Statistical Machine Learning (BE4M33SSU)
Lecture 2: Empirical Risk

Czech Technical University in Prague
V. Franc

Definition of the prediction problem

- $\mathcal{X}$ is a set of input observations/features
- $\mathcal{Y}$ is a set of hidden states
- $(x, y) \in \mathcal{X} \times \mathcal{Y}$ samples randomly drawn from r.v. with p.d.f. $p(x, y)$
- $h\colon \mathcal{X} \to \mathcal{Y}$ is a prediction strategy/hypothesis
- $\ell\colon \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ is a loss function
- Task: find a strategy with the minimal true risk (expected loss)

$$R(h) = \int \sum_{y \in \mathcal{Y}} \ell(y, h(x))\, p(x, y)\, \mathrm{d}x = \mathbb{E}_{(x,y)\sim p}\Big(\ell(y, h(x))\Big)$$

- Bayes predictor h^* attains the minimal risk $R(h^*) = \inf\limits_{h \in \mathcal{Y}^{\mathcal{X}}} R(h)$

Example of a prediction problem

- The statistical model is known:
  - $\mathcal{X} = \mathbb{R}, \quad \mathcal{Y} = \{+1, -1\}, \quad \ell(y, y') = \begin{cases} 0 & \text{if} \quad y = y' \\ 1 & \text{if} \quad y \neq y' \end{cases}$
  - $p(x, y) = p(y) \, \frac{1}{\sqrt{2\pi}\sigma} \, e^{-\frac{1}{2\sigma^2}(x - \mu_y)^2}, \quad y \in \mathcal{Y}.$

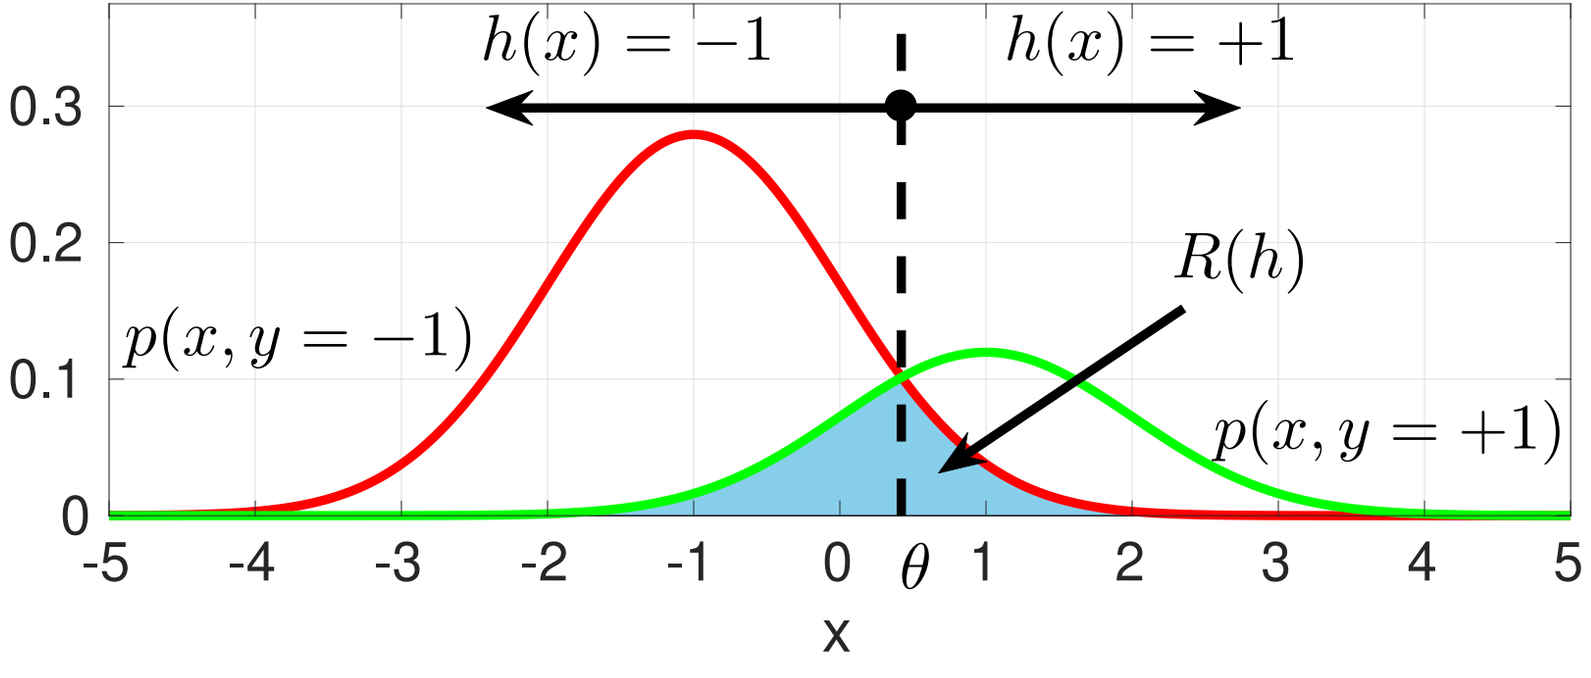

Machine Learning: solving the prediction problem based on examples

- Assumption: we have an access to examples

$$\{(x^1, y^1), (x^2, y^2), \ldots\}$$

drawn from i.i.d. r.v. distributed according to unknown $p(x, y)$.

- 1) Evaluation: estimate true risk $R(h)$ of given $h\colon \mathcal{X} \to \mathcal{Y}$ using test set

$$\mathcal{S}^l = \{(x^i, y^i) \in (\mathcal{X} \times \mathcal{Y}) \mid i = 1, \ldots, l\}$$

drawn i.i.d. from $p(x, y)$.

- 2) Learning: find $h\colon \mathcal{X} \to \mathcal{Y}$ with small $R(h)$ using training set

$$\mathcal{T}^m = \{(x^i, y^i) \in (\mathcal{X} \times \mathcal{Y}) \mid i = 1, \ldots, m\}$$

drawn i.i.d. from $p(x, y)$.

Evaluation: estimation of the expected risk

- Given a predictor $h\colon \mathcal{X} \to \mathcal{Y}$ and a test set $\mathcal{S}^l$ draw i.i.d. from distribution $p(x, y)$, compute the empirical risk

$$R_{\mathcal{S}^l}(h) = \frac{1}{l}\big(\ell(y^1, h(x^1)) + \cdots + \ell(y^l, h(x^l)\big) = \frac{1}{l}\sum_{i=1}^{l} \ell(y^i, h(x^i))$$

  and use it as an estimate of $R(h) = \mathbb{E}_{(x,y)\sim p}(\ell(y, h(x)))$.

- $R_{\mathcal{S}^l}(h)$ is a random number with the variance depending on l.

- We construct a confidence interval such that

$$R(h) \in (R_{\mathcal{S}^{l(h)}} - \varepsilon, R_{\mathcal{S}^{l(h)}} + \varepsilon) \quad \text{with probability (confidence) } \gamma \in (0, 1)$$

where ε is a deviation.

Law of large numbers

- Sample mean (arithmetic average) of the results of random trials gets closer to the expected value as more trials are performed.
- Example: The expected value of a single roll of a fair die is

$$\mu = \mathbb{E}_{z\sim p}(z) = \sum_{z=1}^{6} z\, p(z) = \frac{1+2+3+4+5+6}{6} = 3.5$$

$$\hat{\mu} = \frac{1}{l}\sum_{i=1}^{l} z^i$$

$z^1 = 3$ $z^2 = 1$ $z^3 = 5$ $\quad\cdots\quad$ $z^l = 2$

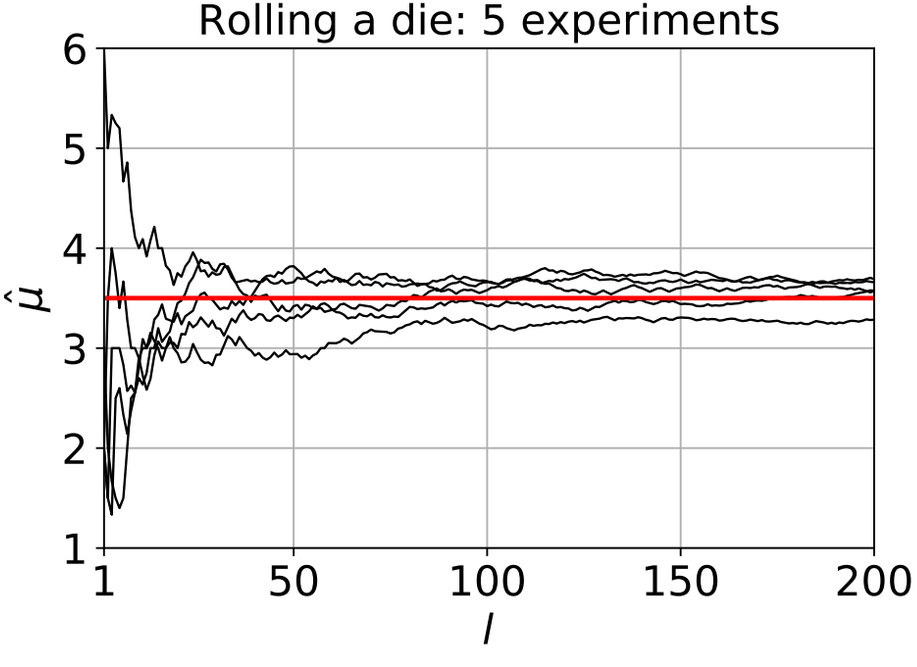

Counting frequency of bad estimates

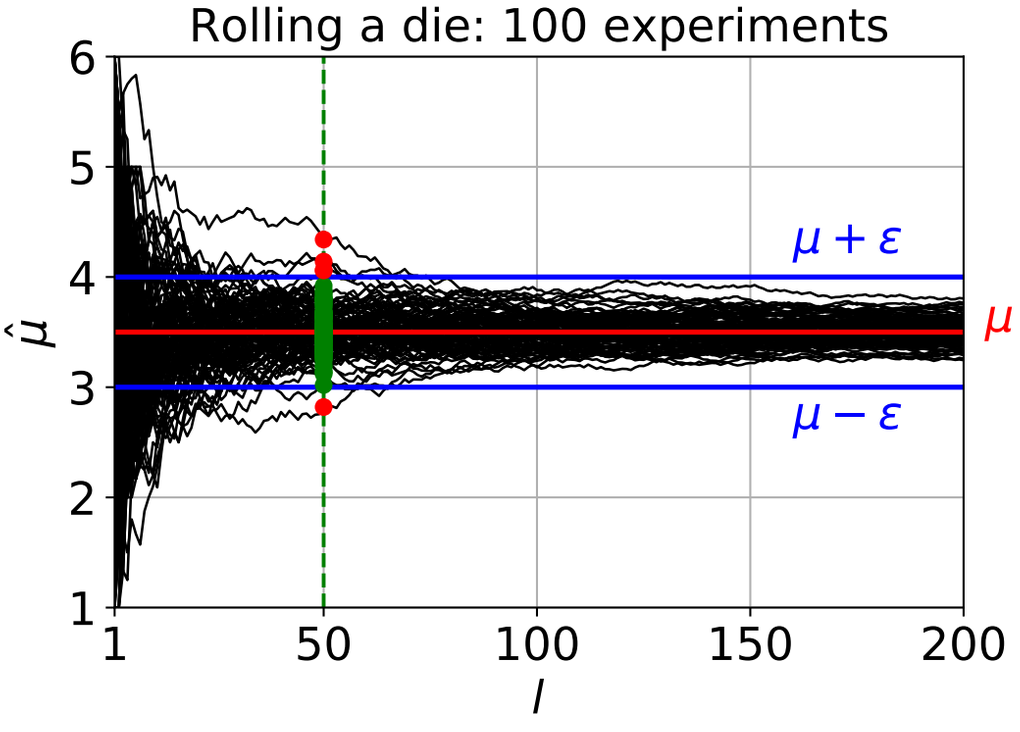

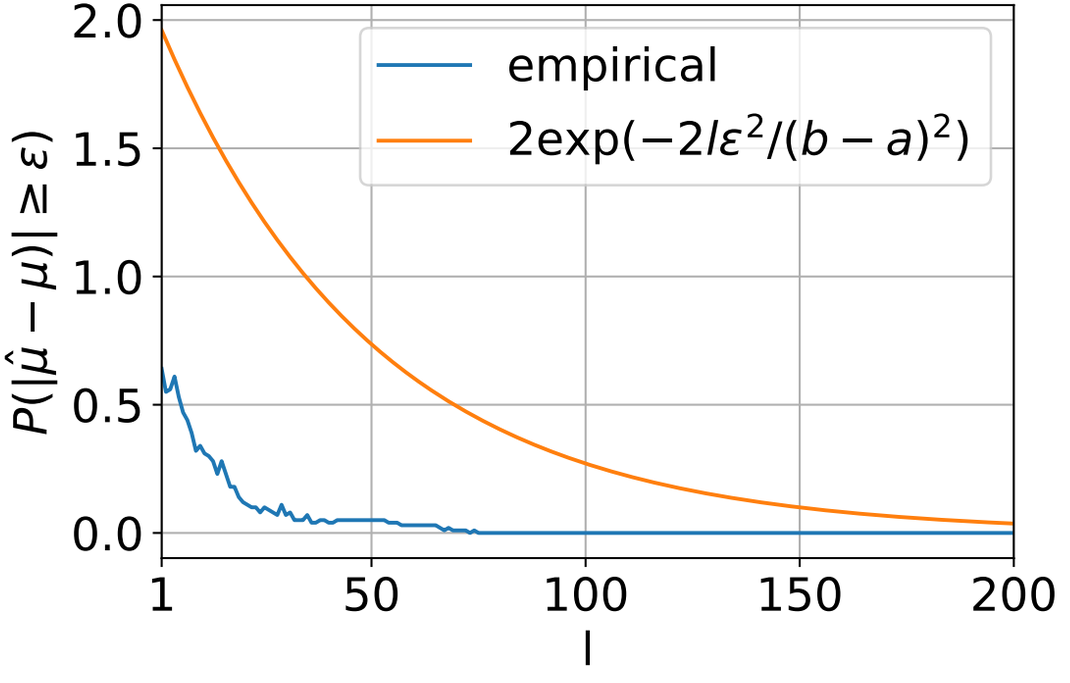

sample size $l = 50$, deviation $\varepsilon = 0.5$

Hoeffding inequality

$$\frac{\#\left(|\hat{\mu} - \mu| \geq \varepsilon\right)}{\#\text{experiments}} = \frac{5}{100} = 0.05 \quad \rightarrow \quad \mathbb{P}\Big(|\hat{\mu} - \mu| \geq \varepsilon\Big) \leq 2e^{-\frac{2l\,\varepsilon^2}{(b-a)^2}}$$

$$a = 1,\ b = 6$$

Hoeffding inequality

Theorem: Let $\{z^1, \ldots, z^l\}$ be a sample from i.i.d. r.v. from $[a, b]$ with expected value μ. Let $\hat{\mu} = \frac{1}{l} \sum_{i=1}^{l} z^i$. Then for any $\varepsilon > 0$ it holds that

$$\mathbb{P}\Big(|\hat{\mu} - \mu| \geq \varepsilon \Big) \leq 2e^{-\frac{2l\,\varepsilon^2}{(b-a)^2}}$$

Properties:

- Conservative: the bound may not be tight.
- General: the bound holds for any distribution.
- Cheap: The bound is simple and easy to compute.

Confidence intervals

$$(l, \gamma) \to \varepsilon$$

- Let $\hat{\mu} = \frac{1}{l}\sum_{i=1}^{l} z^i$ be the sample mean computed from $\{z^1, \ldots, z^l\} \in [a, b]^l$ sampled from r.v. with expected value μ.

- Find ε such that $\mu \in (\hat{\mu} - \varepsilon, \hat{\mu} + \varepsilon)$ with probability at least γ.

  Using the Hoeffding inequality we can write

$$\mathbb{P}\Big(|\hat{\mu} - \mu| < \varepsilon\Big) = 1 - \mathbb{P}\Big(|\hat{\mu} - \mu| \geq \varepsilon\Big) \geq 1 - 2e^{-\frac{2\,l\,\varepsilon^2}{(b-a)^2}} = \gamma$$

  and solving the last equation for ε yields

$$\varepsilon = |b - a|\sqrt{\frac{\log(2) - \log(1 - \gamma)}{2\,l}}$$

Confidence intervals

$$(\varepsilon, \gamma) \rightarrow l$$

- Let $\hat{\mu} = \frac{1}{l} \sum_{i=1}^{l} z^i$ be the sample mean computed from $\{z^1, \ldots, z^l\} \in [a, b]^l$ sampled from r.v. with expected value μ.

- Given a fixed $\varepsilon > 0$ and $\gamma \in (0, 1)$, what is the minimal number of examples l such that $\mu \in (\hat{\mu} - \varepsilon, \hat{\mu} + \varepsilon)$ with probability γ at least ?

  Starting from

$$\mathbb{P}\Big(|\hat{\mu} - \mu| < \varepsilon \Big) = 1 - \mathbb{P}\Big(|\hat{\mu} - \mu| \geq \varepsilon \Big) \geq 1 - 2e^{-\frac{2\, l\, \varepsilon^2}{(b-a)^2}} = \gamma$$

  and solving for l yields

$$l = \frac{\log(2) - \log(1 - \gamma)}{2\, \varepsilon^2} (b - a)^2$$

Evaluation: estimation of the true risk

- Given $h\colon \mathcal{X} \to \mathcal{Y}$ estimate the true risk $R(h) = \mathbb{E}_{(x,y)\sim p}(\ell(y, h(x)))$ by the empirical risk $R_{\mathcal{S}^l}(h) = \frac{1}{l}\sum_{i=1}^{l} \ell(y^i, h(x^i))$ using the test set $\mathcal{S}^l$.
- The incurred losses $z^i = \ell(y^i, h(x^i)) \in [\ell_{\min}, \ell_{\max}]$, $i \in \{1, \ldots, l\}$, are realizations of i.i.d. r.v. with the expected value $\mu = R(h)$.
- According to the Hoeffding inequality, for any $\varepsilon > 0$ the probability of seeing a "bad test set" can be bound by

$$\mathbb{P}\left(\left|R_{\mathcal{S}^l}(h) - R(h)\right| \geq \varepsilon\right) \leq 2e^{-\frac{2l\,\varepsilon^2}{(\ell_{\min}-\ell_{\max})^2}}$$

- Remark: For any $p(x, y)$ and $\ell\colon \mathcal{Y} \times \mathcal{Y} \to [\ell_{\min}, \ell_{\max}]$, the empricial risk $R_{\mathcal{S}^l}(h)$ convergences in probability to the true risk $R(h)$:

$$\forall\, \varepsilon > 0\colon \lim_{l\to\infty} \mathbb{P}\left(\left|R_{\mathcal{S}^l}(h) - R(h)\right| \geq \varepsilon\right) = 0$$

Evaluation: recipe for constructing confidence intervals

- Given $h\colon \mathcal{X} \to \mathcal{Y}$ estimate the true risk $R(h) = \mathbb{E}_{(x,y)\sim p}(\ell(y, h(x)))$ by the empirical risk $R_{\mathcal{S}^l}(h) = \frac{1}{l}\sum_{i=1}^{l} \ell(y^i, h(x^i))$ using the test set $\mathcal{S}^l$.
- Confidence interval:

$$R(h) \in \left(R_{\mathcal{S}^l}(h) - \varepsilon, R_{\mathcal{S}^l}(h) + \varepsilon\right) \quad \text{with probability} \quad \gamma \in (0,1)$$

- For fixed l and $\gamma \in (0,1)$ compute interval width

$$\varepsilon = (\ell_{\max} - \ell_{\min})\sqrt{\frac{\log(2) - \log(1-\gamma)}{2\,l}}\,.$$

- For fixed ε and $\gamma \in (0,1)$ compute number of test examples

$$l = \frac{\log(2) - \log(1-\gamma)}{2\,\varepsilon^2}\,(\ell_{\max} - \ell_{\min})^2$$

Example: confidence intervals for classification error

- The width of $R(h) \in \left(R_{\mathcal{S}^l}(h) - \varepsilon, R_{\mathcal{S}^l}(h) + \varepsilon\right)$ is for $\ell(y, y') = [\![y \neq y']\!]$ given by $\varepsilon = \sqrt{\frac{\log(2) - \log(1-\gamma)}{2\,l}}$

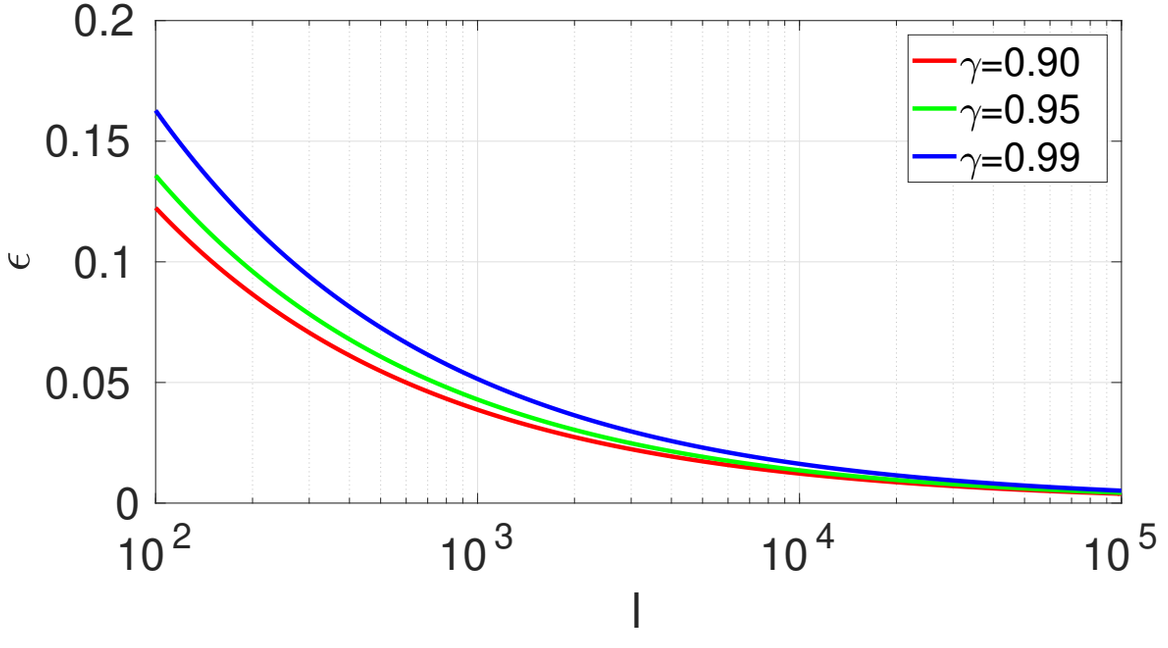

for $\gamma = 0.95$

l	100	1,000	10,000	18,445
ε	0.135	0.043	0.014	0.01

- Example: $l = 10,000$, $R_{\mathcal{S}^l}(h) = 0.162$, then classification error is $16.2 \pm 1.4\%$ with confidence 95%.