University of Amsterdam
Informatics Institute

AMLAB
Amsterdam
Machine Learning Lab

Machine Learning 1

Lecture 1.2 - What is Machine Learning?

Erik Bekkers

(Bishop 1.0 and 1.1)

Slide credits: Patrick Forré and Rianne van den Berg

Image credit: Kirillm | Getty Images

What is machine learning?

“A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E.”

- Tom M. Mitchell

Machine Learning, Tom Mitchell, McGraw Hill, 1997

E: Experience

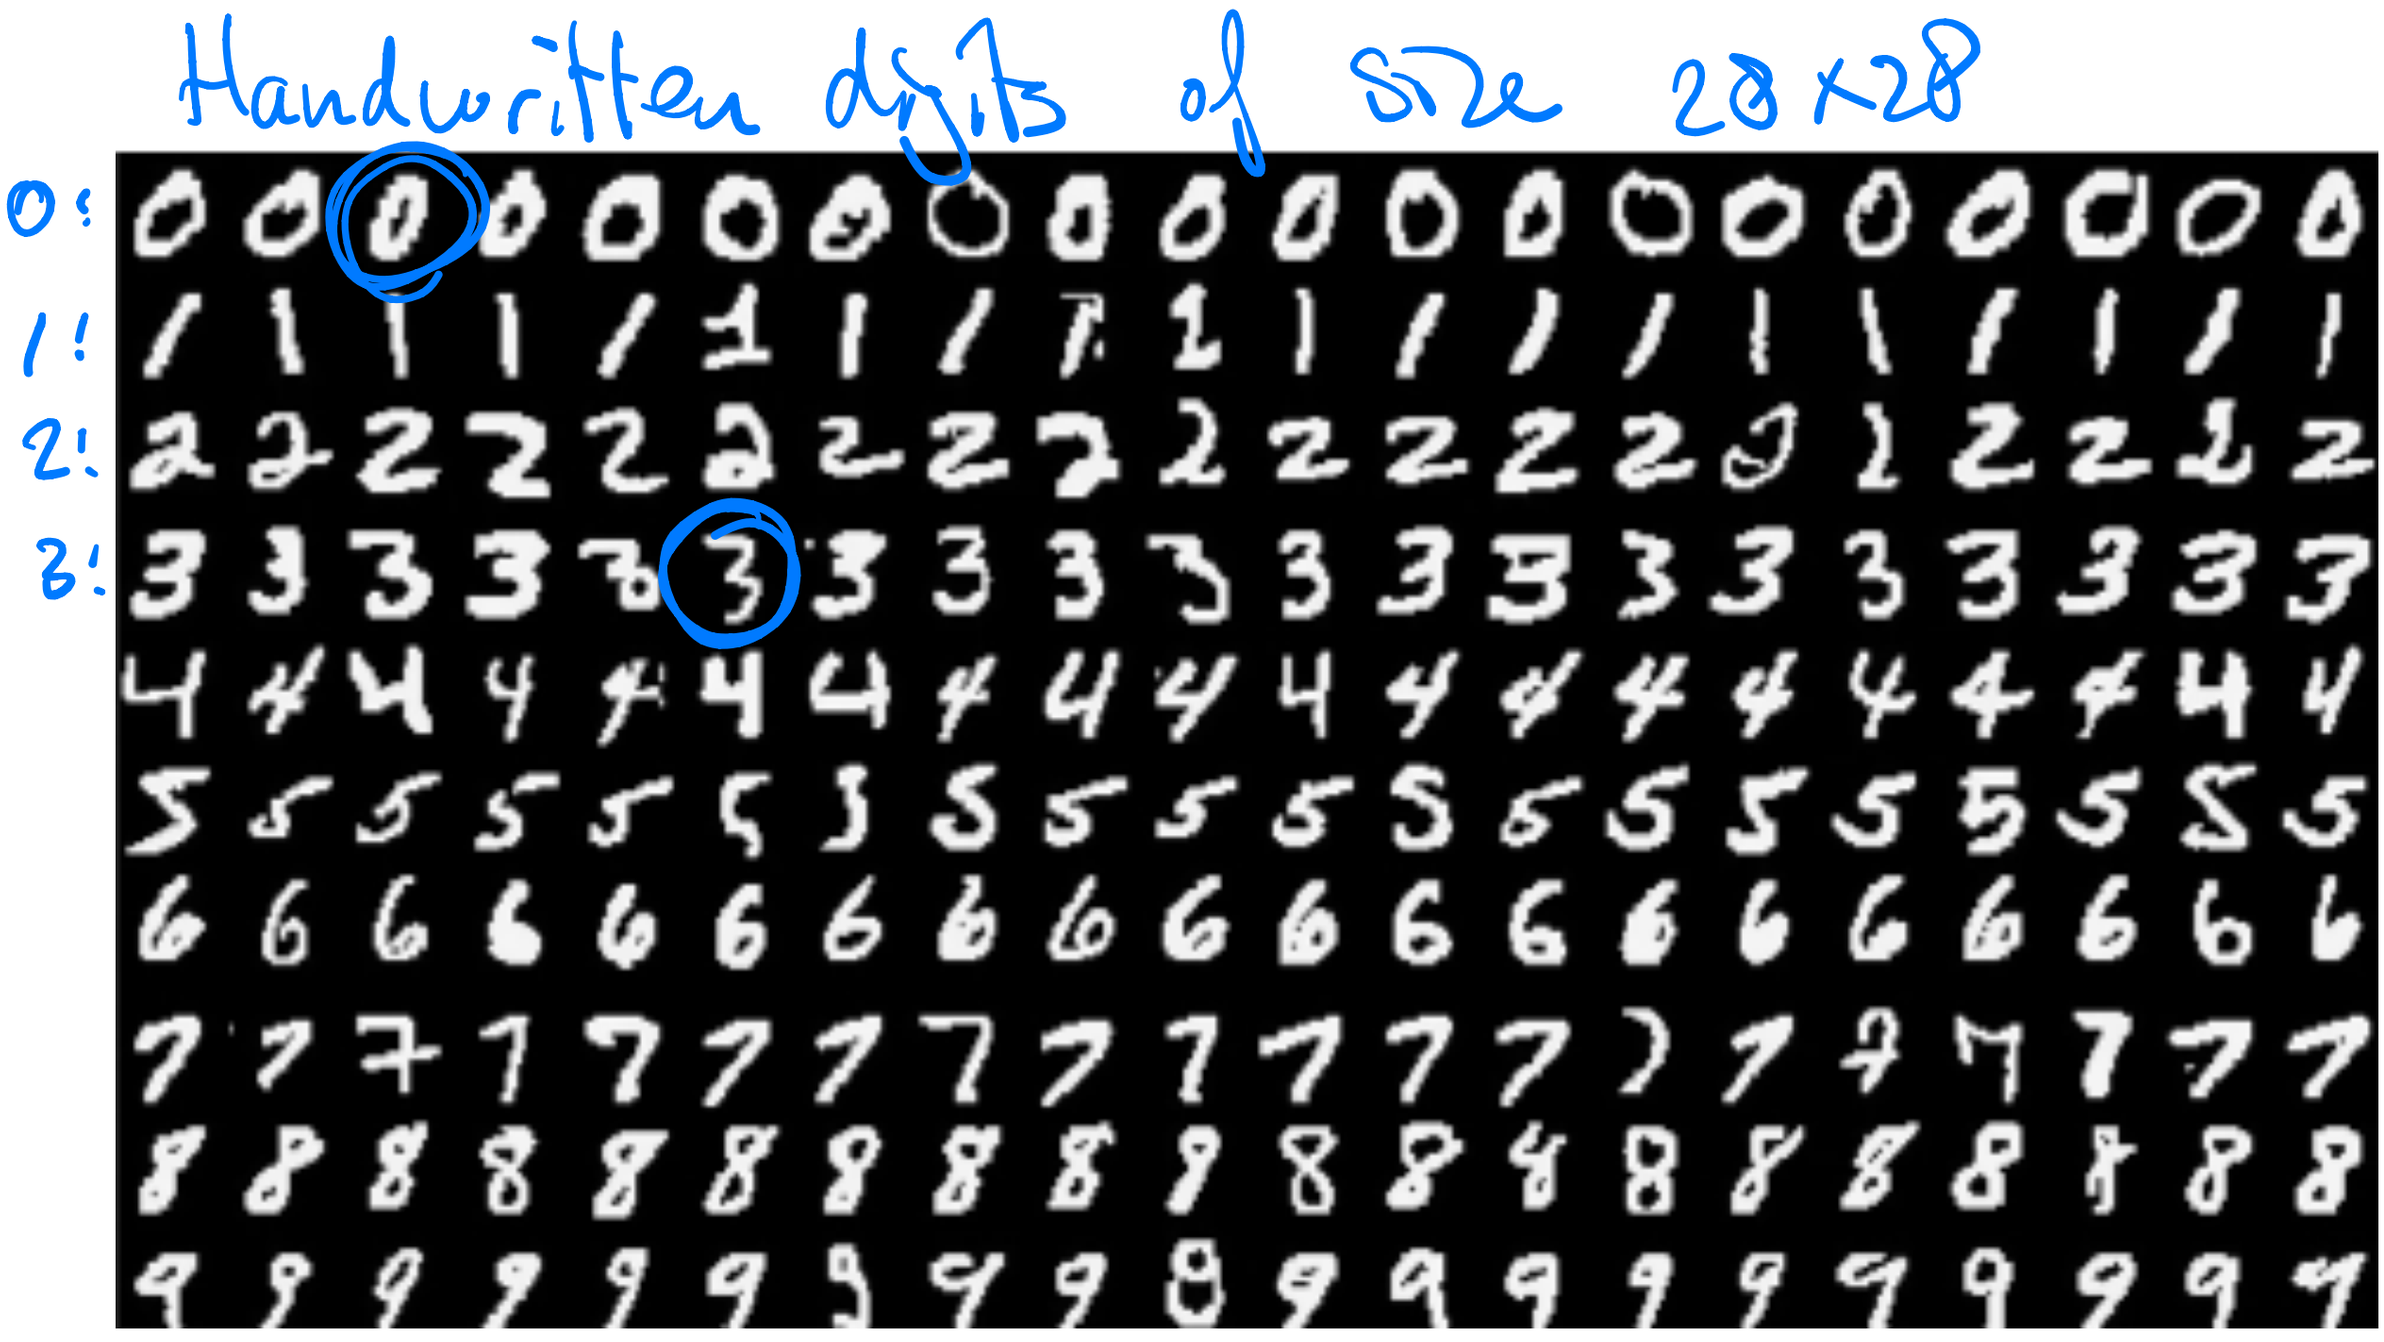

MNIST dataset

E: Experience

heatmap for activity of a gene

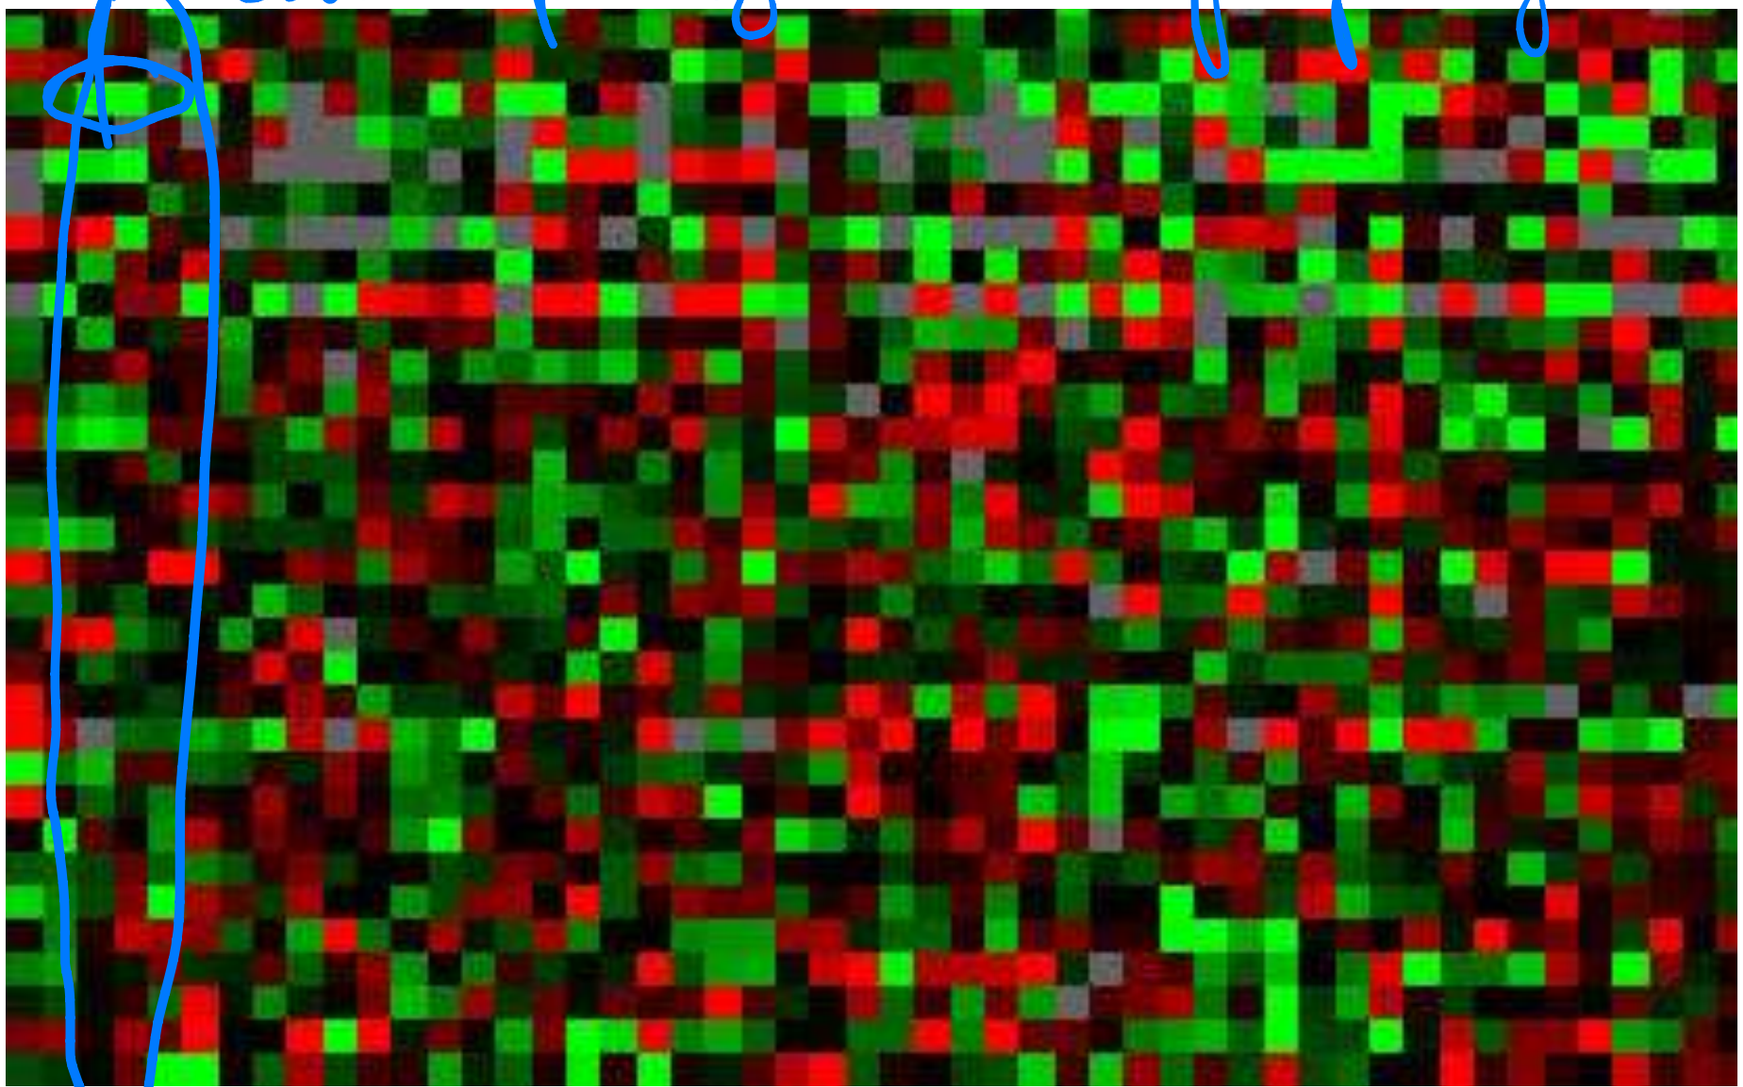

gene

Tumor

Expression matrix of genes (rows) for 64 human tumor samples (columns). [source: ESL 1.3]

E: Experience

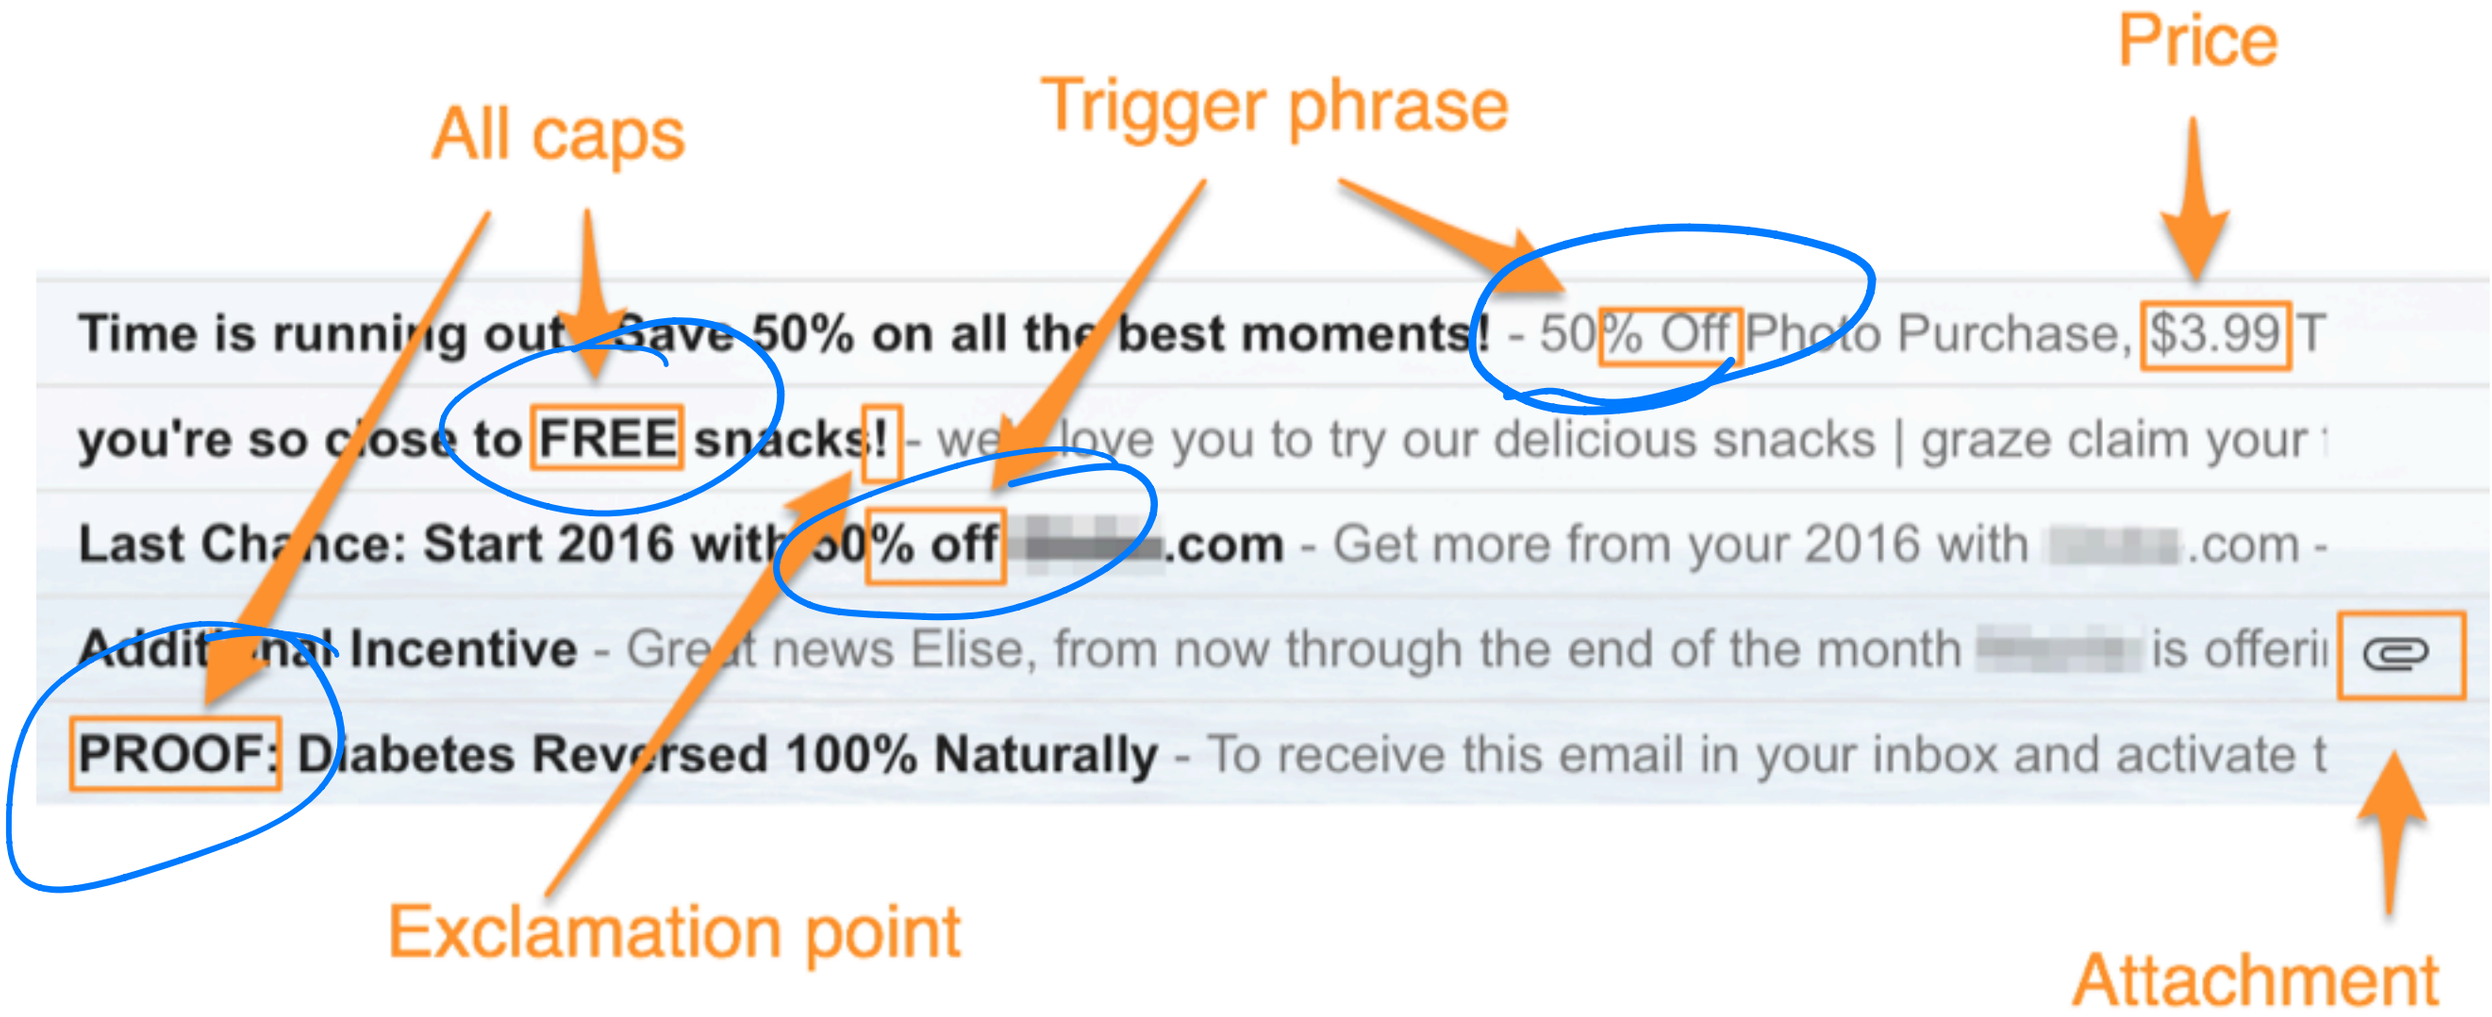

Examples of spam emails. [source: Yesware]

T: Class of tasks

Classification:

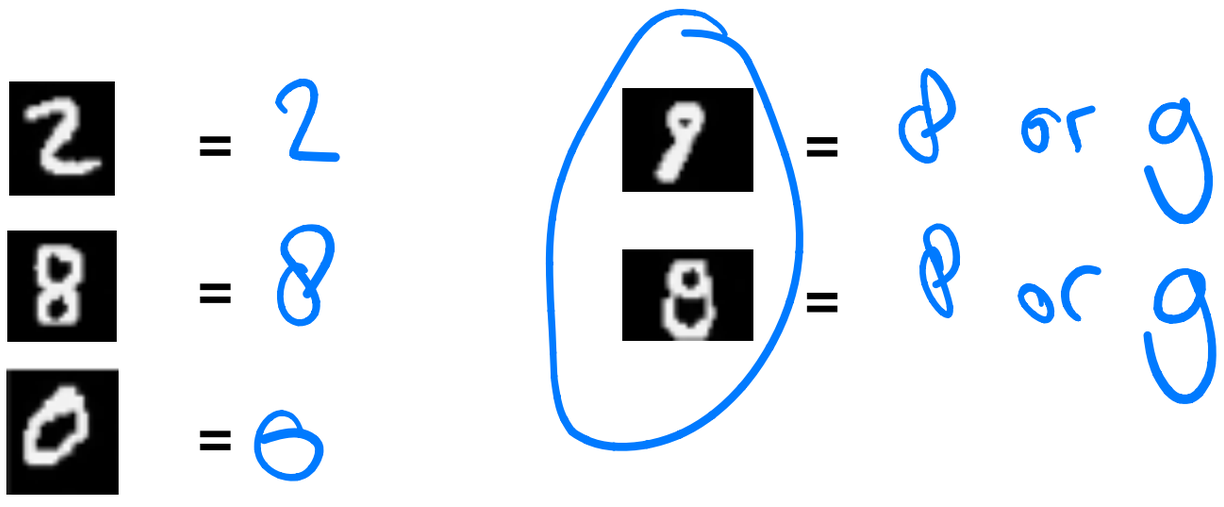

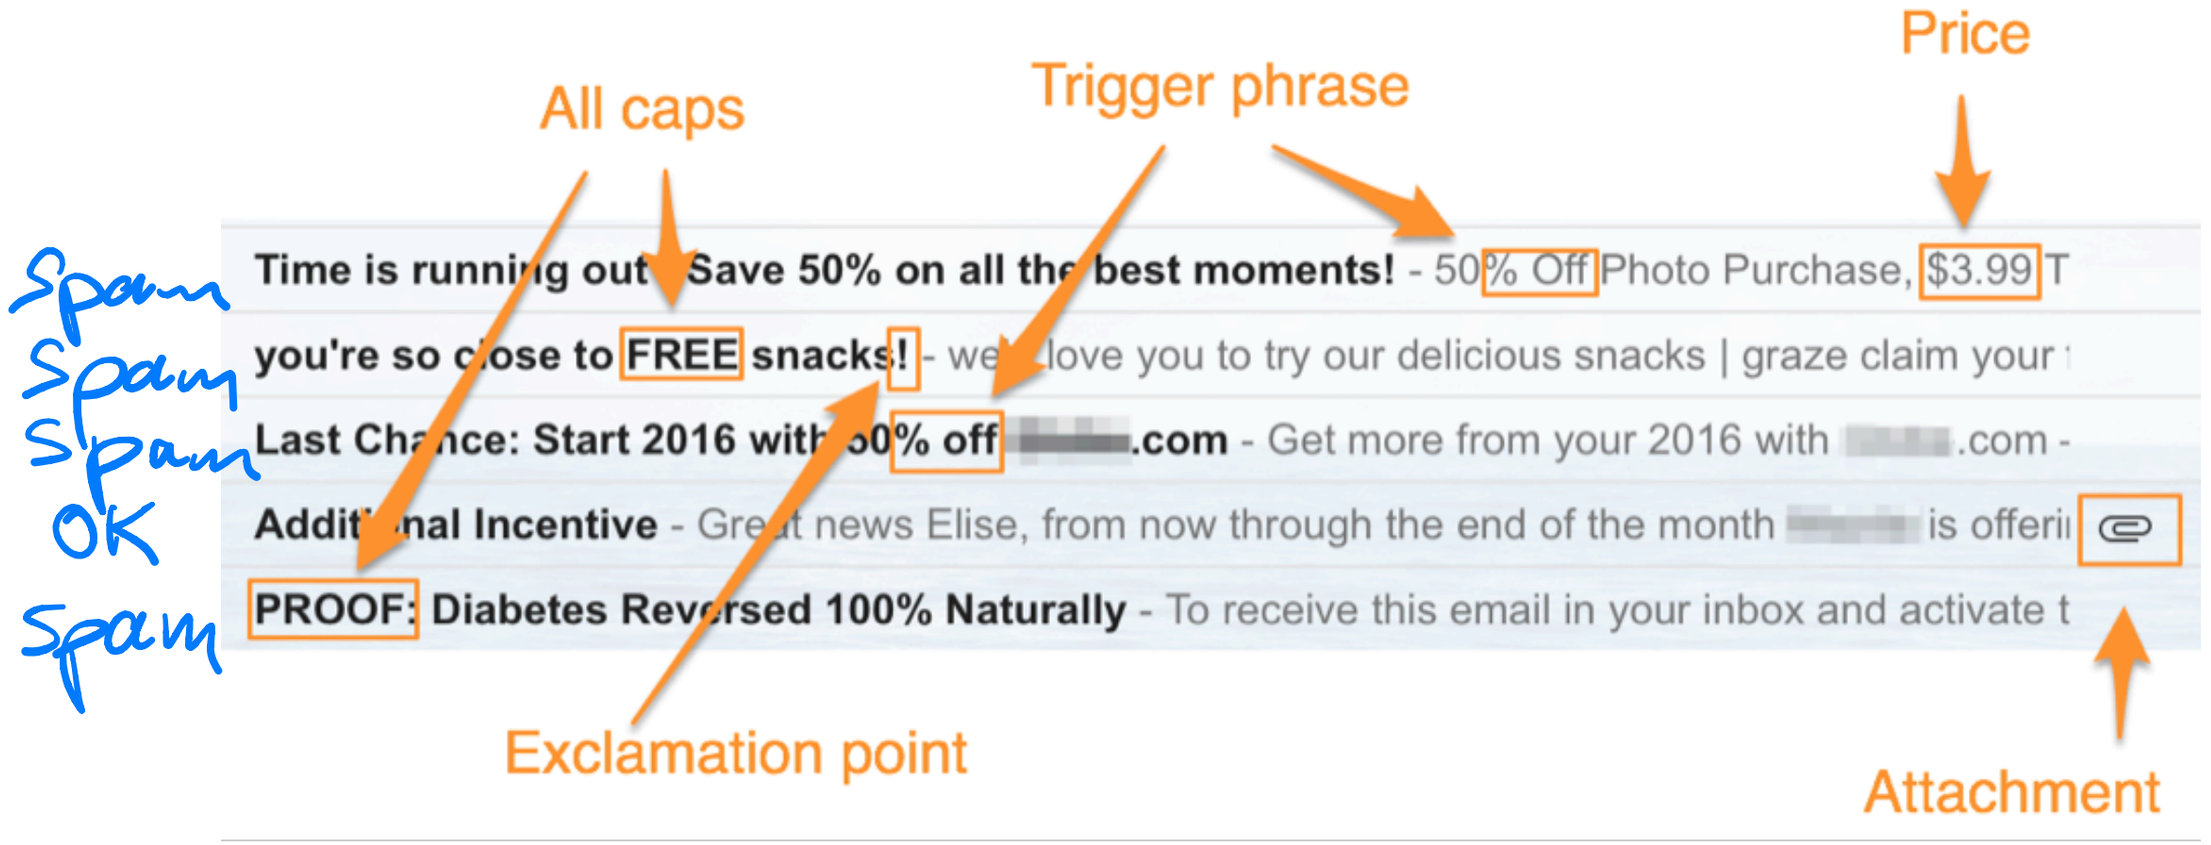

T: Class of tasks

Regression

input : x

target : $t = \sin(2\pi x) + \varepsilon$

noise : $\varepsilon \sim N(0, 1)$

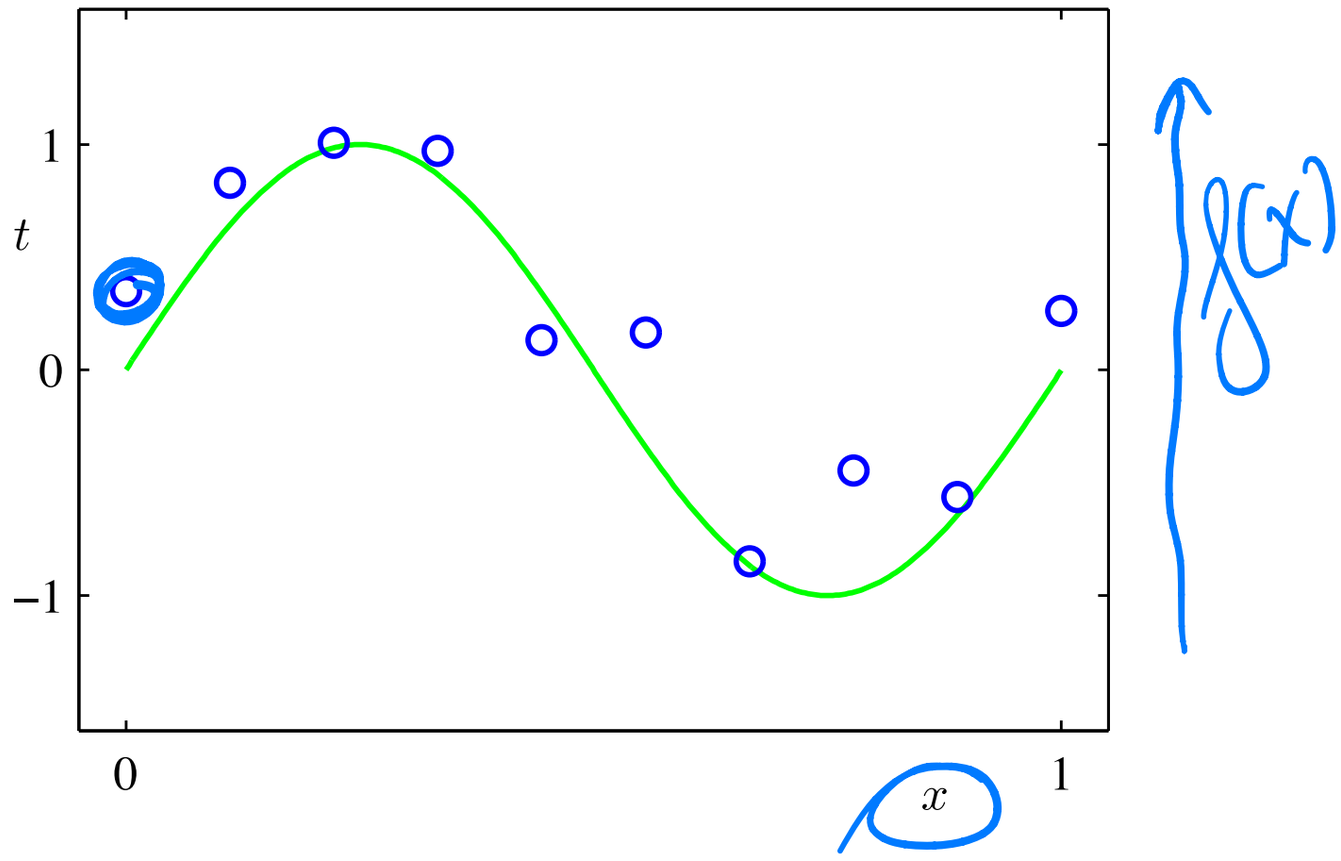

T: Class of tasks

Regression $f(x) = w_0 + w_1 x + w_2 x^2 + \dots + w_M x^M$

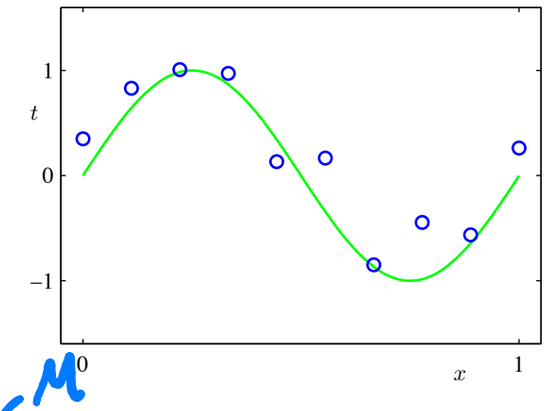

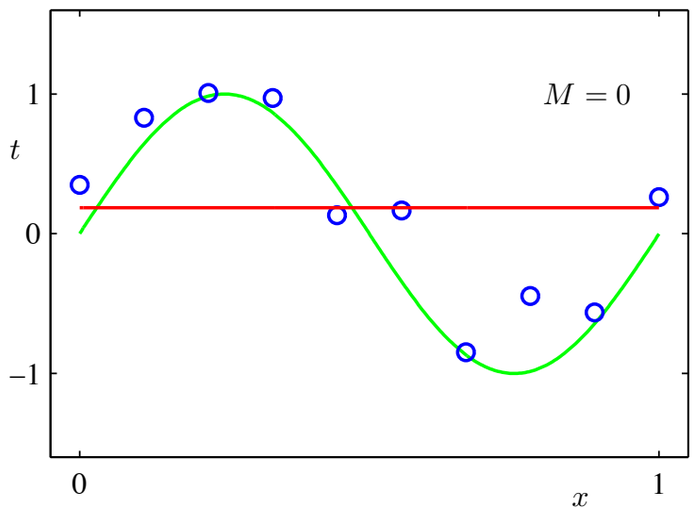

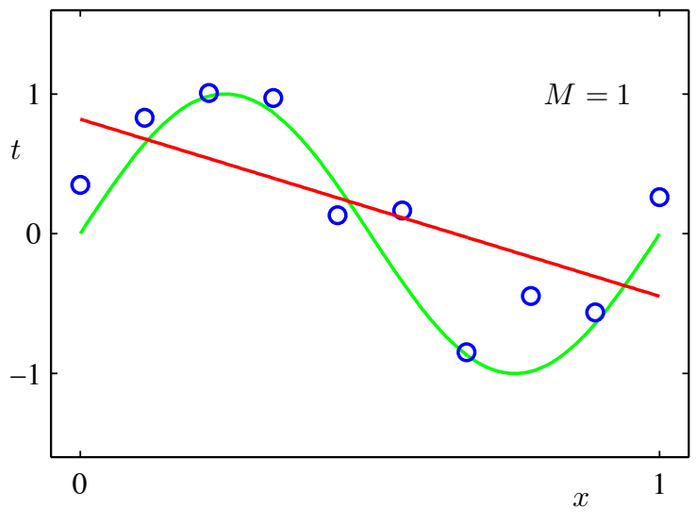

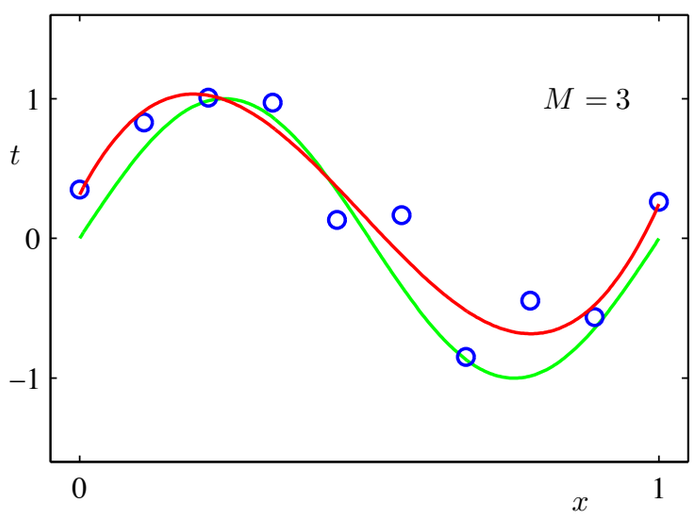

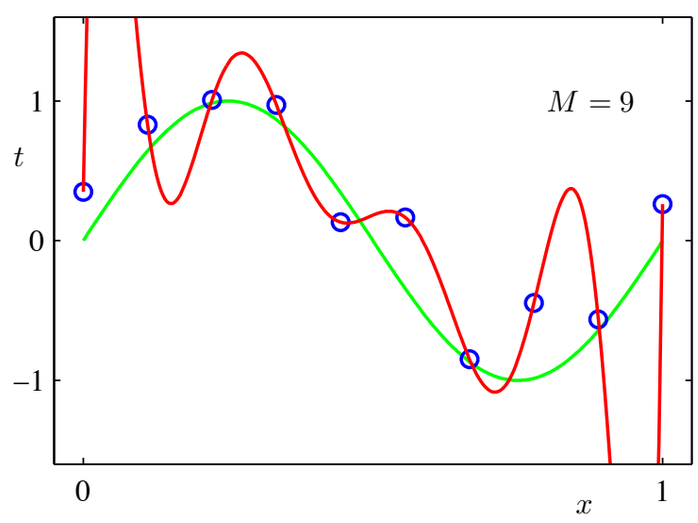

Polynomials of order *M* (red) fit to data constructed as t = sin(2πx) + ε (green)

T: Class of tasks

Clustering

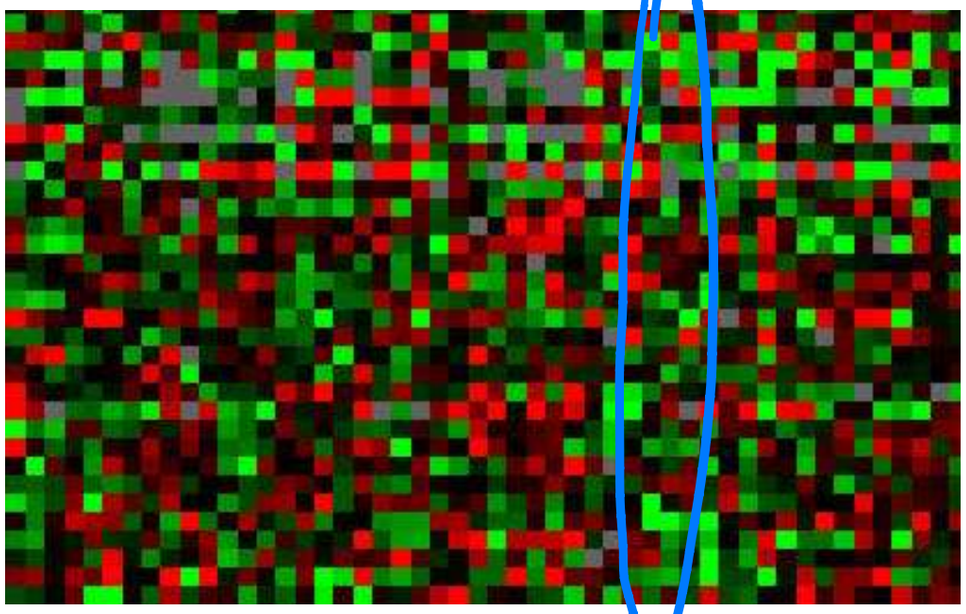

Expression matrix of genes (rows) for 64 human tumor samples (columns). [source: ESL 1.3]

P: Performance measure

Classification

= 2

= 8

= 0

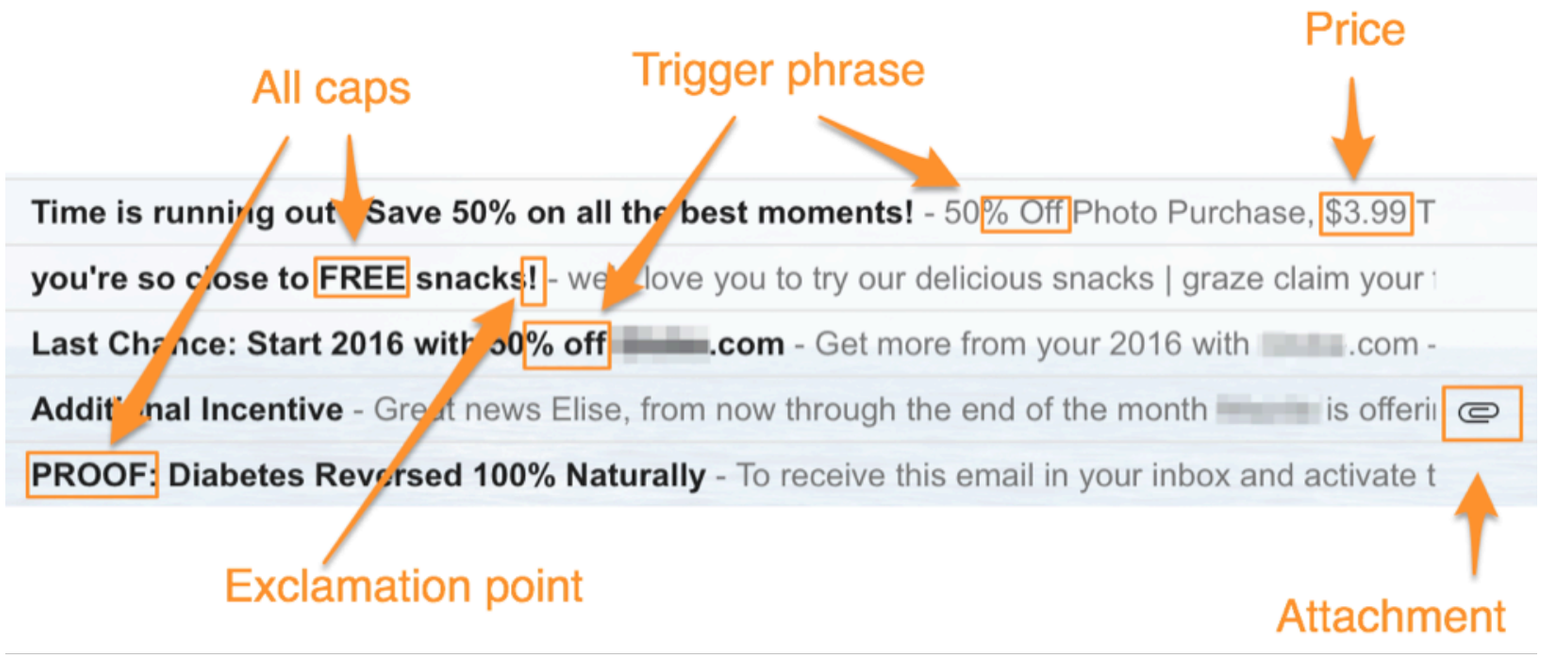

$$\text{accuracy}(y, \hat{y}) = \frac{1}{n_{\text{samples}}} \sum_{i=1}^{n_{\text{samples}}} \mathbb{1}[y_i = \hat{y}_i]$$

indicator funct. = $\begin{cases} 1 & \text{if } y_i = \hat{y}_i \\ 0 & \text{otherwise} \end{cases}$

P: Performance measure

Regression

large MSE

small MSE

MSE = 0

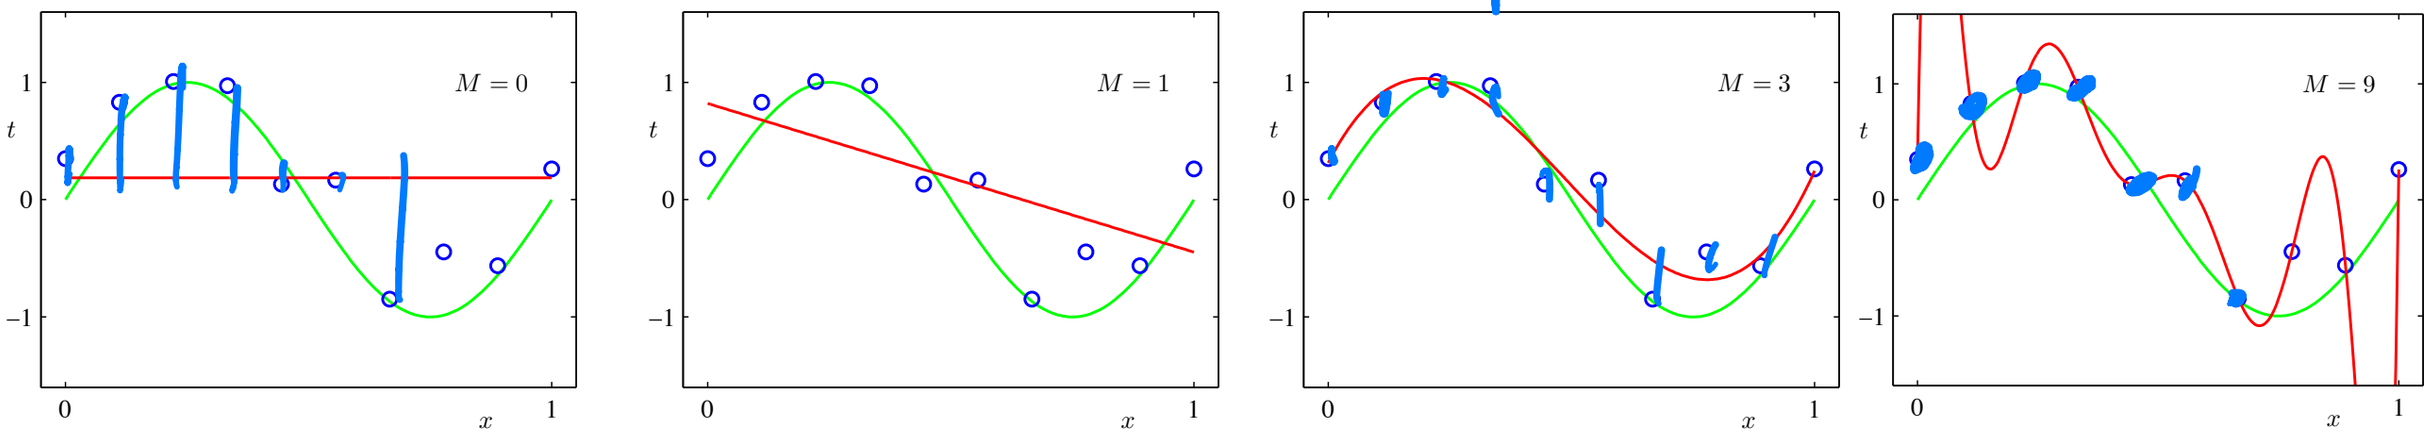

mean squared error

$$\mathrm{MSE}(y, \hat{y}) = \frac{1}{n_{\text{samples}}} \sum_{i=1}^{n_{\text{samples}}} (y_i - \hat{y}_i)^2$$

$\hat{y}_i = f_{\underline{w}}(x_i)$

Polynomials of order *M* (red) fit to data constructed as t = sin(2πx) + ε (green)

P: Performance measure

Clustering

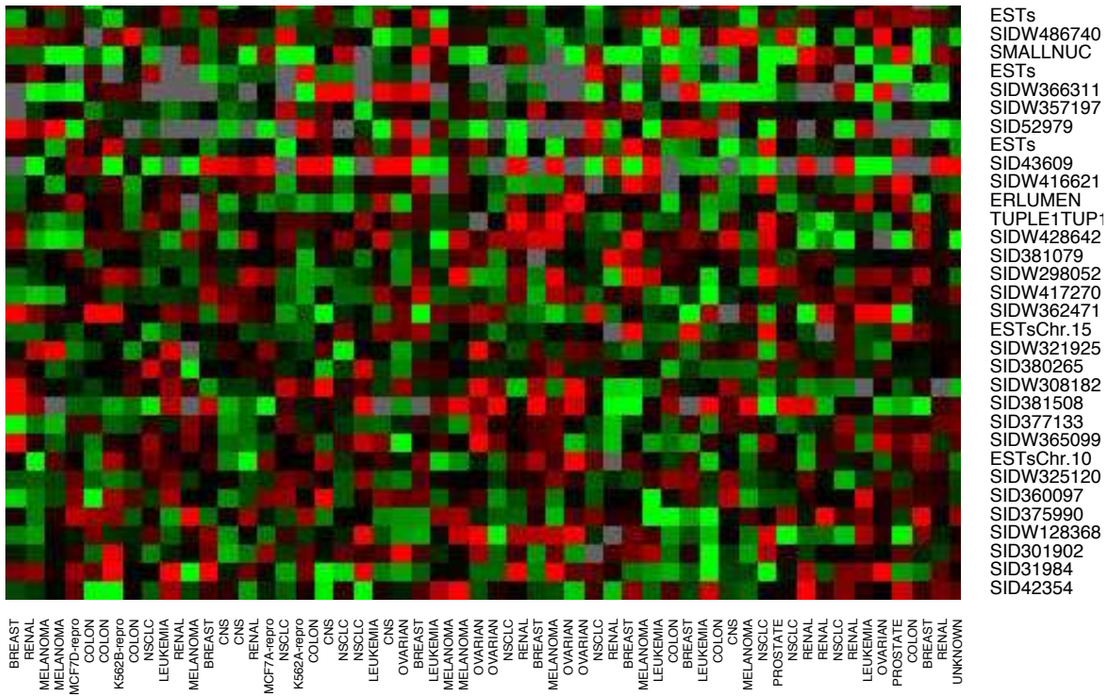

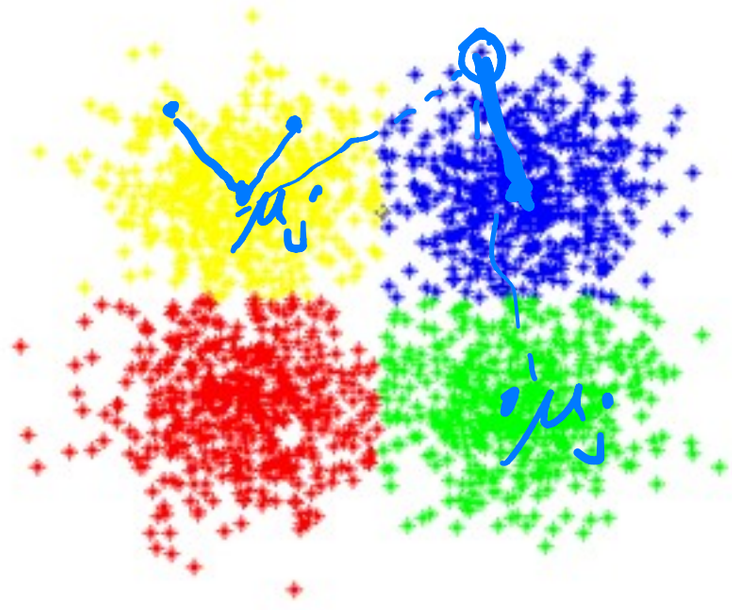

$$\text{within cluster sum of squares} = \sum_{i=1}^{n_{\text{samples}}} \min_{\mu_j \in C} \| \mu_j - x_i \|^2$$

Expression matrix of genes (rows) for 64 human tumor samples (columns). [source: ESL 1.3]

P: Performance measure

$$\mathrm{MSE}(y, \hat{y}) = \frac{1}{n_{\text{samples}}} \sum_{i=1}^{n_{\text{samples}}} (y_i - \hat{y}_i)^2$$

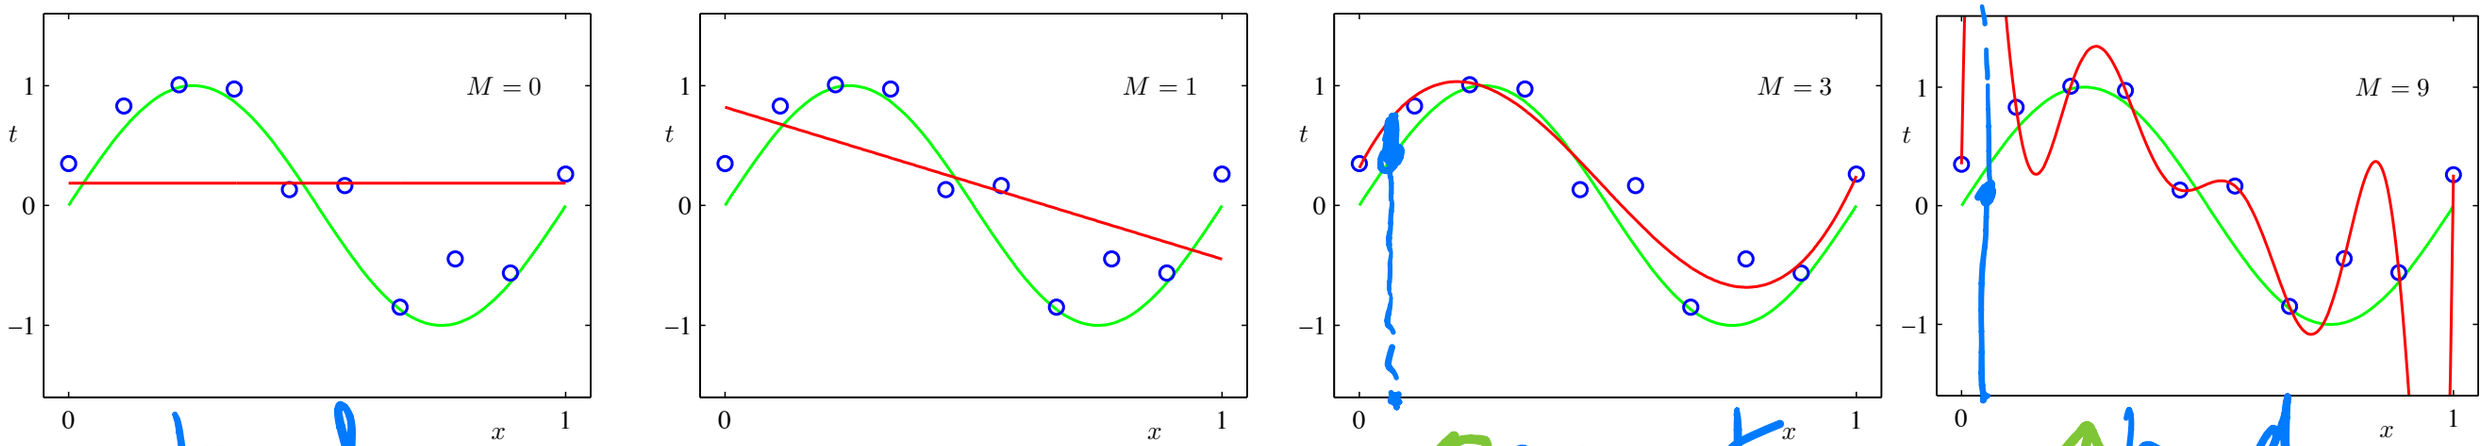

Best performance on training set :

Best performance on new datapoints :

Q: On which datapoints should performance be measured?

Generalisation:

performance should be measured on new data (test data)

What is machine learning?

“A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E.”

- Tom M. Mitchell

Machine Learning, Tom Mitchell, McGraw Hill, 1997