(i)}; \boldsymbol{w}\right)\right)^2$$

- Weight update rule for $g(\boldsymbol{x}; \boldsymbol{w}) = \boldsymbol{w}^T \boldsymbol{x}$:

$$\boldsymbol{w}^{t+1} = \boldsymbol{w}^{t} + \eta \sum_{i=1}^{n} \left(y^{(i)} - {\boldsymbol{w}^{t}}^{T} \boldsymbol{x}^{(i)} \right) \boldsymbol{x}^{(i)}$$

Beyond linear regression

- How to extend the linear regression to non-linear functions?
  - Transform the data using basis functions
  - Learn a linear regression on the new feature vectors (obtained by basis functions)

Generalized linear

- Linear combination of fixed non-linear function of the input vector

$$g(\boldsymbol{x}; \boldsymbol{w}) = w_0 + w_1\phi_1(\boldsymbol{x}) + \dots w_m\phi_m(\boldsymbol{x})$$

$\{\phi_1(\boldsymbol{x}), \dots, \phi_m(\boldsymbol{x})\}$: set of basis functions (or features)

$$\phi_i(\boldsymbol{x}): \mathbb{R}^d \to \mathbb{R}$$

Basis functions: examples

- Linear

If $m = d$, $\phi_i(\mathbf{x}) = x_i$, $i = 1, \ldots, d$, then

$$f(\mathbf{x}; \mathbf{w}) = w_0 + w_1 x_1 + \ldots + w_d x_d$$

- Polynomial (univariate)

If $\phi_i(x) = x^i$, $i = 1, \ldots, m$, then

$$f(x; \mathbf{w}) = w_0 + w_1 x + \ldots + w_{m-1} x^{m-1} + w_m x^m$$

Polynomial regression: example

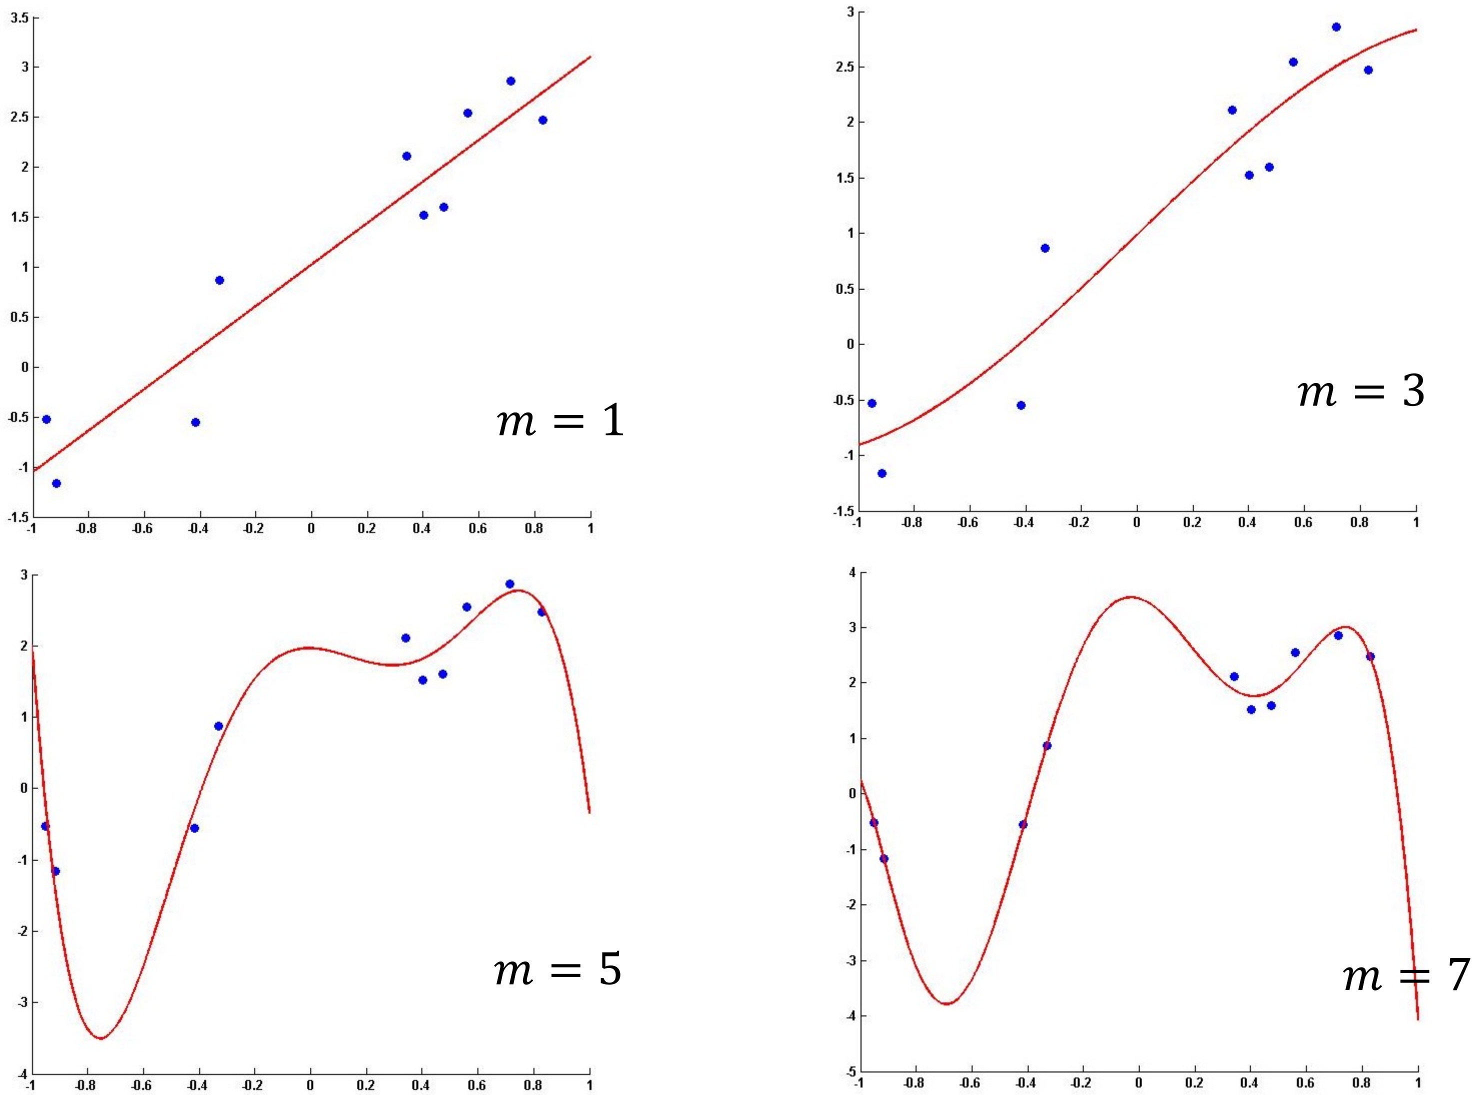

Classification: Not linearly separable data

- Non-linear decision surface: Transform to a new feature space

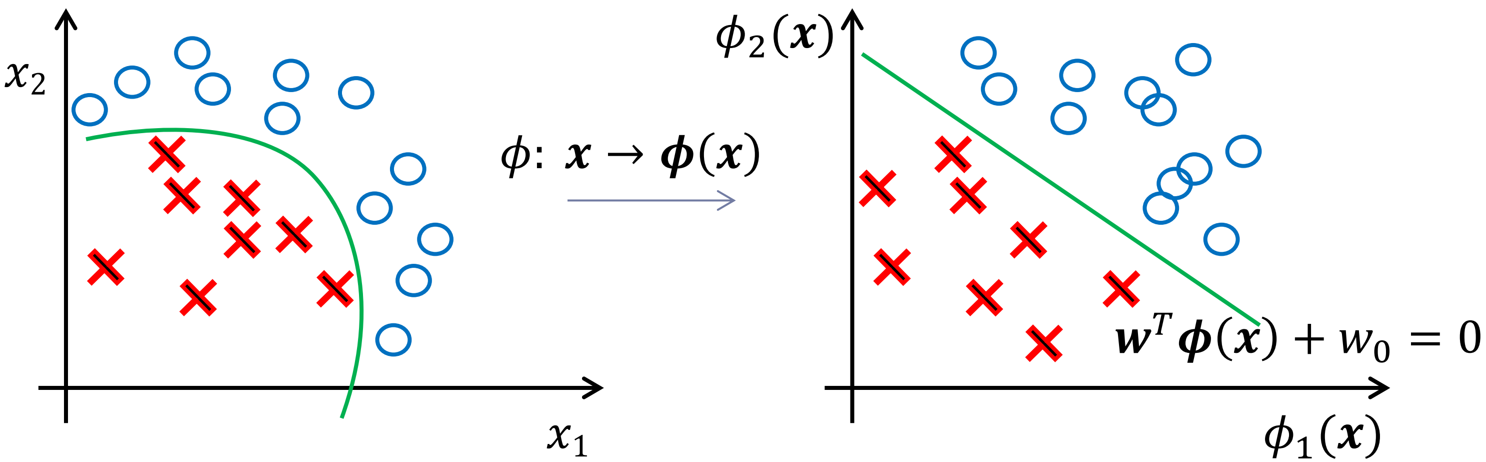

- Quardratic surfaces
  - Two dimensional feature space:

$$\boldsymbol{\phi}(\boldsymbol{x}) = [1, x_1, x_2, x_1^2, x_2^2, x_1 x_2]^T$$

  - d-dimentional feature space

$$\boldsymbol{\phi}(\boldsymbol{x}) = \left[1, x_1, \dots, x_d, x_1^2, \dots, x_d^2, x_1 x_2, \dots, x_1 x_d, x_2 x_3, \dots, x_{d-1} x_d\right]^T$$

Model complexity and overfitting

- With limited training data, models may achieve zero training error but a large test error.

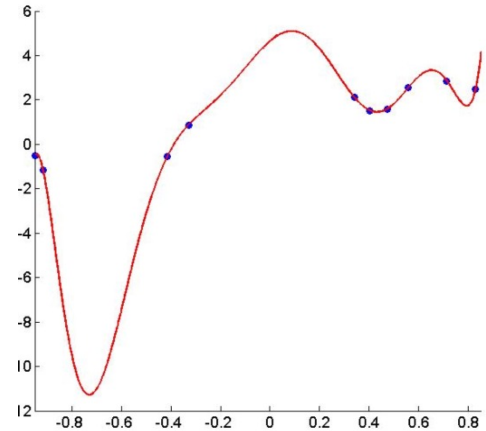

Training (empirical) loss

$$\frac{1}{n}\sum_{i=1}^{n}\left(y^{(i)} - f\left(\boldsymbol{x}^{(i)};\boldsymbol{\theta}\right)\right)^2 \approx 0$$

Expected (true) loss

$$\mathrm{E}_{\mathbf{x},\mathrm{y}}\left\{\left(y - f(\boldsymbol{x};\boldsymbol{\theta})\right)^2\right\} \gg 0$$

- Over-fitting: when the training loss no longer bears any relation to the test (generalization) loss.
  - Fails to generalize to unseen examples.

Polynomial regression

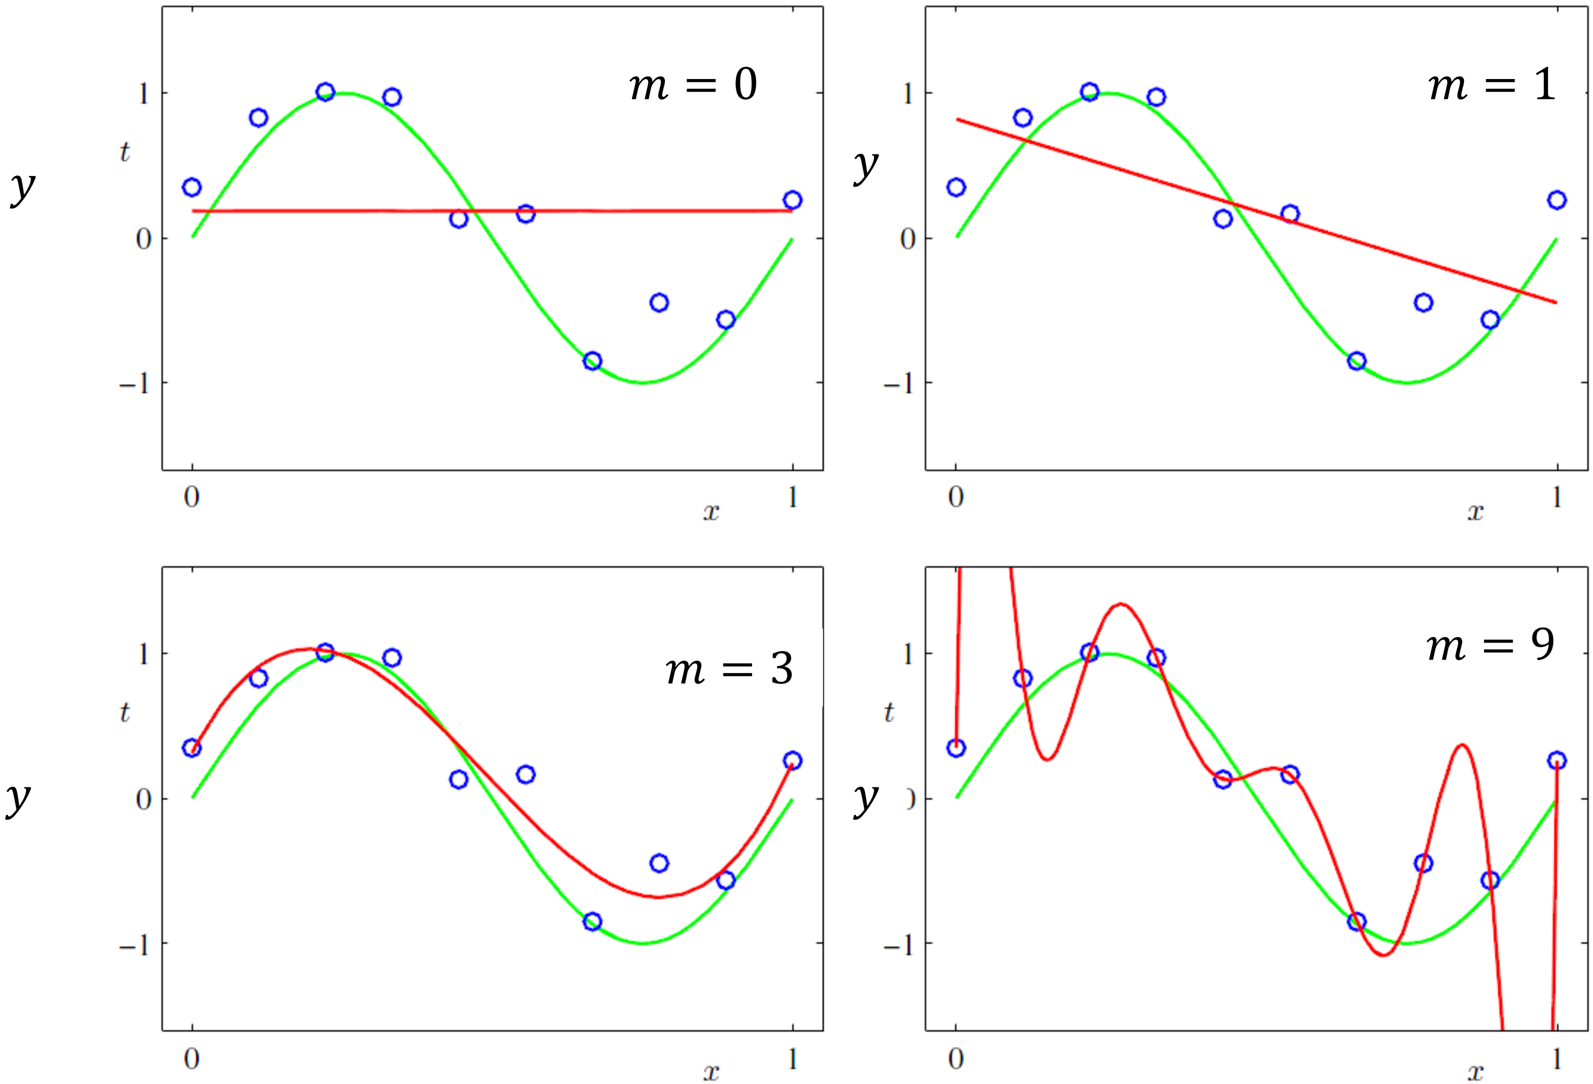

[Bishop]

Over-fitting causes

- Model complexity
  - E.g., Model with a large number of parameters (degrees of freedom)

- Low number of training data
  - Small data size compared to the complexity of the model

Model complexity

- Example:
  - Polynomials with larger m are becoming increasingly tuned to the random noise on the target values.

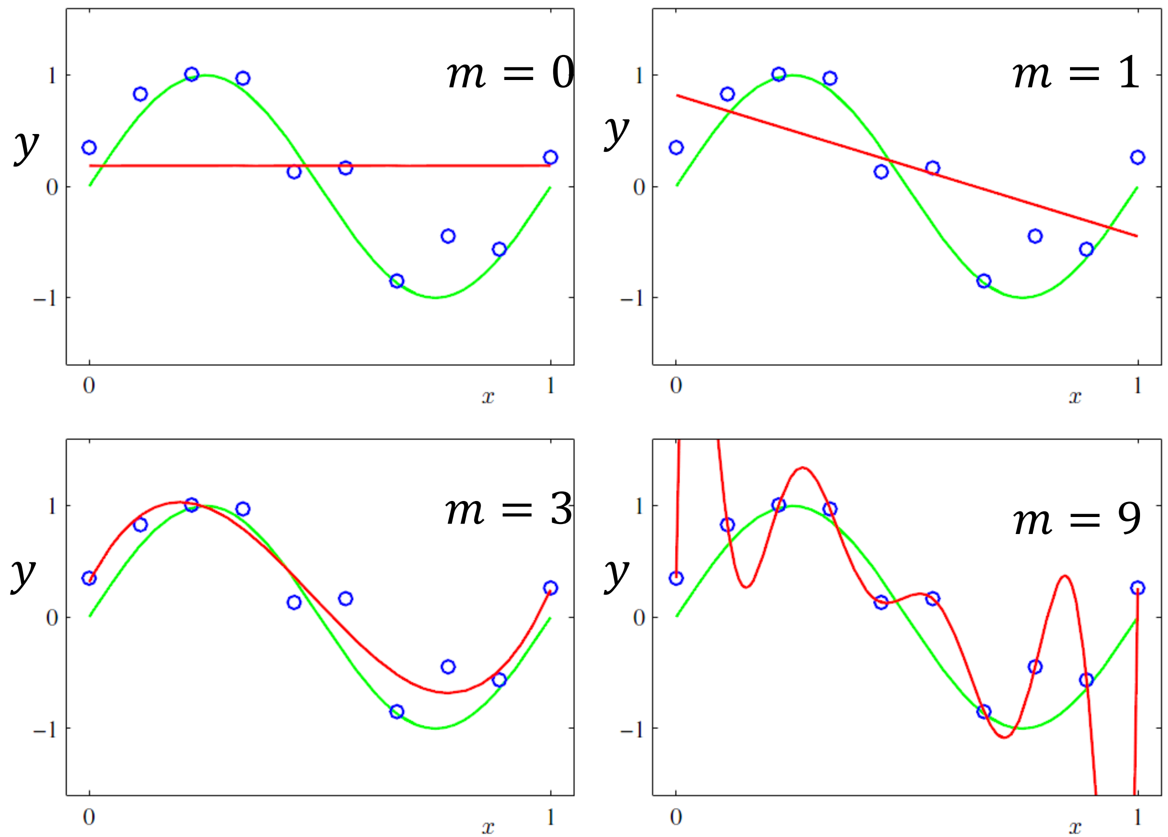

[Bishop]

Number of training data & overfitting

- Over-fitting problem becomes less severe as the size of training data increases.

$m = 9$

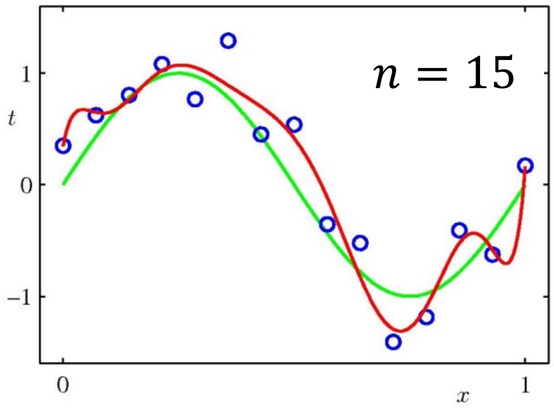

$m = 9$

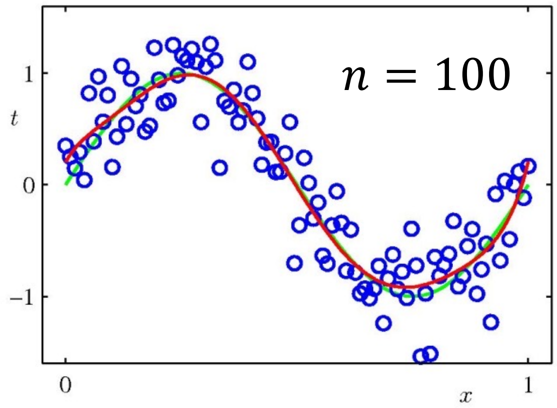

[Bishop]

Avoiding over-fitting

- Determine a suitable value for model complexity (Model Selection)
  - **Simple hold-out method**
  - **Cross-validation**

- Regularization (Occam's Razor)
  - Explicit preference towards simple models
  - Penalize for the model complexity in the objective function

Avoiding over-fitting

- Determine a suitable value for model complexity (Model Selection)
  - **Simple hold-out method**
  - **Cross-validation**

- Regularization (Occam's Razor)
  - Explicit preference towards simple models
  - Penalize for the model complexity in the objective function

Evaluation and model selection

- **Evaluation**:
  - We need to measure how well the learned function can predict the target for unseen examples

- **Model selection:**
  - Most of the time we need to select among a set of models
    - Example: polynomials with different degree m
  - and thus we need to evaluate these models first

Model selection

- **Learning algorithm** defines the data-driven search over the hypothesis space
  - Optimization of parameters

- **Hyper-parameters** are the tunable aspects of the model, that the learning algorithm does *not* select

This slide has been adopted from CMU ML course:
http://www.cs.cmu.edu/~mgormley/courses/10601-s18/

Model selection

- **Model selection** is the process by which we choose the "best" model among a set of candidates
  - assume access to a function capable of measuring the quality of a model
  - typically done "outside" the main training algorithm

- Model selection / hyper-parameter optimization is just another form of learning

This slide has been adopted from CMU ML course:
http://www.cs.cmu.edu/~mgormley/courses/10601-s18/

Simple hold-out: model selection

- Steps:
  - Divide training data into <u>training</u> and <u>validation set</u> v_set
  - Use only the training set to train a set of models
  - Evaluate each learned model on the validation set
    - $J_v(\boldsymbol{w}) = \frac{1}{|v\_set|} \sum_{i \in v\_set} \left(y^{(i)} - f\left(\boldsymbol{x}^{(i)}; \boldsymbol{w}\right)\right)^2$
  - Choose the best model based on the validation set error

Simple hold-out: model selection

- Steps:
  - Divide training data into training and validation set v_set
  - Use only the training set to train a set of models
  - Evaluate each learned model on the validation set
    - $J_v(\boldsymbol{w}) = \frac{1}{|v\_set|} \sum_{i \in v\_set} \left(y^{(i)} - f\left(\boldsymbol{x}^{(i)}; \boldsymbol{w} \right) \right)^2$
  - Choose the best model based on the validation set error

- Usually, too wasteful of valuable training data
  - Training data may be limited.
  - On the other hand, small validation set obtains a relatively noisy estimate of performance.

Simple hold out: training, validation, and test sets

- Simple hold-out chooses the model that minimizes error on validation set.

- $J_v(\hat{\boldsymbol{w}})$ is likely to be an optimistic estimate of generalization error.
  - extra parameter (e.g., degree of polynomial) is fit to this set.

- Estimate generalization error for the test set
  - performance of the selected model is finally evaluated on the test set

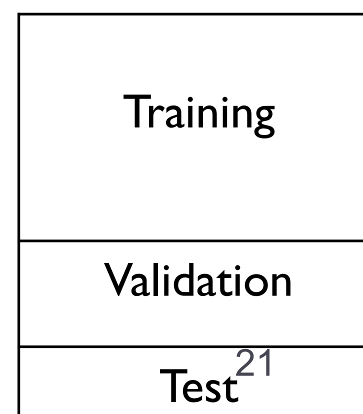

Cross-Validation (CV): evaluation

- k-fold cross-validation steps:
  - Shuffle the dataset and randomly partition training data into k groups of approximately equal size
  - for $i = 1$ to k
    - Choose the i-th group as the held-out validation group
    - Train the model on all but the i-th group of data
    - Evaluate the model on the held-out group

Cross-Validation (CV): evaluation

- k-fold cross-validation steps:
  - Shuffle the dataset and randomly partition training data into k groups of approximately equal size
  - for $i = 1$ to k
    - Choose the i-th group as the held-out validation group
    - Train the model on all but the i-th group of data
    - Evaluate the model on the held-out group
  - Performance scores of the model from k runs are **averaged**.
    - The average error rate as an estimation of the true performance of the model.

Cross-Validation (CV): model selection

- For each model, we first find the average error by CV.

- The model with **the best average performance** is selected.

Cross-validation: polynomial regression example

- 5-fold CV
- 100 runs
  - average

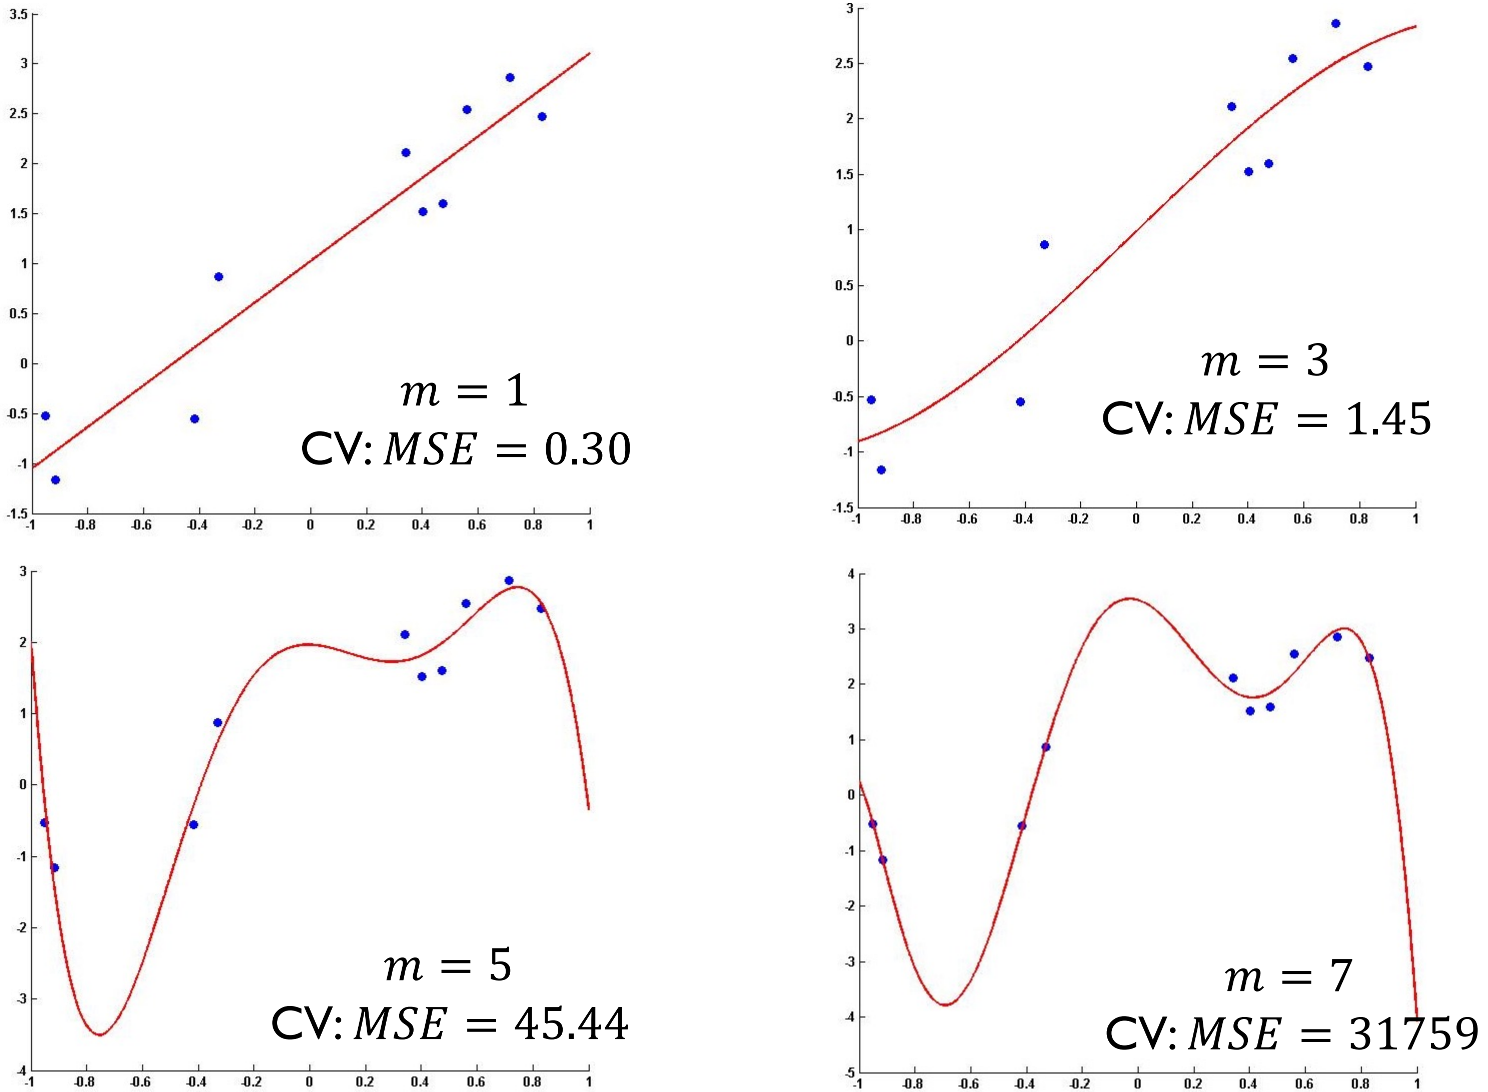

Avoiding over-fitting

- Determine a suitable value for model complexity (Model Selection)
  - **Simple hold-out method**
  - **Cross-validation**

- Regularization (Occam's Razor)
  - Explicit preference towards simple models
  - Penalize for the model complexity in the objective function

Regularization

- Adding a penalty term in the cost function to discourage the coefficients from reaching large values.

Regularization in regression problem

- Adding a penalty term in the cost function to discourage the coefficients from reaching large values.

- **Ridge regression** (weight decay):

$$J(\boldsymbol{w}) = \sum_{i=1}^{n} \left(y^{(i)} - \boldsymbol{w}^T \boldsymbol{\phi}\left(\boldsymbol{x}^{(i)}\right)\right)^2 + \lambda \boldsymbol{w}^T \boldsymbol{w}$$

Regularization in regression problem

- Adding a penalty term in the cost function to discourage the coefficients from reaching large values.

- Ridge regression (weight decay):

$$J(\boldsymbol{w}) = \sum_{i=1}^{n} \left(y^{(i)} - \boldsymbol{w}^T \boldsymbol{\phi}\left(\boldsymbol{x}^{(i)}\right) \right)^2 + \lambda \boldsymbol{w}^T \boldsymbol{w}$$

- Weight update by gradient descent:

$$\boldsymbol{w}_j^{t+1} = \boldsymbol{w}_j^{t} - \eta \nabla_{\boldsymbol{W}} J(\boldsymbol{W}^t)$$

$$\nabla_{\boldsymbol{w}} J(\boldsymbol{w}) = -2 \sum_{i=1}^{n} \left(y^{(i)} - \boldsymbol{w}^T \boldsymbol{\phi}\left(\boldsymbol{x}^{(i)}\right) \right) \boldsymbol{\phi}\left(\boldsymbol{x}^{(i)}\right) + 2\lambda \boldsymbol{w}$$

Regularization in classification problem

- Multi-class logistic regression (i.e., cross entropy loss) with regularization:

$$J(\boldsymbol{W}) = -\sum_{i=1}^{n}\sum_{k=1}^{K} y_k^{(i)} \log\left(g_k\left(\boldsymbol{x}^{(i)}; \boldsymbol{W}\right)\right) + \lambda \sum_{k=1}^{K} \boldsymbol{w}_k^T \boldsymbol{w}_k$$

- Weight Update:

$$\boldsymbol{w}_k^{t+1} = \boldsymbol{w}_k^{t} - \eta \nabla_{\boldsymbol{w}_k} J(\boldsymbol{W}^t)$$

$$\nabla_{\boldsymbol{w}_k} J(\boldsymbol{W}) = -2\sum_{i=1}^{n}\left(y^{(i)} - g_k(\boldsymbol{x}^{(i)}; \boldsymbol{W})\right)\boldsymbol{x}^{(i)} + 2\lambda \boldsymbol{w}_k$$

Regression: polynomial order

- Polynomials with larger m are becoming increasingly tuned to the random noise on the target values.
  - magnitude of the coefficients typically gets larger by increasing m.

	$M=0$	$M=1$	$M=6$	$M=9$
$w_0^\star$	0.19	0.82	0.31	0.35
$w_1^\star$		-1.27	7.99	232.37
$w_2^\star$			-25.43	-5321.83
$w_3^\star$			17.37	48568.31
$w_4^\star$				-231639.30
$w_5^\star$				640042.26
$w_6^\star$				-1061800.52
$w_7^\star$				1042400.18
$w_8^\star$				-557682.99
$w_9^\star$				125201.43

[Bishop]

Regression: regularization parameter

$m = 9$

	$\ln \lambda = -\infty$	$\ln \lambda = -18$	$\ln \lambda = 0$
$\widehat{w}_0$	0.35	0.35	0.13
$\widehat{w}_1$	232.37	4.74	-0.05
$\widehat{w}_2$	-5321.83	-0.77	-0.06
$\widehat{w}_3$	48568.31	-31.97	-0.05
$\widehat{w}_4$	-231639.30	-3.89	-0.03
$\widehat{w}_5$	640042.26	55.28	-0.02
$\widehat{w}_6$	-1061800.52	41.32	-0.01
$\widehat{w}_7$	1042400.18	-45.95	-0.00
$\widehat{w}_8$	-557682.99	-91.53	0.00
$\widehat{w}_9$	125201.43	72.68	0.01

[Bishop]

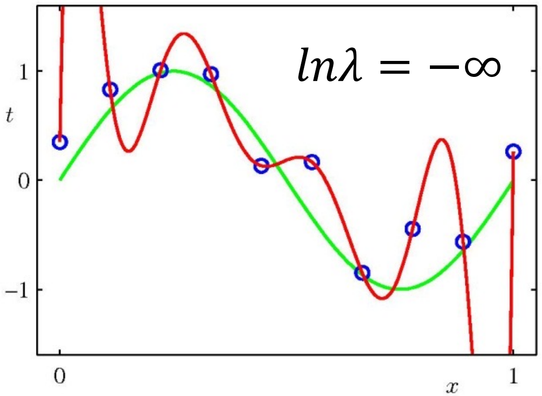

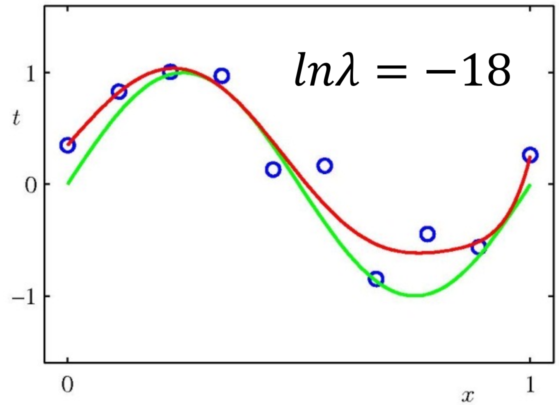

Regularization parameter

- Generalization
  - λ now controls the effective complexity of the model and hence determines the degree of over-fitting

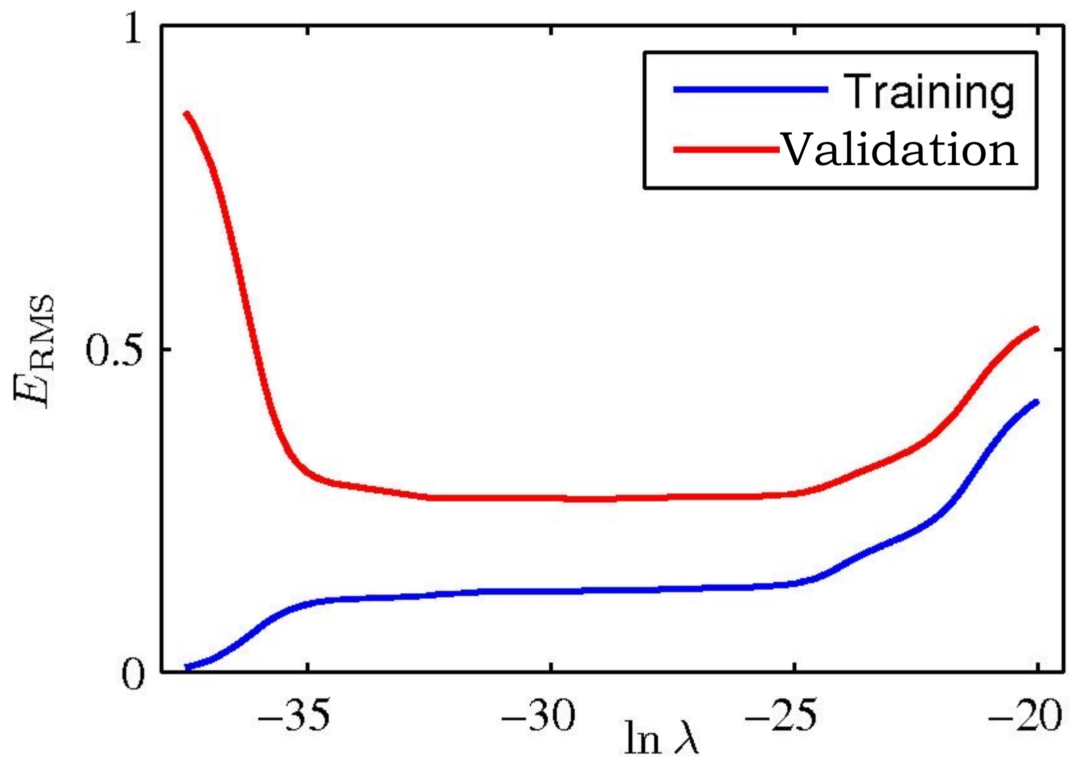

[Bishop]

Choosing the regularization parameter

- A set of models with different values of λ.

- Find $\widehat{\boldsymbol{w}}$ for each model based on training data

- Find $J_v(\widehat{\boldsymbol{w}})$ (or $J_{cv}(\widehat{\boldsymbol{w}})$) for each model
  - $J_v(\boldsymbol{w}) = \frac{1}{n\_v} \sum_{i \in v\_set} \left(y^{(i)} - f\left(x^{(i)}; \boldsymbol{w}\right) \right)^2$

- Select the model with the best $J_v(\widehat{\boldsymbol{w}})$ (or $J_{cv}(\widehat{\boldsymbol{w}})$)

The approximation-generalization trade-off

- Small true error shows good approximation of f out of sample

- More complex $\mathcal{H} \Rightarrow$ better chance of approximating f

- Less complex $\mathcal{H} \Rightarrow$ better chance of generalization out of f

Complexity of hypothesis space: example

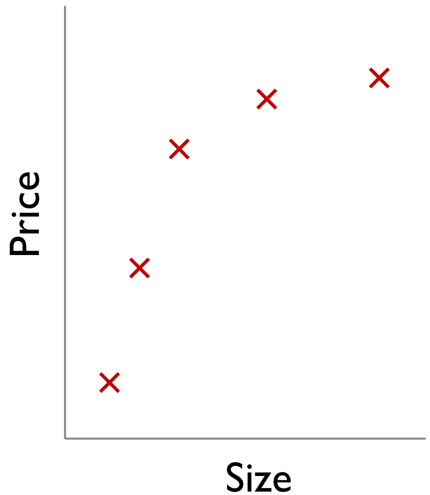

$w_0 + w_1 x$

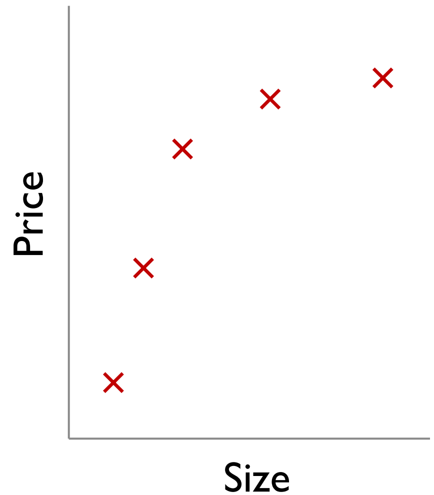

$w_0 + w_1 x + w_2 x^2$

$w_0 + w_1 x + w_2 x^2 + w_3 x^3 + w_4 x^4$

Less complex $\mathcal{H}$

More complex $\mathcal{H}$

This example has been adapted from: Prof. Andrew Ng's slides

Complexity of hypothesis space: example

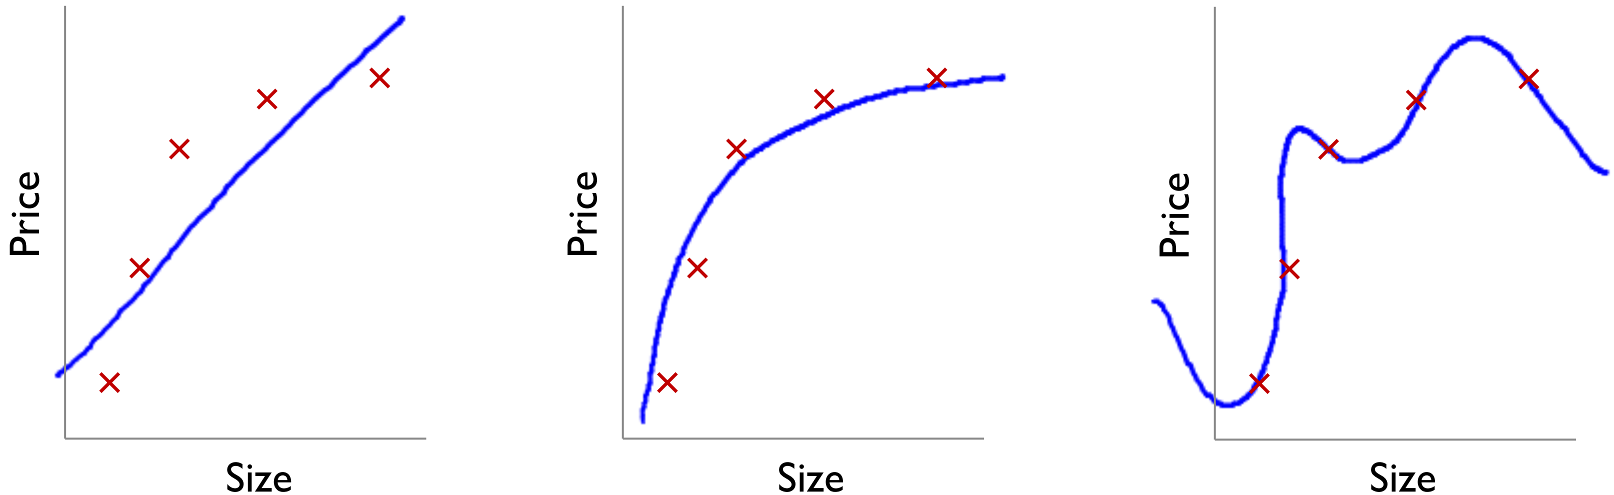

$w_0 + w_1 x$ $\quad$ $w_0 + w_1 x + w_2 x^2$ $\quad$ $w_0 + w_1 x + w_2 x^2 + w_3 x^3 + w_4 x^4$

Underfitting $\quad$ Overfitting

This example has been adapted from: Prof. Andrew Ng's slides

Complexity of hypothesis space: example

$$J_v(\boldsymbol{w}) = \frac{1}{n\_v} \sum_{i \in val\_set} \left(y^{(i)} - f\left(\boldsymbol{x}^{(i)}; \boldsymbol{w}\right) \right)^2$$

$$J_{train}(\boldsymbol{w}) = \frac{1}{n\_train} \sum_{i \in train\_set} \left(y^{(i)} - f\left(\boldsymbol{x}^{(i)}; \boldsymbol{w}\right) \right)^2$$

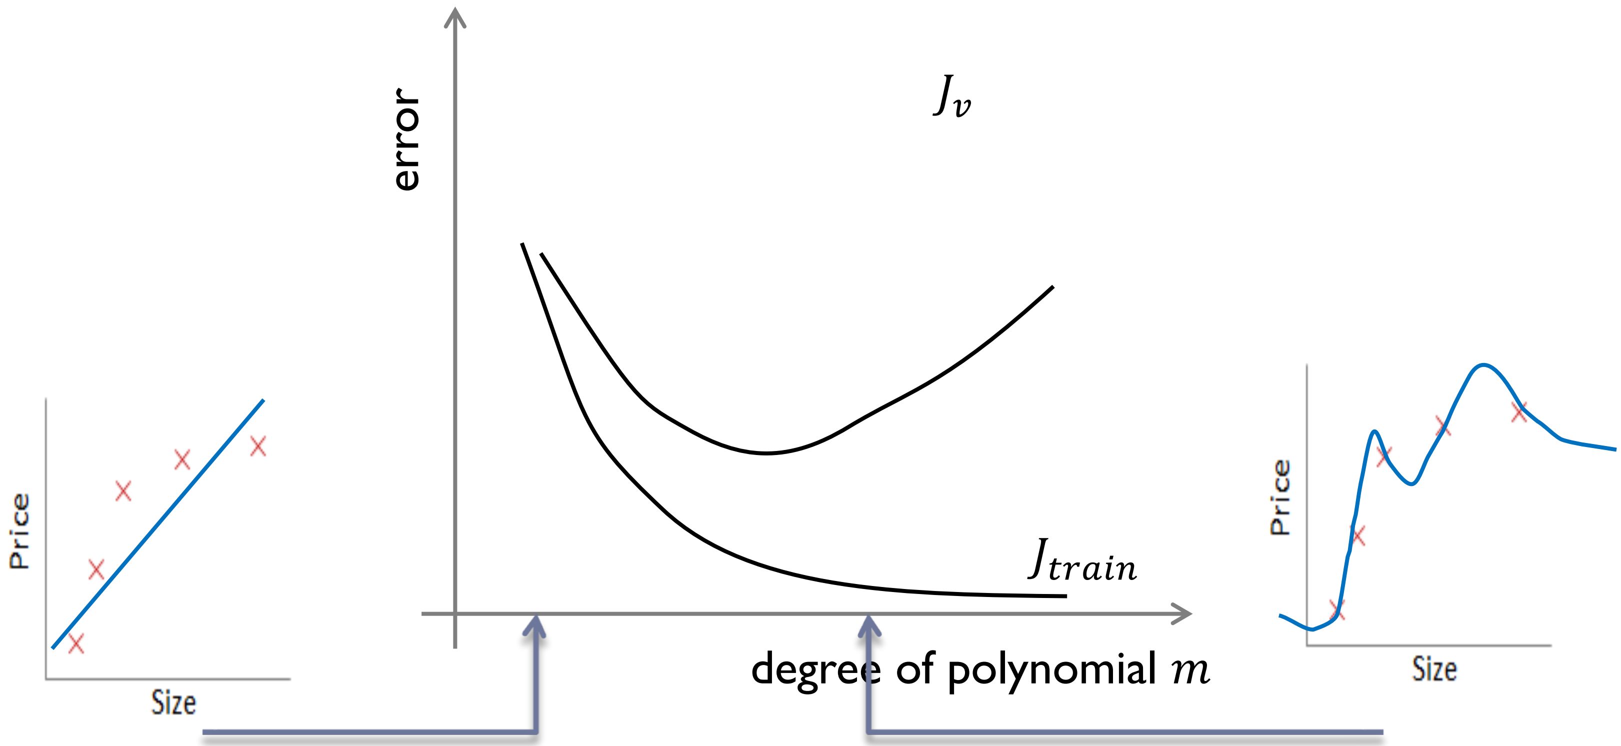

Complexity of hypothesis space

- Less complex $\mathcal{H}$:
  - $J_{train}(\widehat{\boldsymbol{w}}) \approx J_v(\widehat{\boldsymbol{w}})$ and $J_{train}(\widehat{\boldsymbol{w}})$ is very high

- More complex $\mathcal{H}$:
  - $J_{train}(\widehat{\boldsymbol{w}}) \ll J_v(\widehat{\boldsymbol{w}})$ and $J_{train}(\widehat{\boldsymbol{w}})$ is low

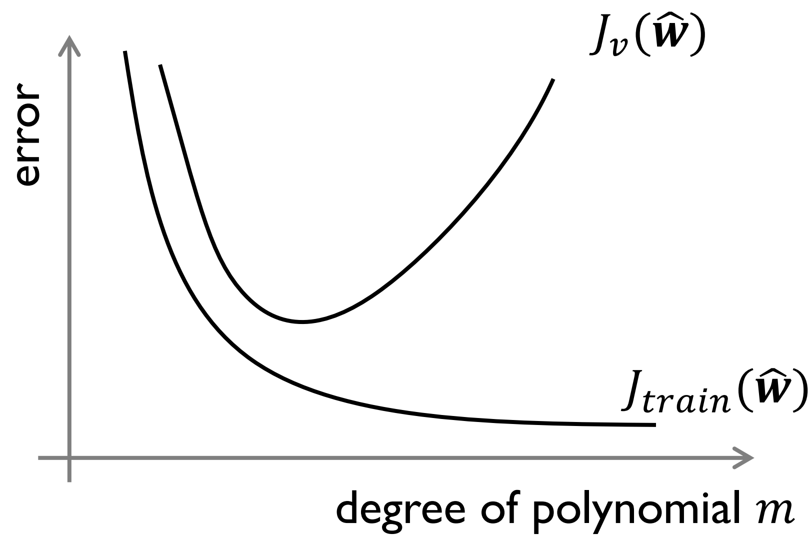

Size of training set

$$J_v(\boldsymbol{w}) = \frac{1}{n\_v} \sum_{i \in val\_set} \left(y^{(i)} - f\left(x^{(i)}; \boldsymbol{w} \right) \right)^2$$

$$J_{train}(\boldsymbol{w}) = \frac{1}{n\_train} \sum_{i \in train\_set} \left(y^{(i)} - f\left(x^{(i)}; \boldsymbol{w} \right) \right)^2$$

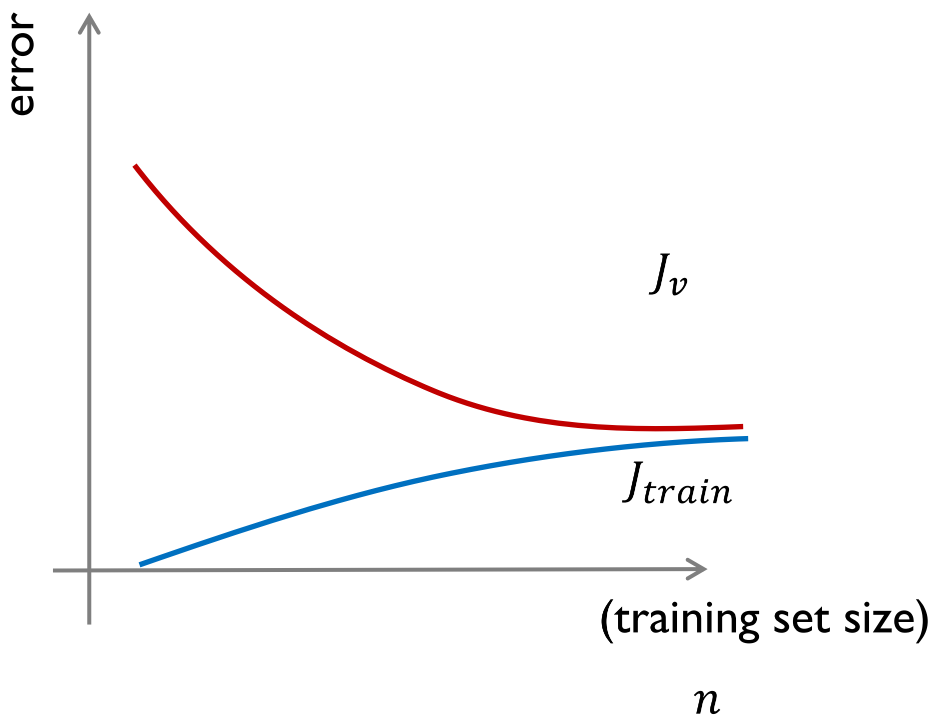

$$f(x; \boldsymbol{w}) = w_0 + w_1 x + w_2 x^2$$

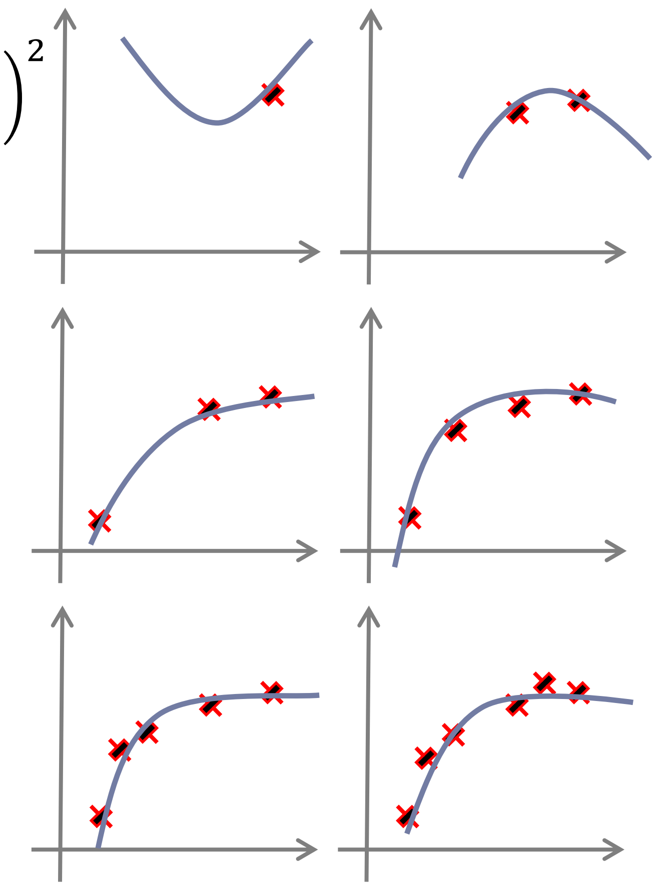

This slide has been adapted from: Prof. Andrew Ng's slides

Less complex $\mathcal{H}$

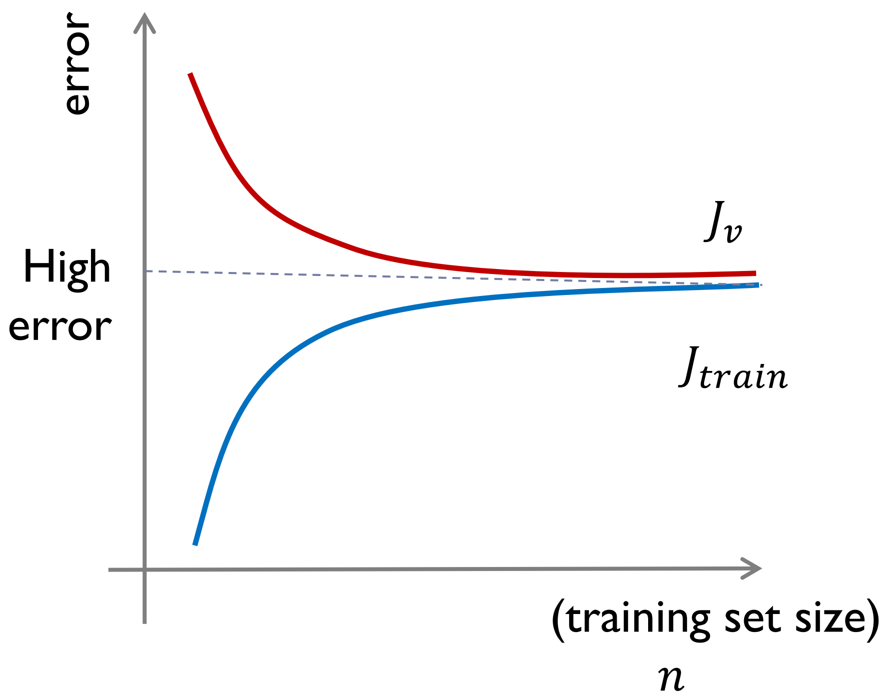

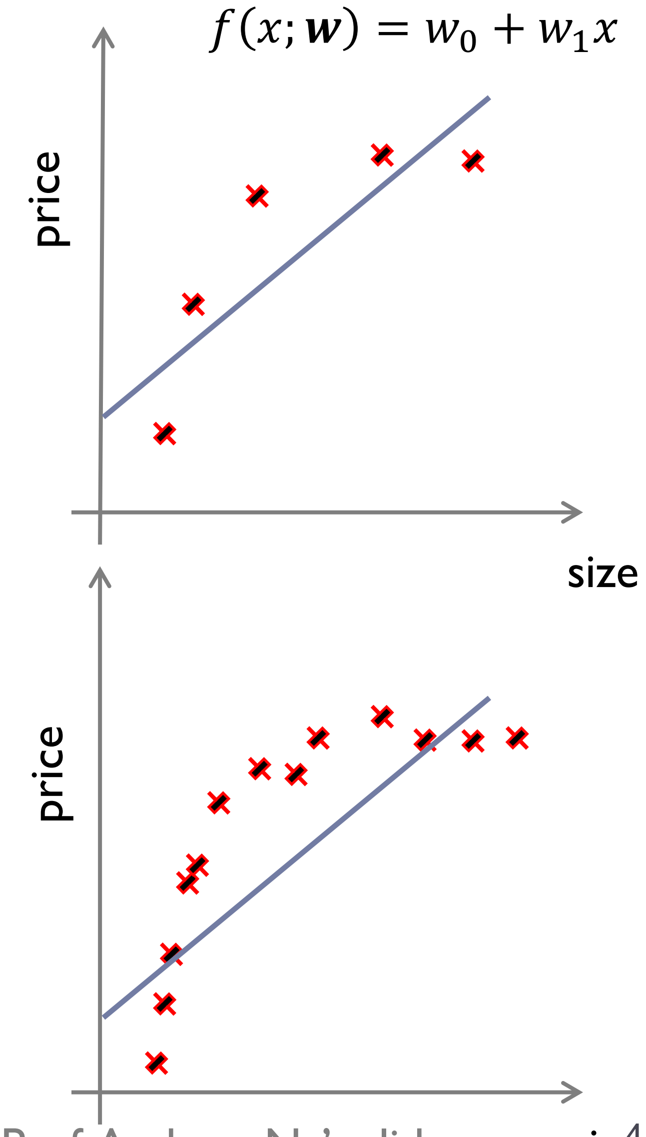

If model is very simple, getting more training data will not (by itself) help much.

This slide has been adapted from: Prof. Andrew Ng's slides

More complex $\mathcal{H}$

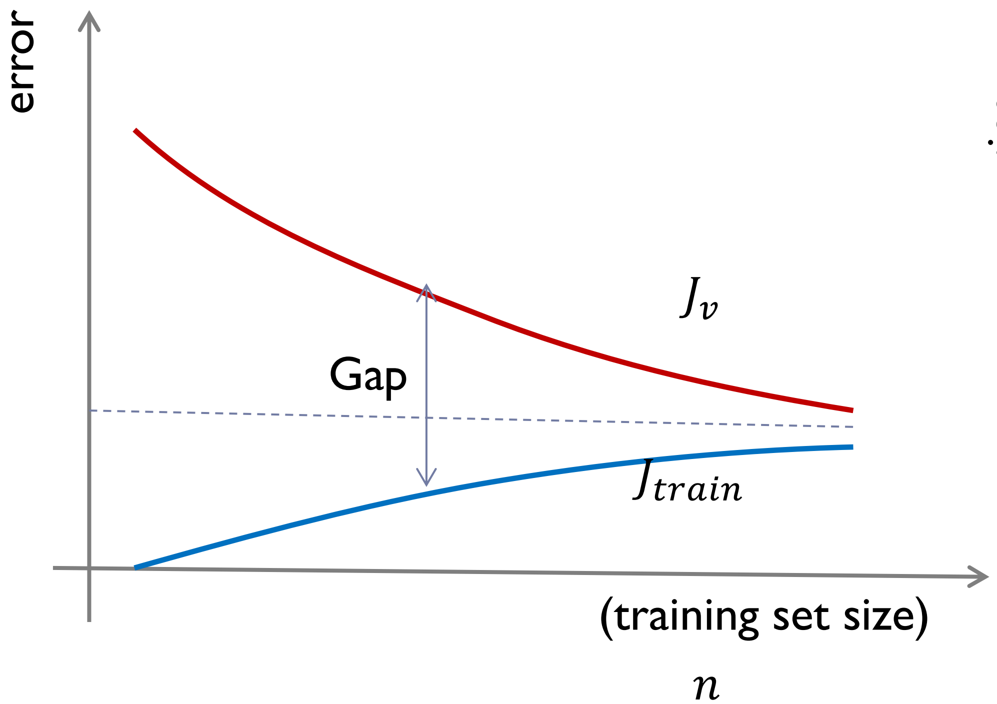

$$f(x; \boldsymbol{w}) = w_0 + w_1 x + \cdots w_{10} x^{10}$$

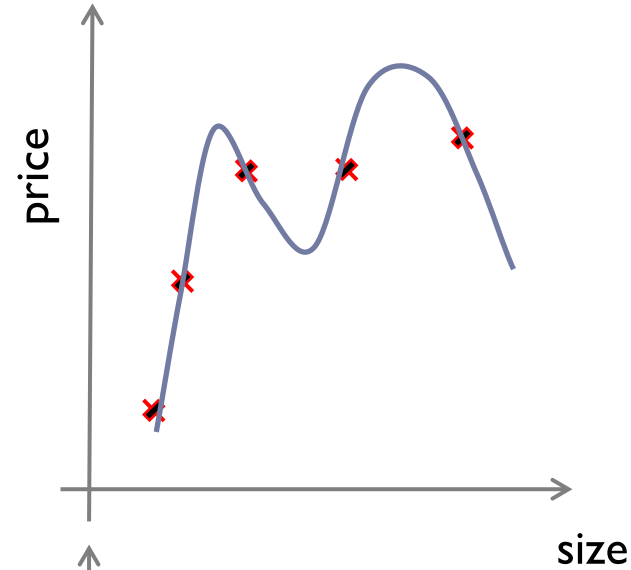

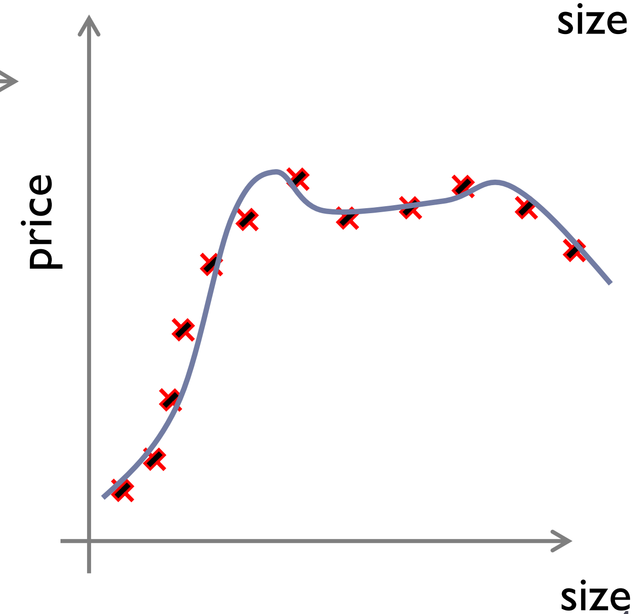

For more complex models, getting more training data is usually helps.

This slide has been adapted from: Prof. Andrew Ng's slides

Regularization: example

$$f(x;\boldsymbol{w}) = w_0 + w_1 x + w_2 x^2 + w_3 x^3 + w_4 x^4$$

$$J(\boldsymbol{w}) = \frac{1}{n}\left(\sum_{i=1}^{n}\left(y^{(i)} - f\left(x^{(i)};\boldsymbol{w}\right)\right)^2 + \lambda \boldsymbol{w}^T \boldsymbol{w}\right)$$

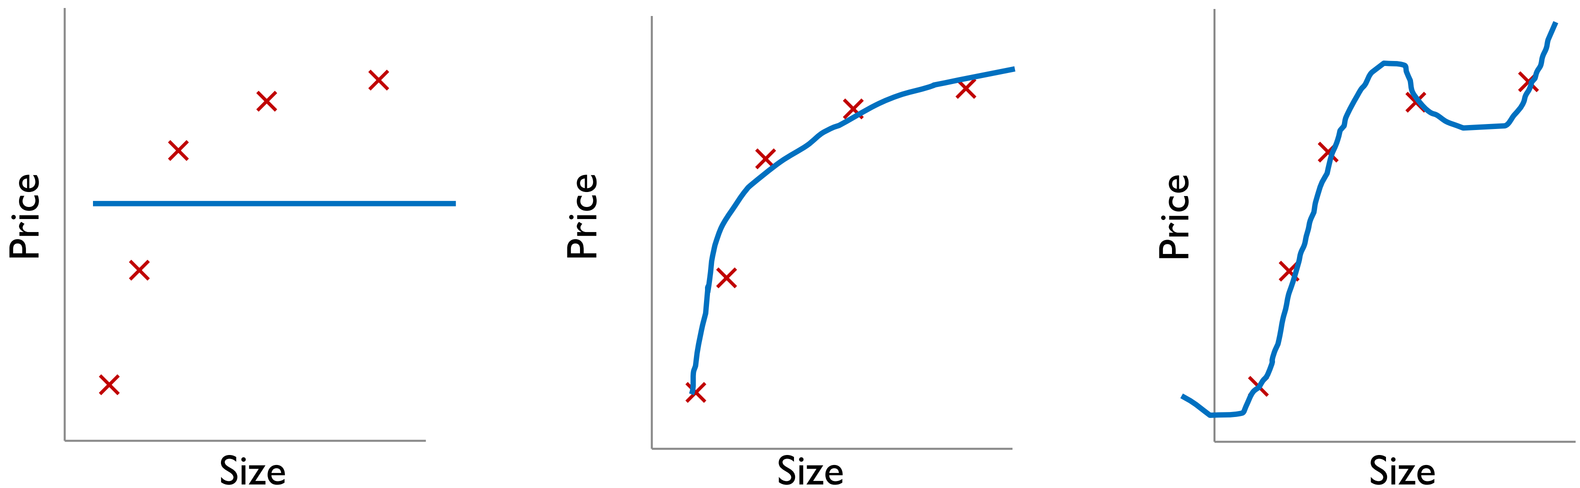

Large λ	Intermediate λ	Small λ
(Prefer to more simple models)		(Prefer to more complex models)
$w_1 = w_2 \approx 0$		$\lambda = 0$

This example has been adapted from: Prof. Andrew Ng's slides

Model complexity: Bias-variance trade-off

- Least squares can lead to severe over-fitting if complex models are trained using data sets of limited size.

- A frequentist viewpoint of the model complexity issue, known as the *bias-variance* trade-off.

The learning diagram: deterministic target

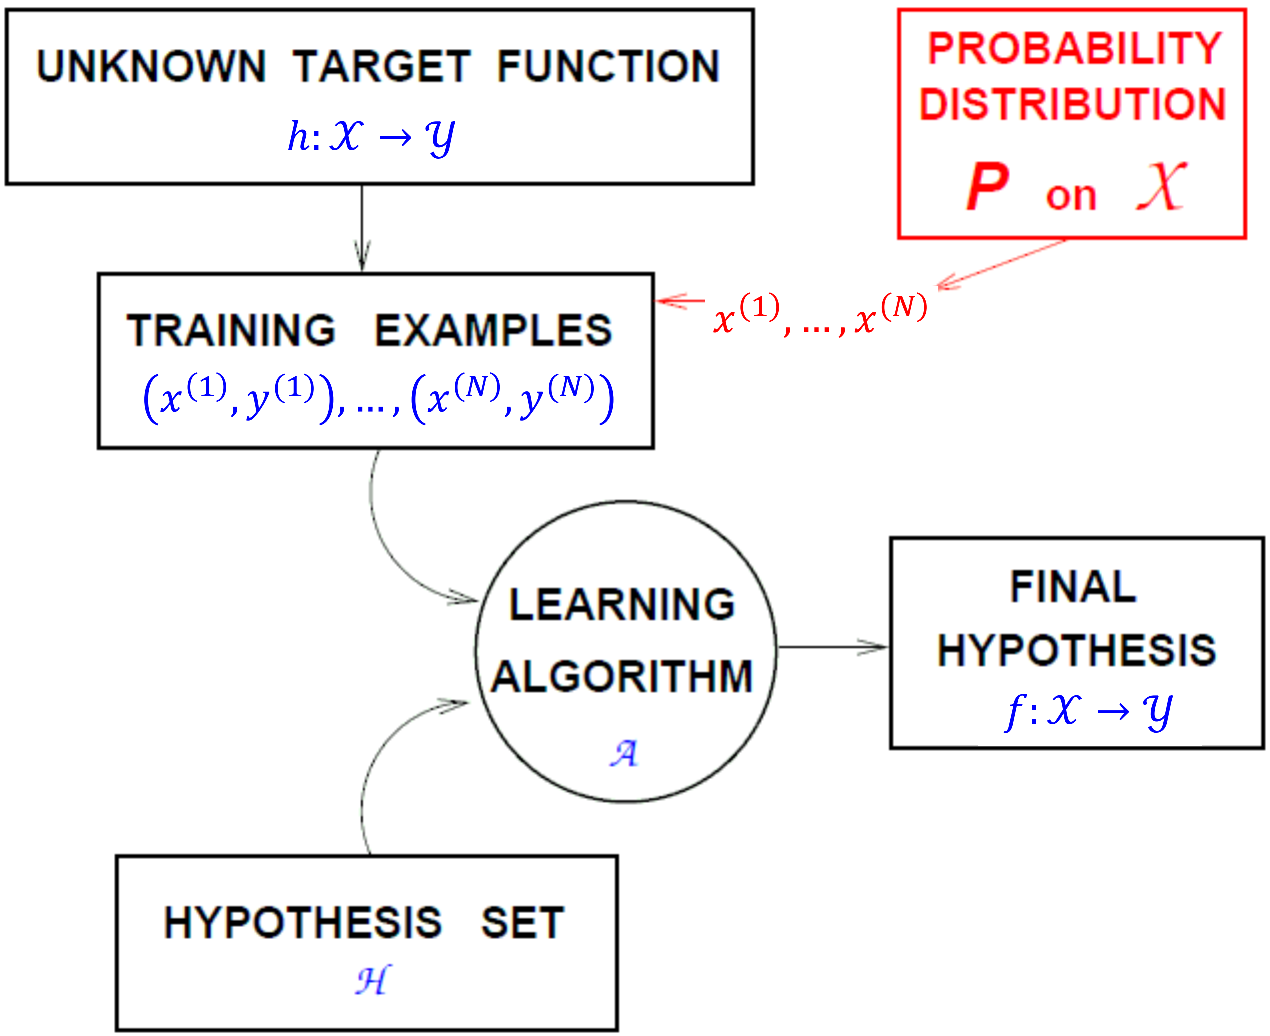

[Y.S. Abou Mostafa, 2012]

The learning diagram including noisy target

- Type

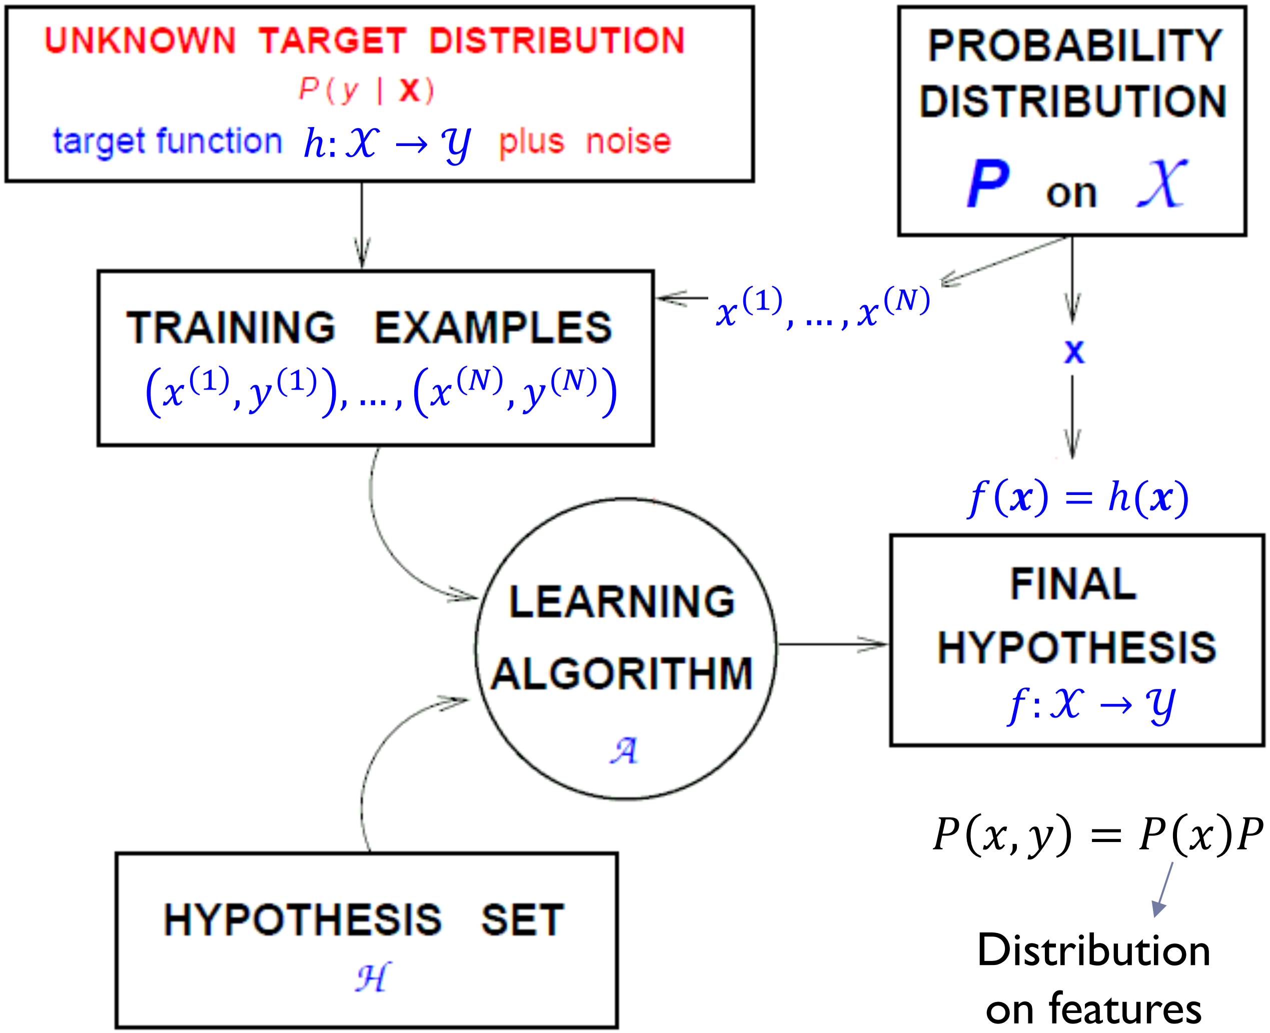

$$P(x,y) = P(x)P(y|x)$$

Distribution on features

Target distribution

[Y.S. Abou Mostafa, 2012]

Expectation of true error

$(\boldsymbol{x}, y) \sim P$

$h(\boldsymbol{x})$: minimizes the expected loss

$$E_{true}\left(f_{\mathcal{D}}(\boldsymbol{x})\right) = \mathbb{E}_{\boldsymbol{x},y}[(f_{\mathcal{D}}(\boldsymbol{x}) - y)^2]$$
$$= \mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right] + noise$$

$$\mathbb{E}_{\mathcal{D}}\left[\mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]\right]$$
$$= \mathbb{E}_{\boldsymbol{x}}\left[\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]\right]$$

We now want to focus on $\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$.

The average hypothesis

$$\bar{f}(\boldsymbol{x}) \equiv E_{\mathcal{D}}[f_{\mathcal{D}}(\boldsymbol{x})]$$

$$\bar{f}(\boldsymbol{x}) \approx \frac{1}{K}\sum_{k=1}^{K} f_{\mathcal{D}^{(k)}}(\boldsymbol{x})$$

K training sets (of size N) sampled from $P(\boldsymbol{x}, y)$: $\mathcal{D}^{(1)}, \mathcal{D}^{(2)}, \dots, \mathcal{D}^{(K)}$

Using the average hypothesis

$$\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x}) + \bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

Using the average hypothesis

$$\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x}) + \bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2 + \left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2 + 2\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)\left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)\right]$$

Using the average hypothesis

$$\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x}) + \bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2 + \left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2 + 2\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)\left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2\right] + \left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2$$

Bias and variance

$$\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right] = \underbrace{\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2\right]}_{\mathrm{var}(\boldsymbol{x})} + \underbrace{\left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2}_{\mathrm{bias}(\boldsymbol{x})}$$

$$\mathbb{E}_{\boldsymbol{x}}\left[\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]\right] = \mathbb{E}_{\boldsymbol{x}}[\mathrm{var}(\boldsymbol{x}) + \mathrm{bias}(\boldsymbol{x})]$$

$$= \mathrm{var} + \mathrm{bias}$$

Bias-variance trade-off

$$\text{var} = \mathbb{E}_{\boldsymbol{x}}\left[\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2\right]\right]$$

$$\text{bias} = \mathbb{E}_{\boldsymbol{x}}\left[\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right]$$

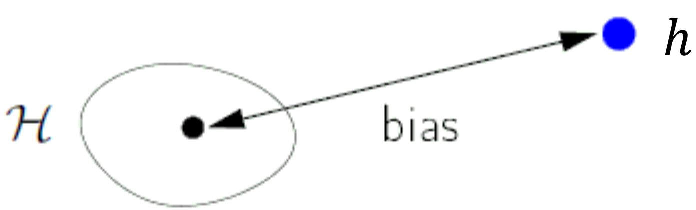

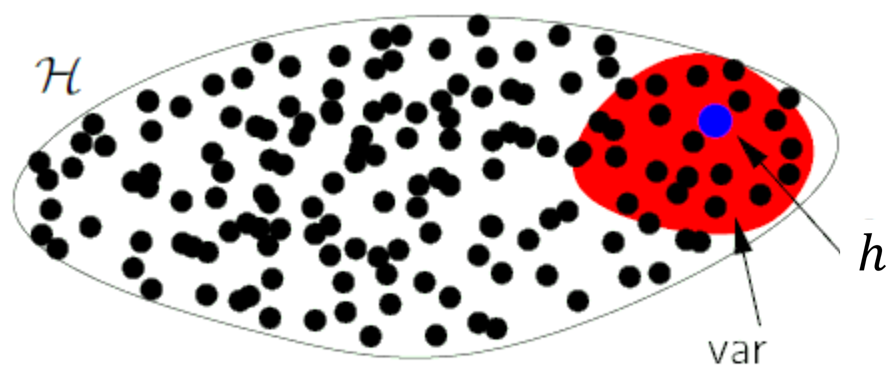

More complex $\mathcal{H}$ $\Rightarrow$ lower bias but higher variance

[Y.S. Abou Mostafa, 2012]

Example I: sin target

- Only two training example $N = 2$

- Two models used for learning:
  - $\mathcal{H}_0 : f(x) = b$
  - $\mathcal{H}_1 : f(x) = ax + b$

- Which is better $\mathcal{H}_0$ or $\mathcal{H}_1$?

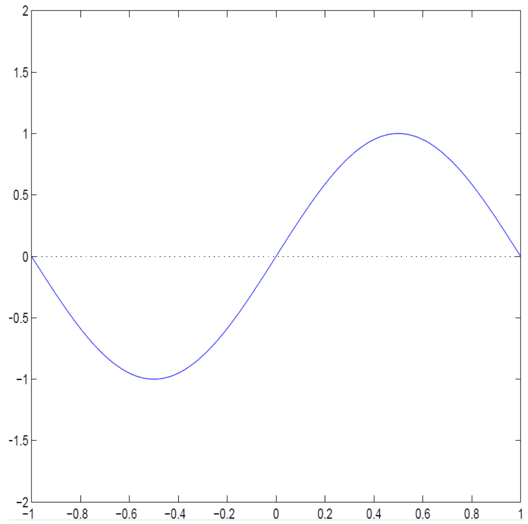

[Y.S. Abou Mostafa, 2012]

Example I: learning from a training set

$\mathcal{H}_0$

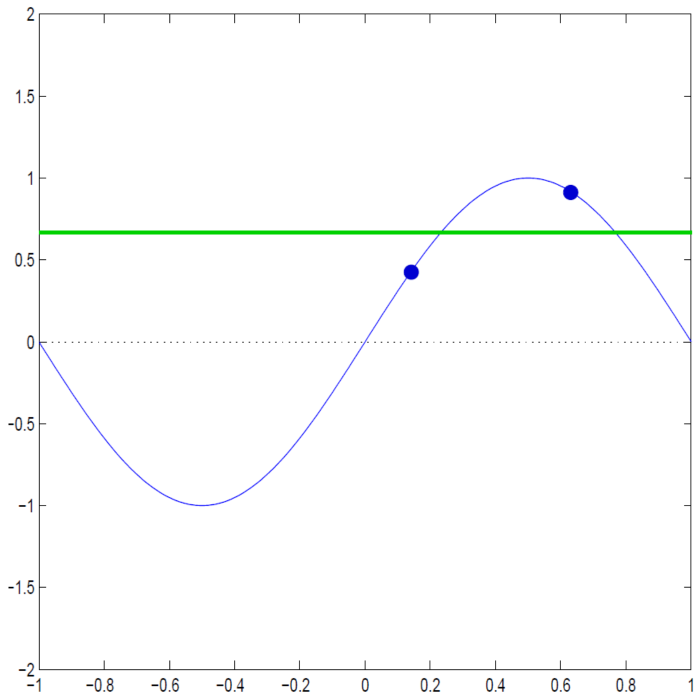

$\mathcal{H}_1$

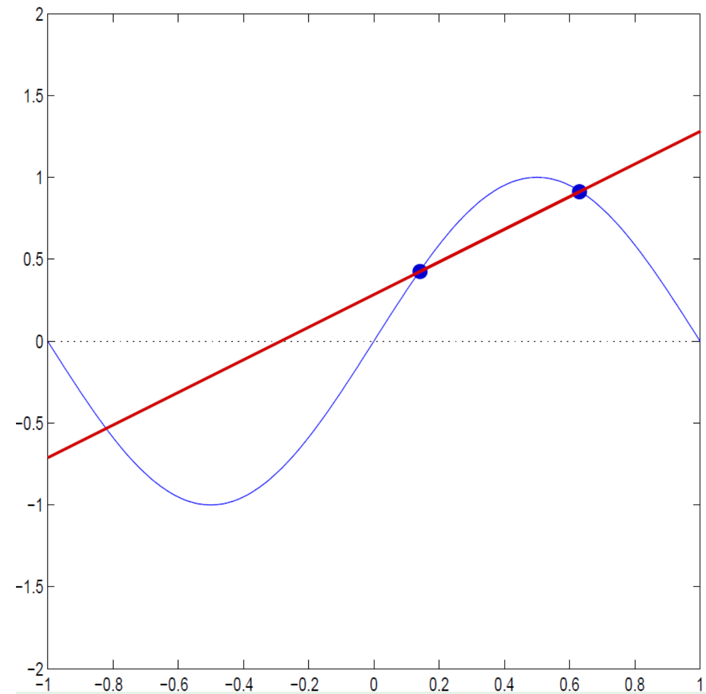

[Y.S. Abou Mostafa, 2012]

Example I: variance $\mathcal{H}_0$

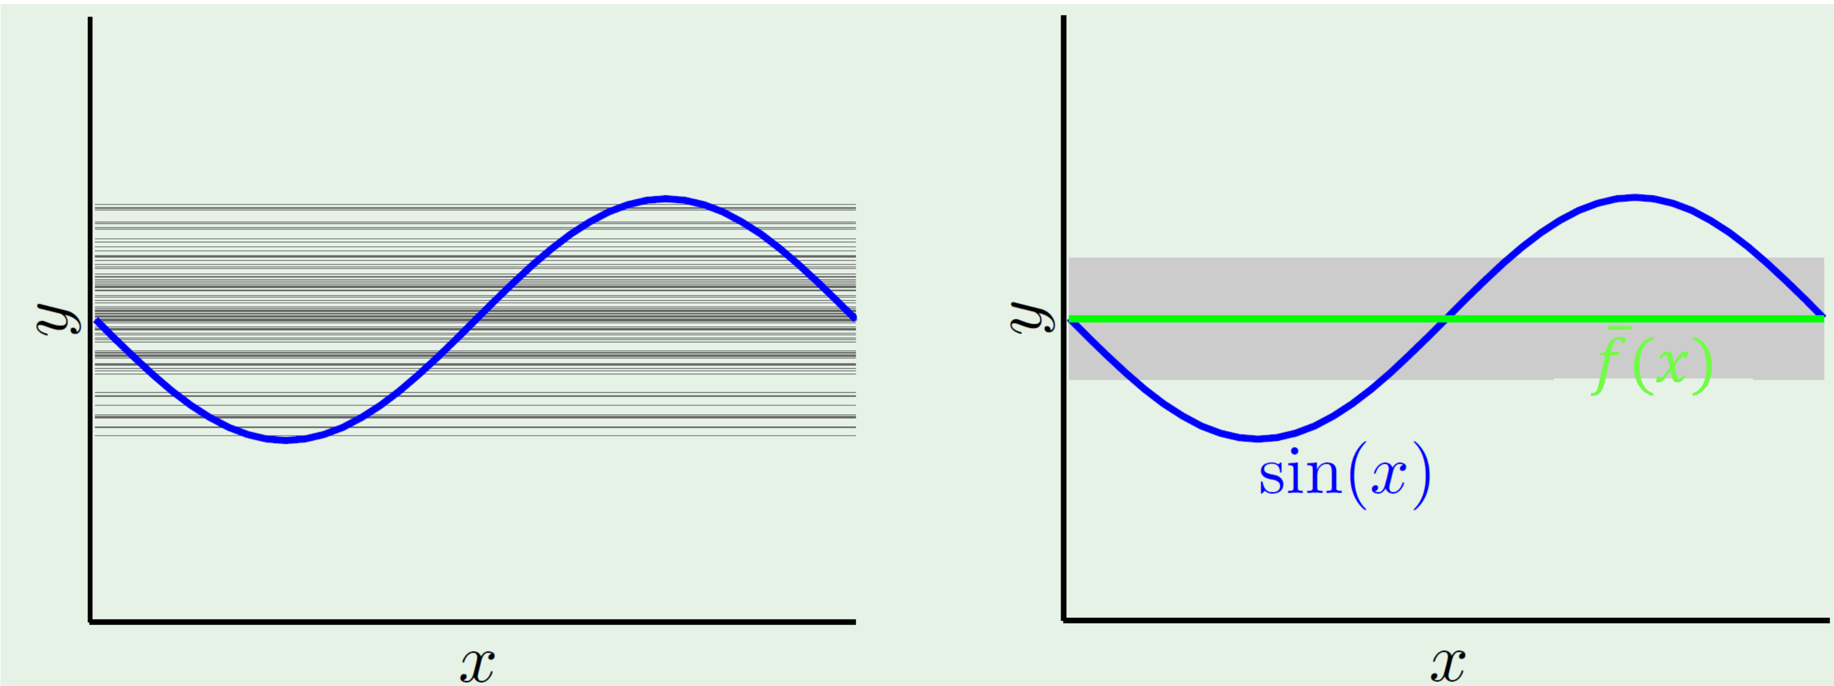

[Y.S. Abou Mostafa, et. al]

Example I: variance $\mathcal{H}_1$

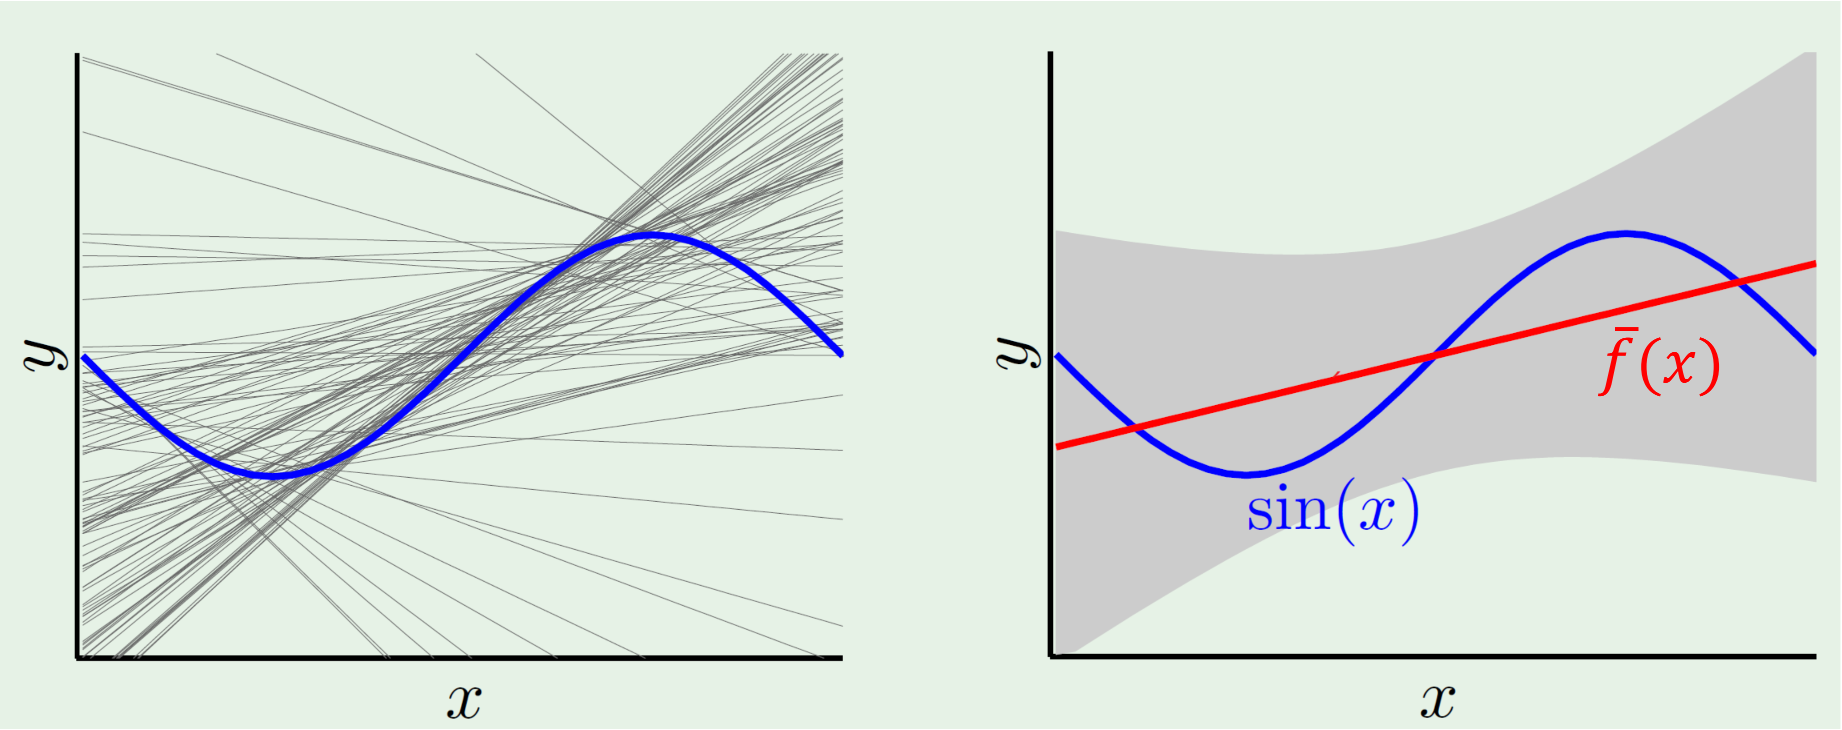

[Y.S. Abou Mostafa, et. al]

Example I: which is better?

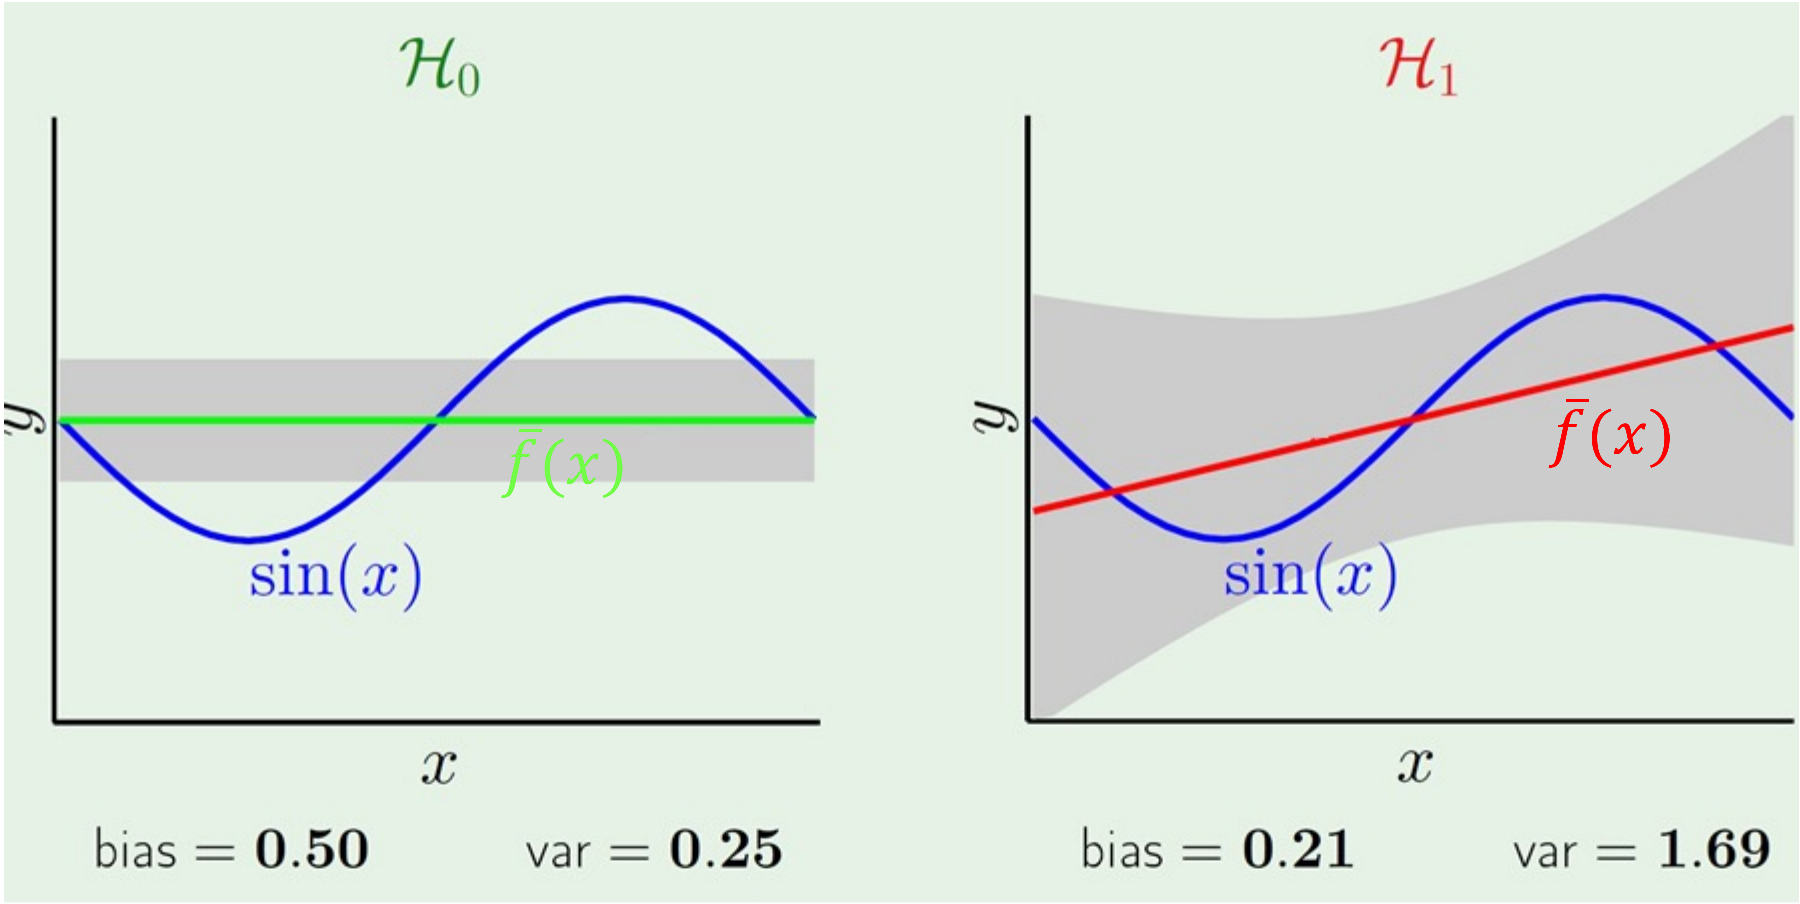

[Y.S. Abou Mostafa, 2012]

Lesson

Match the **model complexity**

to the data sources

not to the complexity of the target function.

Example I: regularization

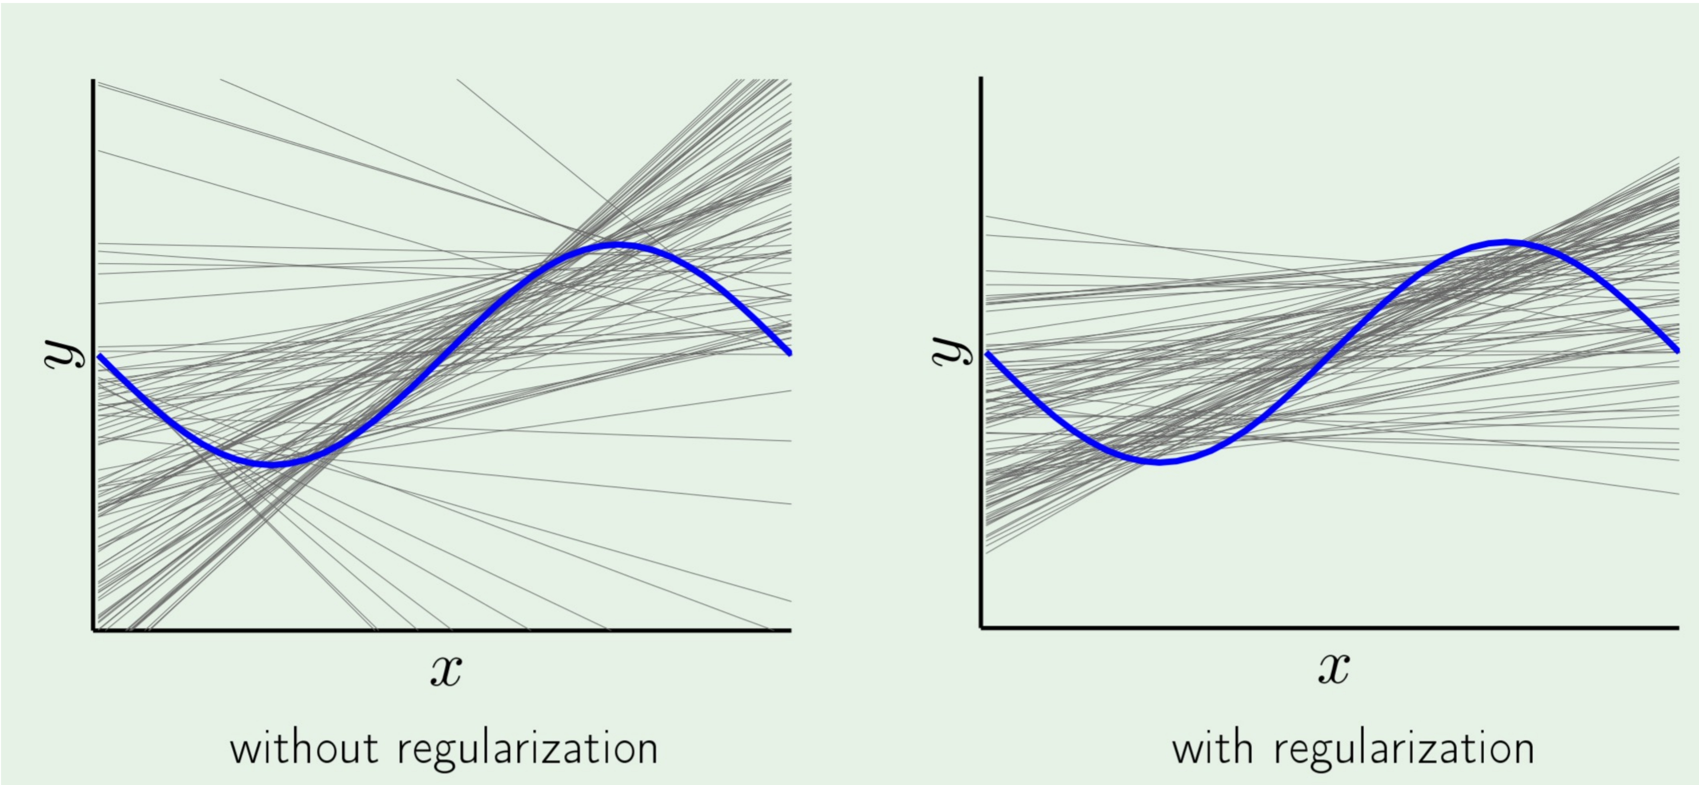

[Y.S. Abou Mostafa, 2012]

Example II: regularization & bias-variance

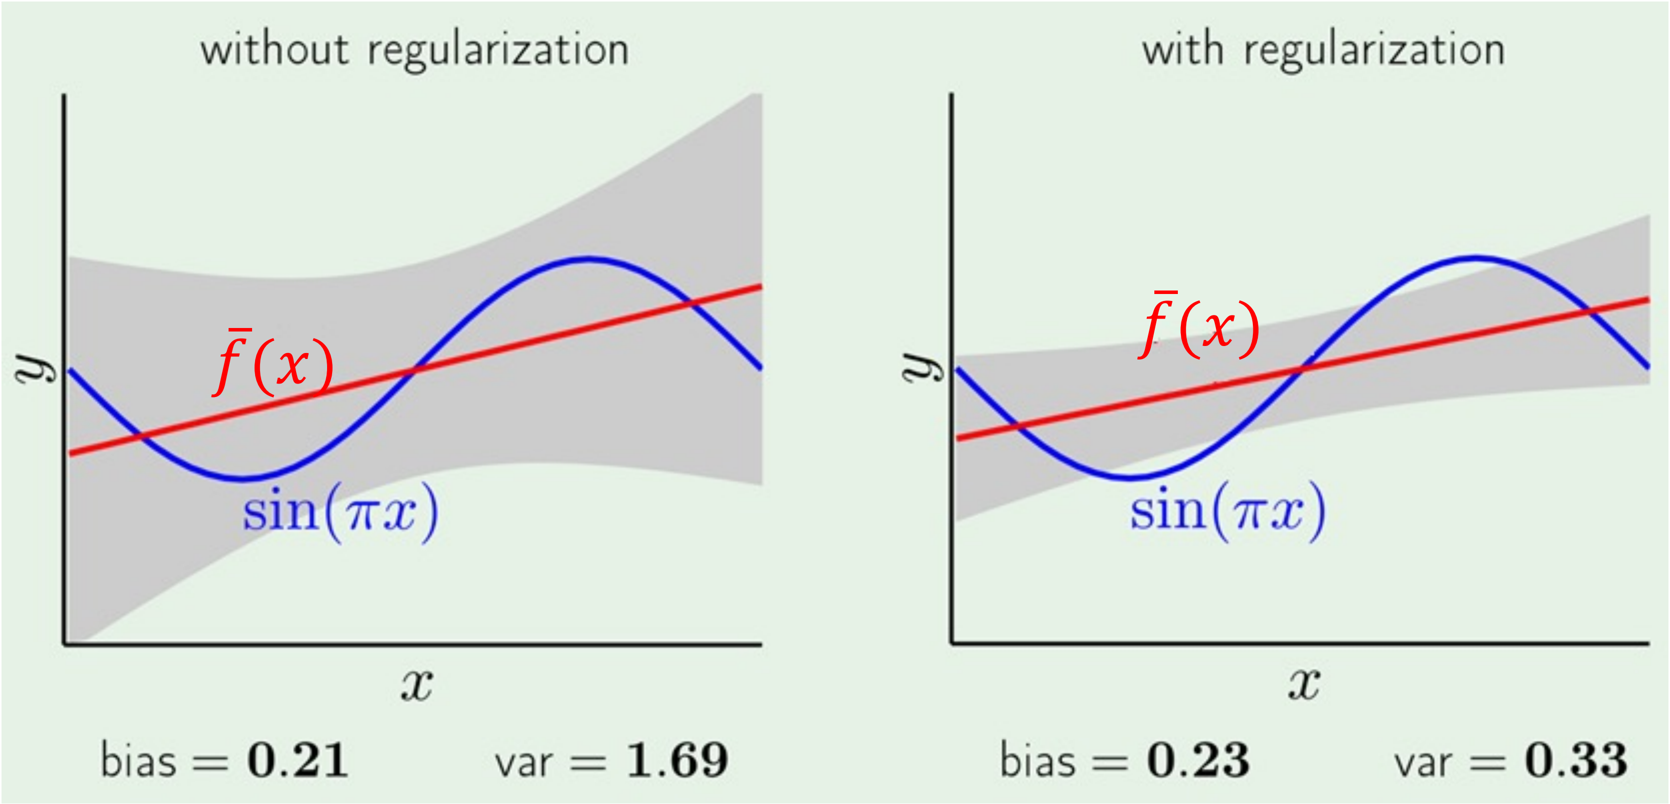

[Y.S. Abou Mostafa, 2012]

Winner of $\mathcal{H}_0$, $\mathcal{H}_1$, and $\mathcal{H}_1$ with regularization

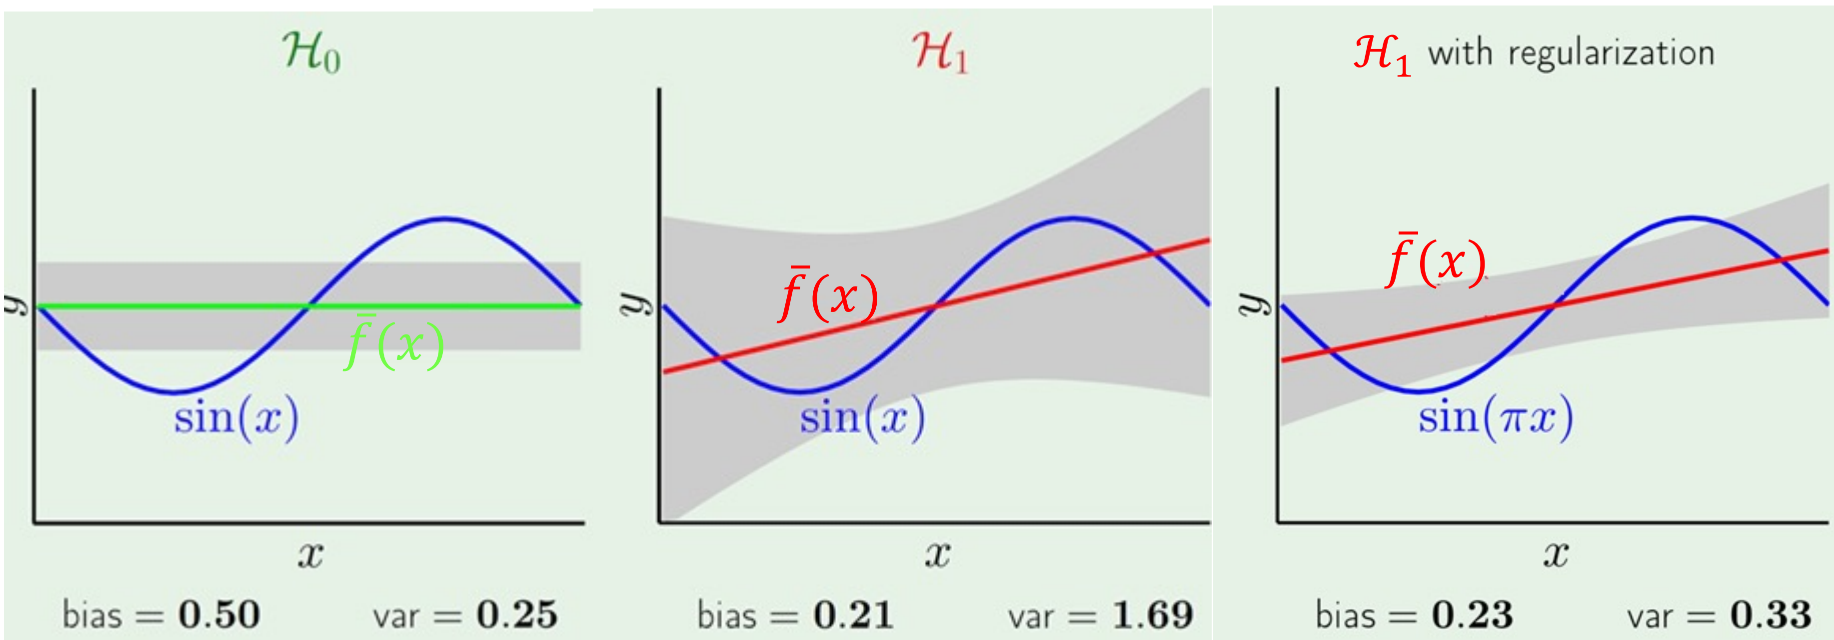

[Y.S. Abou Mostafa, 2012]

Example II: regularization & bias/variance

λ is large

λ is intermediate

λ is small

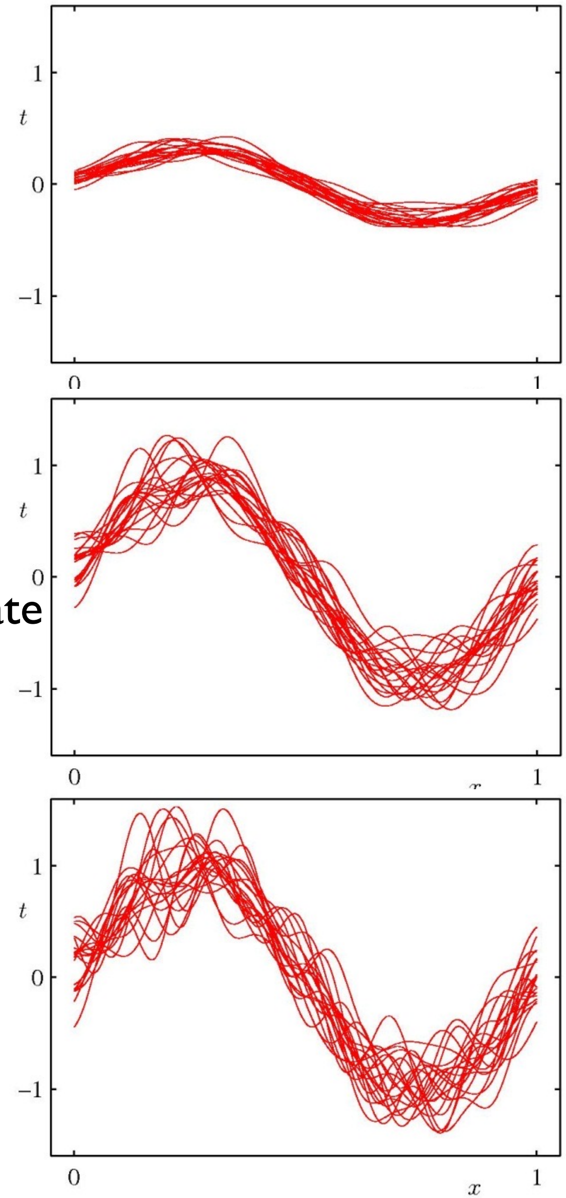

$L = 100$ data sets
$N = 25$
$m = 25$

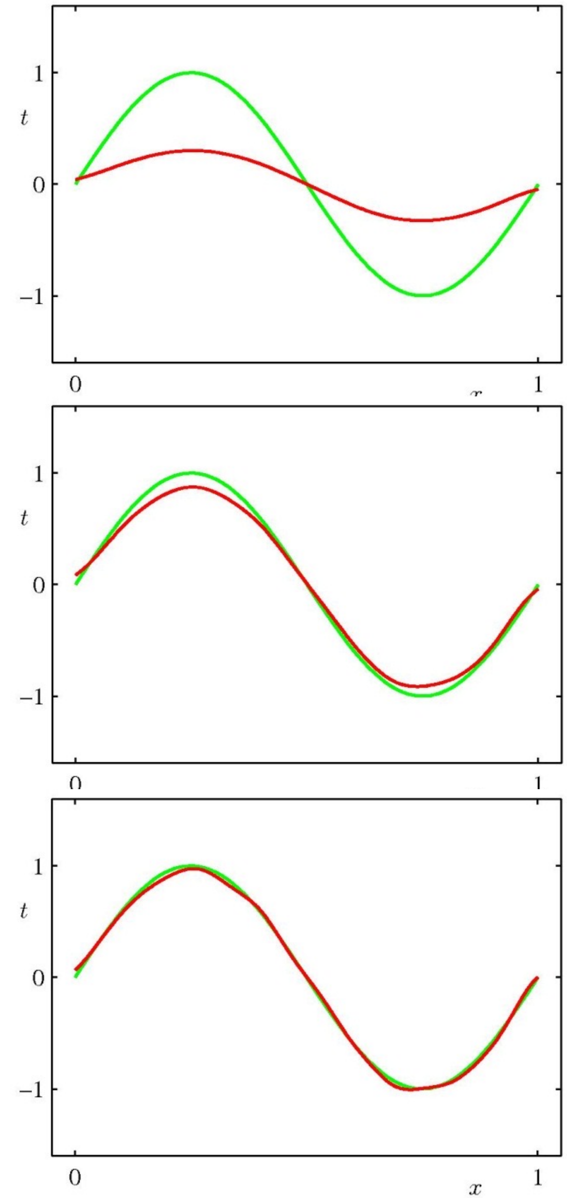

[Bishop]

Example II: Learning curves of bias, variance, and noise

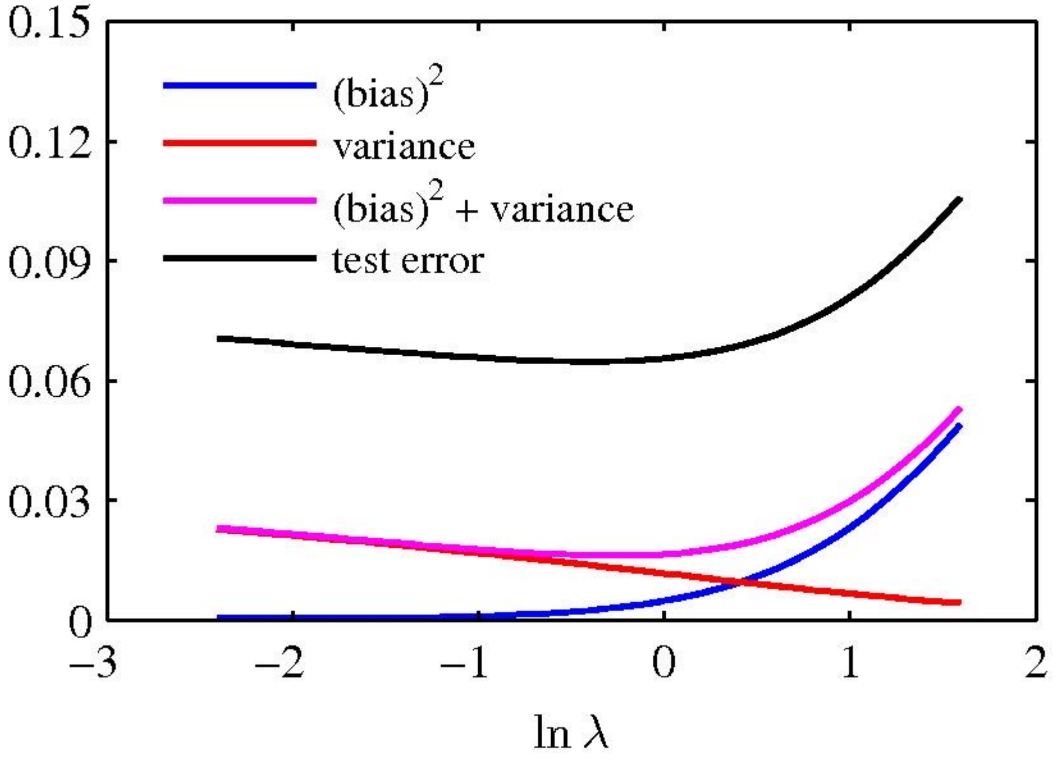

[Bishop]

Summary

- Generalized models
- Overfitting problem & how to avoid it
  - Evaluation and model selection
  - Regularization
- Bias-variance trade-off in regression problem

Leave-One-Out Cross Validation (LOOCV)

- When data is particularly scarce, cross-validation with $k = N$
  - Leave-one-out treats each training sample in turn as a test example and all other samples as the training set.

- Use for small datasets
  - When training data is valuable
  - LOOCV can be time expensive as N training steps are required.

Best unrestricted regression function

- If we know the joint distribution $P(\boldsymbol{x}, y)$ and no constraints on the regression function?
  - cost function: mean squared error

$$h^* = \underset{h:\mathbb{R}^d \to \mathbb{R}}{\operatorname{argmin}} \mathbb{E}_{\boldsymbol{x},y} \left[\left(y - h(\boldsymbol{x}) \right)^2 \right]$$

$$h^*(\boldsymbol{x}) = \mathbb{E}_{y|\boldsymbol{x}}[y]$$

Best unrestricted regression function: Proof

$$\mathbb{E}_{\boldsymbol{x},y}\left[\left(y - h(\boldsymbol{x})\right)^2\right] = \iint \left(y - h(\boldsymbol{x})\right)^2 p(\boldsymbol{x}, y)\,d\boldsymbol{x}\,dy$$

Best unrestricted regression function: Proof

$$\mathbb{E}_{\boldsymbol{x},y}\left[\left(y-h(\boldsymbol{x})\right)^2\right]=\iint \left(y-h(\boldsymbol{x})\right)^2 p(\boldsymbol{x},y)d\boldsymbol{x}dy$$

- For each $\boldsymbol{x}$, separately minimize loss since $h(\boldsymbol{x})$ can be chosen independently for each different $\boldsymbol{x}$:

$$\frac{\delta \mathbb{E}_{\boldsymbol{x},y}\left[\left(y-h(\boldsymbol{x})\right)^2\right]}{\delta h(\boldsymbol{x})}=-\int 2\left(y-h(\boldsymbol{x})\right)p(\boldsymbol{x},y)dy=0$$

Best unrestricted regression function: Proof

$$\mathbb{E}_{\boldsymbol{x},y}\left[\left(y - h(\boldsymbol{x})\right)^2\right] = \iint \left(y - h(\boldsymbol{x})\right)^2 p(\boldsymbol{x}, y) d\boldsymbol{x} dy$$

- For each $\boldsymbol{x}$, separately minimize loss since $h(\boldsymbol{x})$ can be chosen independently for each different $\boldsymbol{x}$:

$$\frac{\delta \mathbb{E}_{\boldsymbol{x},y}\left[\left(y - h(\boldsymbol{x})\right)^2\right]}{\delta h(\boldsymbol{x})} = -\int 2\left(y - h(\boldsymbol{x})\right) p(\boldsymbol{x}, y) dy = 0$$

$$\Rightarrow h(\boldsymbol{x}) = \frac{\int y p(\boldsymbol{x}, y) dy}{\int p(\boldsymbol{x}, y) dy} = \frac{\int y p(\boldsymbol{x}, y) dy}{p(\boldsymbol{x})} = \int y p(y|\boldsymbol{x}) dy = \mathbb{E}_{y|\boldsymbol{x}}[y]$$

Best unrestricted regression function: Proof

$$\mathbb{E}_{\boldsymbol{x},y}\left[\left(y - h(\boldsymbol{x})\right)^2\right] = \iint \left(y - h(\boldsymbol{x})\right)^2 p(\boldsymbol{x}, y) d\boldsymbol{x} dy$$

- For each $\boldsymbol{x}$, separately minimize loss since $h(\boldsymbol{x})$ can be chosen independently for each different $\boldsymbol{x}$:

$$\frac{\delta \mathbb{E}_{\boldsymbol{x},y}\left[\left(y - h(\boldsymbol{x})\right)^2\right]}{\delta h(\boldsymbol{x})} = -\int 2\left(y - h(\boldsymbol{x})\right) p(\boldsymbol{x}, y) dy = 0$$

$$\Rightarrow h(\boldsymbol{x}) = \frac{\int y p(\boldsymbol{x}, y) dy}{\int p(\boldsymbol{x}, y) dy} = \frac{\int y p(\boldsymbol{x}, y) dy}{p(\boldsymbol{x})} = \int y p(y|\boldsymbol{x}) dy = \mathbb{E}_{y|\boldsymbol{x}}[y]$$

$$\Longrightarrow h^*(\boldsymbol{x}) = \mathbb{E}_{y|\boldsymbol{x}}[y]$$

Error decomposition

$(\boldsymbol{x}, y) \sim P$

$h(\boldsymbol{x})$: minimizes the expected loss

$$E_{true}\big(f_{\mathcal{D}}(\boldsymbol{x})\big) = \mathbb{E}_{\boldsymbol{x},y}[(f_{\mathcal{D}}(\boldsymbol{x}) - y)^2]$$ Expected loss

Error decomposition

$(\boldsymbol{x}, y) \sim P$

$h(\boldsymbol{x})$: minimizes the expected loss

$$E_{true}\big(f_{\mathcal{D}}(\boldsymbol{x}\,)\big) = \mathbb{E}_{\boldsymbol{x},y}[(f_{\mathcal{D}}(\boldsymbol{x}\,) - y)^2]$$

Expected loss

$$= \mathbb{E}_{\boldsymbol{x},y}\,[(f_{\mathcal{D}}(\boldsymbol{x}\,) - h(\boldsymbol{x}) + h(\boldsymbol{x}) - y)^2]$$

Error decomposition

$(\boldsymbol{x}, y) \sim P$

$h(\boldsymbol{x})$: minimizes the expected loss

$$E_{true}\left(f_{\mathcal{D}}(\boldsymbol{x})\right) = \mathbb{E}_{\boldsymbol{x},y}[(f_{\mathcal{D}}(\boldsymbol{x}) - y)^2]$$

Expected loss

$$= \mathbb{E}_{\boldsymbol{x},y}\,[(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x}) + h(\boldsymbol{x}) - y)^2]$$

$$= \mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right] + \mathbb{E}_{\boldsymbol{x},y}[(h(\boldsymbol{x}) - \boldsymbol{y})^2]$$
$$+2\mathbb{E}_{\boldsymbol{x},y}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)(h(\boldsymbol{x}) - y)\right]$$

Error decomposition

$(\boldsymbol{x}, y) \sim P$

$h(\boldsymbol{x})$: minimizes the expected loss

$$E_{true}\big(f_{\mathcal{D}}(\boldsymbol{x})\big) = \mathbb{E}_{\boldsymbol{x},y}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - y\right)^2\right]$$

Expected loss

$$= \mathbb{E}_{\boldsymbol{x},y}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x}) + h(\boldsymbol{x}) - y\right)^2\right]$$

$$= \mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right] + \mathbb{E}_{\boldsymbol{x},y}\left[\left(h(\boldsymbol{x}) - \boldsymbol{y}\right)^2\right]$$
$$+2\underbrace{\mathbb{E}_{\boldsymbol{x},y}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)\left(h(\boldsymbol{x}) - y\right)\right]}_{\mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)\mathbb{E}_{y|\boldsymbol{x}}\left[\left(h(\boldsymbol{x}) - y\right)\right]\right]}$$

Error decomposition

$(\boldsymbol{x}, y) \sim P$

$h(\boldsymbol{x})$: minimizes the expected loss

$$E_{true}\left(f_{\mathcal{D}}(\boldsymbol{x})\right) = \mathbb{E}_{\boldsymbol{x},y}[(f_{\mathcal{D}}(\boldsymbol{x}) - y)^2]$$

Expected loss

$$= \mathbb{E}_{\boldsymbol{x},y}\left[(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x}) + h(\boldsymbol{x}) - y)^2\right]$$

$$= \mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right] + \mathbb{E}_{\boldsymbol{x},y}[(h(\boldsymbol{x}) - \boldsymbol{y})^2]$$
$$+2\underbrace{\mathbb{E}_{\boldsymbol{x},y}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)(h(\boldsymbol{x}) - y)\right]}_{\mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)\underbrace{\mathbb{E}_{y|\boldsymbol{x}}[(h(\boldsymbol{x}) - y)]}_{0}\right]}$$

Error decomposition

$(\boldsymbol{x}, y) \sim P$

$h(\boldsymbol{x})$: minimizes the expected loss

$$E_{true}\big(f_{\mathcal{D}}(\boldsymbol{x})\big) = \mathbb{E}_{\boldsymbol{x},y}[(f_{\mathcal{D}}(\boldsymbol{x}) - y)^2]$$

$$= \mathbb{E}_{\boldsymbol{x},y}\,[(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x}) + h(\boldsymbol{x}) - y)^2]$$

$$= \mathbb{E}_{\boldsymbol{x}}\left[\big(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\big)^2\right] + \underbrace{\mathbb{E}_{\boldsymbol{x},y}[(h(\boldsymbol{x}) - \boldsymbol{y})^2]}_{\text{noise}}$$

$$+\,0$$

- Noise shows the irreducible minimum value of the loss function

Expectation of true error

$$E_{true}\big(f_{\mathcal{D}}(\boldsymbol{x})\big) = \mathbb{E}_{\boldsymbol{x},y}[(f_{\mathcal{D}}(\boldsymbol{x}) - y)^2]$$

$$= \mathbb{E}_{\boldsymbol{x}}\left[\big(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\big)^2\right] + noise$$

Expectation of true error

$$E_{true}\left(f_{\mathcal{D}}(\boldsymbol{x})\right) = \mathbb{E}_{\boldsymbol{x},y}[(f_{\mathcal{D}}(\boldsymbol{x}) - y)^2]$$
$$= \mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right] + noise$$

$$\mathbb{E}_{\mathcal{D}}\left[\mathbb{E}_{\boldsymbol{x}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]\right]$$
$$= \mathbb{E}_{\boldsymbol{x}}\left[\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]\right]$$

We now want to focus on $\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$.

The average hypothesis

$$\bar{f}(\boldsymbol{x}) \equiv E_{\mathcal{D}}[f_{\mathcal{D}}(\boldsymbol{x})]$$

$$\bar{f}(\boldsymbol{x}) \approx \frac{1}{K}\sum_{k=1}^{K} f_{\mathcal{D}^{(k)}}(\boldsymbol{x})$$

K training sets (of size N) sampled from $P(\boldsymbol{x}, y)$: $\mathcal{D}^{(1)}, \mathcal{D}^{(2)}, \dots, \mathcal{D}^{(K)}$

Using the average hypothesis

$$\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x}) + \bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

Using the average hypothesis

$$\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x}) + \bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2 + \left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2 + 2\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)\left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)\right]$$

Using the average hypothesis

$$\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x}) + \bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2 + \left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2 + 2\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)\left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)\right]$$

$$= \mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2\right] + \left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2$$

Bias and variance

$$\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right] = \underbrace{\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2\right]}_{\mathrm{var}(\boldsymbol{x})} + \underbrace{\left(\bar{f}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2}_{\mathrm{bias}(\boldsymbol{x})}$$

$$\mathbb{E}_{\boldsymbol{x}}\left[\mathbb{E}_{\mathcal{D}}\left[\left(f_{\mathcal{D}}(\boldsymbol{x}) - h(\boldsymbol{x})\right)^2\right]\right] = \mathbb{E}_{\boldsymbol{x}}[\mathrm{var}(\boldsymbol{x}) + \mathrm{bias}(\boldsymbol{x})]$$

$$= \mathrm{var} + \mathrm{bias}$$